



The Vector Approach to Data Center Power Planning

How to avoid unplanned obsolescence in the power distribution infrastructure

Abstract

Many data centers, including most of those built before 2001, are at risk of outstripping their capacity to power and cool their IT systems. Already, data centers consume 10–30 times more energy per square foot than the typical office building—a figure that has doubled in the last five years. Energy costs represent the single largest component of operating expense, and a potential barrier to future expansion.

Does IT really have a handle on this trend?

More regularly and frequently, organizations are hitting fixed limits in their power systems—even systems that were designed and deployed fairly recently. With the volatile rate of change in IT technologies, power demands can quickly exceed established barriers in a legacy distribution system, such as the performance potential of existing amperage/voltage ratings, UPSs, cabling and connectors. The cost of upgrading, augmenting or replacing the power architecture can be astronomical.

The costs often could have been minimized or avoided if the power planning process had simply been more forward-looking and holistic in the first place. This white paper describes an approach that considers the major milestones and thresholds in data center power requirements—and how planners should adjust their strategies and recommendations for data centers as they pass through different evolutionary stages.

If you don't want to get caught short, read on.

Contents

- Traditional thinking in a transitional world 3**
- Why are traditional power planning approaches falling short?..... 3**
 - Myth #1. “We can keep adding on to what we have with more of the same.” 4
 - Myth #2. “The low-cost solution will save us money.” 4
 - Myth #3. “The power infrastructure should be built to our average watts per square foot.” ... 5
 - Myth #4: “We can always scale the power distribution to match IT growth later.” 5
- The Vector Approach to power distribution planning 6**
 - Type A: The legacy data center 7
 - Type B: The transitional data center..... 8
 - Type C: The next-generation data center..... 10
 - Type D: The next-generation, high-density data center 12
- Build flexibility and visibility into enclosure-level power distribution 14**
- Closing thoughts..... 15**
- For more information..... 15**

Blade servers, virtualization, unified computing... IT systems and the design strategies for deploying them are changing fast. What about the supporting infrastructure—the power distribution system that is required to run it all? Power planning methodologies and assumptions often lag behind the evolutionary curve of technology innovation and data center growth. Consider these recent, real-world examples:

The Reality

The national retail chain built its data center power infrastructure to a five-year-old specification, supplying 5,000 watts of redundant power to each enclosure. Trouble was, IT enclosures now required 10,000 watts of redundant power. The company deployed dual 5,000-watt power distribution on A and B sources for each enclosure, instead of one 10,000-watt infrastructure. This architecture used up twice as many pole positions and maxed out panelboards far below their actual power capacity.

The Missed Opportunity

Had the infrastructure been designed and built with higher amperage connections, the company would have saved \$5,000 per enclosure and reduced cabling and panelboard pole positions by 50 percent—for a total savings of more than \$500,000.

The Reality

The software services company deployed 30-ampere (30A), three-phase distribution to enclosures in its 20,000-square-foot data center—more than enough power for one blade server per enclosure. Trouble was, the company modified its hardware strategy and now wanted *two* blade servers per enclosure. Without major changes in the power infrastructure, they faced overloads and tripped breakers that would bring servers down.

The Missed Opportunity

Had the infrastructure been built with 60A/208V power distribution from the start, the data center would be okay. Instead, everything from enclosure-level power distribution to cabling to panelboard breakers had to be replaced or reconfigured at a cost of nearly \$250,000.

The Reality

The financial services institution designed and deployed a data center power distribution system they believed could support three blade servers per enclosure. By their calculations, 50A, three-phase power to each enclosure was more than enough. Trouble was, they didn't realize that newer dual-corded servers actively draw power from both sources at once. They had sized the power infrastructure based on ammeter readings from one power source, not both. Since each blade server actually draws about 5.2 kVA, the 50A infrastructure (14.4 kVA) wasn't up to the task.

The Missed Opportunity

The company had invested so much money in the 50A infrastructure, they didn't want to rip it out and replace it with 60A power distribution (17 kVA). So they decided to deploy only two blade servers per enclosure. The lower-density arrangement meant 33 percent more enclosures, 33 percent more rack-level power systems, 33 percent more pole positions, 33 percent more raised-floor space and more stress on the HVAC system—adding \$400,000 to the total cost of the data center.



Traditional thinking in a transitional world

All of these nightmare stories have a common theme; the power distribution systems were planned based on old rules, old metrics, old server designs, and old assumptions for power density per U and per enclosure.

With the proliferation of blade servers and virtualization strategies, power consumption keeps rising—up to 600–1000 watts per U and growing. Power consumption in high-performance computing applications may soon reach up to 40 kW per rack. Furthermore, power demand can easily double or triple during peak periods, and it fluctuates with every move, add or change. Adding a 1U or 2U server used to mean drawing 200–300 more watts from the branch circuit; a new blade server consumes *20 times* as much current.

That means the power distribution system is more easily stressed by even the simplest changes in your data center. How much current are your servers drawing right now? Are electrical circuits approaching capacity, ready to trip a breaker if transaction processing rises or a new component is added? Would you be able to see trouble coming?

In a startling number of cases, there isn't much spare capacity for normal evolution of the data center, and there isn't much visibility into the power distribution system the enclosure level. In Data Center Resources experience conducting power audits, we find that the power infrastructure is often an expensive bottleneck to much-needed IT expansion. Approximately 80 percent of data centers are facing serious and insidious problems that could cause unplanned downtime.

It is time to drastically rethink how data center power infrastructures are planned—the metrics and approaches used to create the design, and the best way to look at capital expense versus total cost of ownership.

Why are traditional power planning approaches falling short?

In a word, density. Only two years ago, a typical 2U server required about 370 watts of power. A 2U server purchased today would pull closer to 530 watts. Power demands have escalated 30–50 percent in the same footprint in a very short time. Since it could take 12–18 months to plan and build a new data center, the new power design could actually be obsolete by the time you flip the switch.

When you start talking about blade servers, the picture is magnified. A blade server consumes as much power as a typical home's electric oven. Imagine trying to make your home's wiring accommodate two, three or four ovens in your kitchen. You can't force-fit, you can only upgrade. The electrician's bill will be a shock, compared to what the cost would have been if you had planned ahead for four ovens when building the house.

This is the scene, magnified many times over with data center power designs. Data centers are changing far more rapidly than they used to, yet many are based on power systems that weren't designed for that rate of change. The average turnover of a server is about three years for enterprises, about five to six years for small- to mid-sized businesses (SMBs). That means every few months; a notable share of the data center's hardware is being exchanged for more power-hungry equivalents in the same footprint.

Amid these realities, power planning has often been guided (or rather, *misguided*) by four prevailing myths:

Myth #1. “We can keep adding on to what we have with more of the same.”

Traditional power components don’t match very well with next-generation, high-density IT devices. For example, a traditional electrical distribution panelboard can handle 64,850 watts of power, only enough to support 12 blade servers. There is clearly a major mismatch between traditional electrical components and IT hardware.

Even if you have deployed circuits with receptacles large enough to handle the power requirements of each rack, you may have a restriction at the panelboards main breaker, which then becomes the infrastructure’s weak link.

The power infrastructure that meets minimum specs probably lacks other essential merits, such as visibility. Legacy power distribution systems offer little or no view into power draw and power quality at critical points in the power chain. Even if monitoring data is available, it can be difficult to aggregate that data and get a true picture of power consumption at the detail level (for load balancing) and at the summary level (to understand overall utilization and energy efficiency).

Myth #2. “The low-cost solution will save us money.”

Choosing the least-cost power distribution option at the enclosure level can mean far greater expense upstream—more cabling, more connectors, more breakers and panelboards, more elements to install, monitor and maintain.

The chart below shows several different approaches for delivering 5 kW, 7.5 kW, 10 kW and 15 kW to an enclosure—ranging from low-end 30A/120V single-phase input to high-end 60A/208V three-phase input. Total cost of ownership was calculated based on the cost of the enclosure power distribution unit, cables, allocation of the rack power panel (RPP, assumed to cost \$15,000), incidental labor and implications for airflow/cooling.

The result may be surprising and counterintuitive; the cumulative cost of the “less expensive” low-density option is actually higher. Deploying low-power solutions can end up costing 40–50 percent more than higher-power solutions for the same application.

Total number of IT racks: 50

							Power per rack / roof		
Power	Specification		RPP max circuits	Cable + Installation	RPP per circuit	Cost per ePDU	5 kVA/250 kVA		7.5 kVA
Input Plug	Type	VA					Circuits	Cost	Circuits
L5-30P (1Φ)	30A/120V	2,880	168	\$1,000	\$68	\$529	100	\$159,650	150
L6-30P (1Φ)	30A/208V	5,000	80	\$1,050	\$135	\$579	50	\$88,213	100
L21-20P (3Φ)	20A/208V	5,700	56	\$1,100	\$243	\$579	50	\$96,075	100
L21-30P (3Φ)	30A/208V	8,600	56	\$1,150	\$243	\$779	50	\$108,575	50
IEC309 (1Φ)	60A/208V	10,000	80	\$1,300	\$135	\$999	50	\$121,713	50
IEC309 (3Φ)	60A/208V	17,000	56	\$1,400	\$243	\$1,349	50	\$149,575	50
							10 kVA / 500kVA		15 kVA
							Circuits	Cost	Circuits
L5-30P (1Φ)	30A/120V	2,880	168	\$1,000	\$68	\$529	200	\$330,120	300
L6-30P (1Φ)	30A/208V	5,000	80	\$1,050	\$135	\$579	100	\$187,245	150
L21-20P (3Φ)	20A/208V	5,700	56	\$1,100	\$243	\$579	100	\$202,970	150
L21-30P (3Φ)	30A/208V	8,600	56	\$1,150	\$243	\$779	100	\$227,970	100
IEC309 (1Φ)	60A/208V	10,000	80	\$1,300	\$135	\$999	50	\$132,533	100
IEC309 (3Φ)	60A/208V	17,000	56	\$1,400	\$243	\$1,349	50	\$160,395	50

NOTES: Includes material and labor for 75' cable (NY City)

Table 1. Low-power solutions can actually be the higher-cost strategy.

Myth #3. “The power infrastructure should be built to our average watts per square foot.”

Watts per square foot is no longer the basis for the best floor designs (although it is still a basic measure when discussing the data center as a whole). IT equipment varies so markedly in power consumption that an average figure for the whole data center could be off-base for every enclosure in the space. You could have racks of 120V network equipment consuming 3000 watts, racks of new 2U servers at 5000+ watts, and racks of blade servers consuming 15,000 watts. Do you design power distribution for the average across all racks—say, 10,000 watts—thereby overbuilding for some and starving others?

Even a watts-per-enclosure measure could be meaningless. Is the enclosure full or half-empty? Is IT equipment running at full processing capacity and power appetite? Or is it on standby or <50 percent loaded in an N+1 or 2N redundant configuration? Is it midday or midnight? The same equipment draws very different power for light work or heavy processing. The same configuration, doing the same work, can also yield a very different watts-per-square-foot figure when enclosures are different sizes, as shown in Table 2.

5 kVA load, 30A/208V input with L6-30 connection			
Width (in)	Depth (in)	Area (ft ²)	Watts/ft ²
24	40	6.7	900
24	44	7.3	818
30	48	10.0	600

Table 2. Watts per enclosure square foot can vary widely based on enclosure dimensions.

Ultimately, it is not productive to say, “I have 5000 watts here and 15,000 watts over there, so the power infrastructure should be sized for 10,000 watts,” or even, “A fully loaded rack will usually need 7,500 watts.” To optimize the power infrastructure, you need to tailor the strategy in a more granular and dynamic way—and balance a variety of other present and future variables, not just power capacity.

Myth #4: “We can always scale the power distribution to match IT growth later.”

Sure, but usually at excessive cost. For instance, if the data center was wired with 12-gauge wire (the standard for 20-amp circuits), the evolution to 30-amp circuits would mean a complete rewiring job. If power drops to enclosures have L6-30 input connectors, the matching enclosure power distribution units (PDUs) would become obsolete when the enclosure needs >5 kVA. If 225A panelboards are standard, the data center could be pressed into upgrading to 400A panelboards as blade servers are introduced.

With least-cost, low-power solutions, you also have little or no visibility into power conditions at the enclosure level. Such lack of insight might be acceptable when IT equipment in the enclosure draws only 3–5 kVA, but do you want to be blind about power conditions when the enclosure is drawing 10-20 kVA and up? Tripped circuits would be an ever-present threat to reliability.

Organizations that subscribe to these four popular myths tend to:

- Deploy low-density power solutions at the enclosure level to save money, while inadvertently increasing their costs for upstream power equipment.
- Reach bottlenecks in power distribution as the data center expands the scope of its operations or gradually replaces lower-density, legacy servers with newer, high-density ones.

- Deploy one standard power distribution configuration for all enclosures, which leads to overbuilding the design and incurring unnecessary expense.
- Miss out on capabilities that could significantly improve energy efficiency and reliability, such as remote monitoring at the branch circuit and enclosure levels.

The Vector Approach to power distribution planning

In the real world, data centers undergo a lot of change and evolution—and they can reach critical milestones, transition points where one power distribution model must be supplanted by another. The stories described earlier show how difficult it can be to bridge that transition if power planning was based on traditional approaches. The rate of change in the IT environment is just too great, too exponential, too volatile, to plan based on a near-term horizon and rear-view assumptions.

Power planning must move away from flat metrics (such as average watts per square foot) and a relatively static view of the data center. DCR has created a more dynamic approach to power planning—one that mirrors the transitional nature of evolving data centers. We call it the Vector Approach, because it factors the magnitude and direction of change into planning processes.

The Vector Approach defines four different stages of data center evolution:

- **Type A—the legacy data center**—is typical of small businesses and collocation facilities with a large percentage of older, single-corded IT devices that run on 120V power sources.
- **Type B—the transitional data center**—is typically found in larger organizations that sustain a mix of old and new IT equipment, 120V and 208V, single- and dual-corded.
- **Type C—the next-generation data center**—has a high count of newer, more power-hungry 1U and 2U servers running on 208V power.
- **Type D—the next-generation, high-density data center**—uses blade servers, virtualization and/or unified networks that demand a lot of power per square foot.

The fundamental assumptions that go into power planning will be different at each level:

- For one, some key power distribution components have fixed limits. A panelboard of a certain rating offers only so many circuit breakers and delivers x amount of total power. Power cables, plugs and receptacles of different types can handle only so much current. Enclosure-based power distribution units have widely varying ratings and features.
- Second, features that were optional for a Type A or B data center—capabilities such as load balancing, monitoring and remote management—could be baseline requirements for a more advanced data center.
- The first step therefore is to identify where your data center stands on the evolutionary scale—now, 18 months from now, 24 months from now, and so on—and to modify the planning mentality accordingly.
- Planning must be performed in full context, not limited to standard metrics or rules of thumb. Why? The strategy that provides the best value and performance for a Type A data center would be an expensive bottleneck for an emerging Type B data center—and a disaster for a Type C data center. The strategy that satisfies a Type D data center would be unnecessarily costly for a Type A or B data center. Knowing where



you stand on the scale, wise choices can be made to manage capital outlay today while setting the stage for anticipated evolution later.

Power distribution options for evolving data centers

The number of servers that can be supported by each power drop depends on the power rating of the drop. For example, at the low end, a 15A/120V power drop offers 1.4 kW of available power, enough to support four 1U servers or two 2U servers. On the higher end, a 60A/208V three-phase power drop provides 17.3 kW of power, enough to support 49 1U servers, 27 2U servers or three to four blade servers.

Looking at it a different way, to support a fully populated 42U rack with dual-corded 1U servers, you would need either:

- Ten 20A/208V single-phase feeds/power strips, or
- Six 30A/208V single-phase feeds/power strips, or
- Four 30A/208V three-phase feeds/power strips, or
- Two 60A/208V three-phase feeds/power strips.

With the rise in computing density, three-phase power drops are becoming more common. Blade servers, in particular, are driving the need for 208V, three-phase power drops to the rack. The actual type and number of power drops and power strips for your data center will depend on the type of IT equipment in each rack and planned changes over the next few years. Considerations include:

- Cost and availability of each power drop and rack power strip
- Cable management needs inside the rack
- Cost to add or change power drops to the rack

Let's take a look at some key planning considerations for each type of data center.

Type A: The legacy data center

Typical attributes

- Any size data center that has a relatively slow technology adoption rate
- Enclosures containing a high-percentage of single-corded IT equipment running on 120V power
- IT equipment drawing less than 3000W per enclosure
- Small- to mid-sized business (SMB) that replaces servers on a five- to six-year schedule
- Collocation facilities that sell service based on the power input to each customer's enclosure

Key issues

- Planning the phase-out of legacy servers for newer servers
- Providing the 208V power recommended for dual-corded servers that will soon be added
- Providing the 400W–500W power (3–8.6 kVA per enclosure) typically required for new 2U servers

Power planning considerations

In most legacy data centers, the power distribution system is not ready for the transition to newer, dual-corded servers with their demands for higher wattage, 208V power and redundant power feeds. Power distribution from the panelboard to enclosures is typically 20A/120V (delivering 1920 VA) or 30A/120V (2880 VA), with 12 or 24 outlets in the enclosure.

With this architecture, even a moderate count of new servers could put the data center at risk for overload conditions and tripped circuits. The incremental power required to each rack will require additional panelboard

breaker positions and more cable drops. The infrastructure quickly becomes unwieldy and complex. Enclosure-level PDUs would also be ineffective. Newer servers draw 30-50 percent more power, so the legacy 20A/120V power distribution architecture would be quickly undersized—with only a few outlets useable on each PDU.

The path forward

Upgrade from 20A/120V to 30A/208V power. It might be tempting to choose the lower-cost option of evolving to 20A/208V, but at 3300 VA, that option will provide only modest gains in power capacity and no net gain at the panelboard.

In contrast, 30A/208V power provides 5000 VA (5 kVA) per rack. If you select this option, you should also upgrade to 400A panelboards to make sure you have the capacity at the panelboard for the upsized circuits. You don't want to have half-filled panelboards in the future.

One side benefit of migrating from 120V to 208V distribution is that the servers usually become more efficient at the higher voltage, some by as much as two percent, which reduces day-to-day operating cost.

Use L14-30 connections for power input to enclosures. Unlike the L6-30 inputs commonly used for 30A/208V power, L14-30 connectors support both 120V (for older routers, hubs and other equipment) as well as 208V for servers.

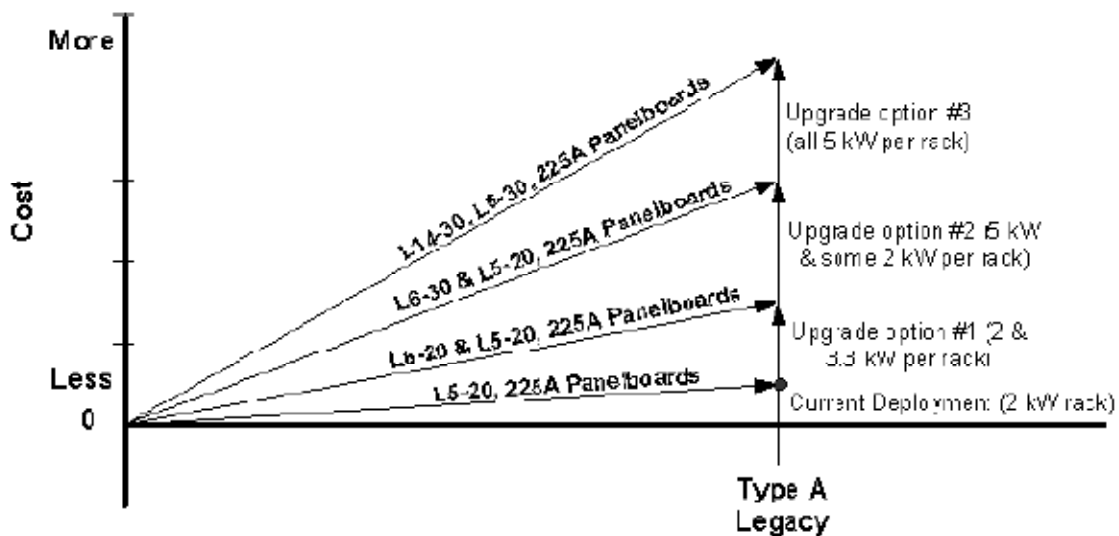


Figure 1. The Vector Approach—Type A data center

Type B: The transitional data center

Typical attributes

- Top 20 percent of SMBs, collocation facilities or a large enterprise data center
- Adopts next-generation servers on a frequent schedule (average three-year server rotation)
- Majority of hardware is current, may have a large network switch or one or more blade servers
- No more than 5–10 percent of IT equipment running only on 120V power
- Some enclosures requiring 3–5 kVA, others at 5–8.6 kVA
- Data center infrastructure established for 30A/208V single-phase power
- May have begun using 30A/208V *three-phase* power to support blade servers or network switches

Key issues

- Managing the transition away from 120V equipment
- Supplying power to enclosures at a faster rate of increase than before
- Determining whether single-phase or three-phase power will be optimal
- Considering the issues of load balancing and isolation, in addition to power consumption

Power planning considerations

At this stage of evolution, the benefits of 208V power have become clear, most notably the ability to support more hardware on a single circuit. But should that be single-phase or three-phase power? The best answers will depend on the mix of IT equipment and anticipated growth.

Smaller loads could be served by single-phase power; higher-density enclosures will generally require three-phase power. For general capacity planning purposes for a typical organization, you can assume a growth rate of 3–4 percent per quarter for the next 18–24 months.

Most data centers in this category use 30A/208V single-phase power at the enclosure. While this infrastructure provides 5 kVA, you can only fill the enclosure to about 50 or 66 percent of its capacity. To support more IT equipment, you would have to run additional circuits, which begins to adversely affect the future migration of the electrical architecture.

There are some limitations even with 30A/208V three-phase power:

- Only 14 circuits—or 14 enclosures—can be connected to a standard 42-pole panelboard. That means some power available at the panelboard level may be stranded.
- The connection to the enclosure can only support one blade server. If future plans call for higher-density computing, consider planning as though the data center is already at the next higher stage of evolution.

The path forward

Assess present and potential power consumption for enclosures.

- If enclosures will require no more than 3–5 kVA, 30A/208V single-phase power (5 kVA) with an L6-30 input connection will be sufficient.
- If some 120V connections are required (such as for routers and hubs), you will need a neutral wire in the power system. An L14-30 input will better utilize the upstream panelboard, compared to the more commonly used L21-20 configuration.
- If loads could expand to 8–9 kVA in the planning horizon, consider 30A/208V three-phase power (delivering 8.6 kVA) and 400A panelboards. You may not need all this power immediately, but you will reap the benefits of adding one more phase/wire and a larger breaker when the time comes. Also, at this point, load balancing and individual circuit monitoring can be just as important as power capacity.

Determine when 120V equipment will be phased out. If 120V equipment will be required, seek to independently isolate it from the majority of hardware. You could support the odd pieces of 120V equipment with a 208Y/120V

three-phase receptacles (L21-30R), but the addition of a neutral wire throughout a facility adds unnecessary cost to the overall build-out budget.




Figure 2. The Vector Approach—Type B data center

C: The next-generation data center

Typical attributes

- Large standalone or corporate data center
- All computer hardware running on 208V sources, no remaining 120V equipment
- Enclosures generally 80–100 percent filled
- No enclosure with more than two blade servers
- Power consumption at 5–10 kVA per enclosure

Key issues

- Need for greater visibility into power conditions, especially potential overload conditions
- Supplying the needed power capacity without consuming excessive panelboard pole positions
- Dealing with the inrush current generated when rebooting high-density IT equipment

Power planning considerations



When a single enclosure starts to draw up to 10 kVA of power, each enclosure's PDU represents a greater portion of the panelboard's total power. (A 225A panelboard delivers approximately 65 kVA total; a 400A panelboard delivers 130 kVA.) A fault at one enclosure now has a potentially greater impact on the critical mission. As a result, power monitoring becomes more important than ever.

As each enclosure now draws more power, planners must be more aware of how well panelboards are used. You don't want to strand power capacity at the panelboard because you have maxed out pole positions well below the capacity of the panelboard. Likewise, you need to plan for how much power each circuit could draw, so you don't leave the panelboard half full.

The path forward

Use a 60A/208V single-phase configuration to deliver power to enclosures. Since there is no 120V equipment in the data center, the loads can be efficiently supported by a higher density single-phase solution. Unlike the delta connection of a 30A, three-phase solution, the 60A single-phase configuration delivers approximately 15 percent more power, and does it more productively, on only one double-pole breaker.

Choose enclosure PDUs with onboard ammeters for load balancing. You need to know which breaker-protected segment of a PDU should be used to plug in a piece of IT equipment, to balance the load across the unit. An imbalance at the enclosure level can cause negative effects all the way up to the UPS. The PDU should also have 20A double-pole breakers for load isolation and reliability, and remote access to enable proactive preventive maintenance. (A 60A/208V PDU can support two blade servers.)

Consider switchable outlets on the enclosure PDU, because of the potentially harmful effect of inrush current. When powering up, IT equipment temporarily draws a large inrush current that can last for 2–10 ms and be as much as 10–60 times the normal operating current. If the reboot was triggered by a power outage, the IT equipment draws extra current to recharge internal capacitors. As loads in an enclosure approach 10 kVA, these conditions could trip a circuit or cause undesirable domino effects in the power chain.

Install 400A panelboard (130 kVA). With 60A/208V single-phase power distribution, you can physically connect 20 enclosures to a 42 pole panelboard—21 on column-style panelboards. (If you deployed a 30A/208V three-phase solution, you would only be able to connect 14 enclosures.)

However, you couldn't realistically max out all 20 enclosures. Consider that 20 enclosure PDUs at 60A/208V, 80 percent derated, adds up to 200 kVA of power consumption, while a fully rated 400A panelboard provides 130 kVA. This limitation clarifies the importance of monitoring power to prevent overload conditions.

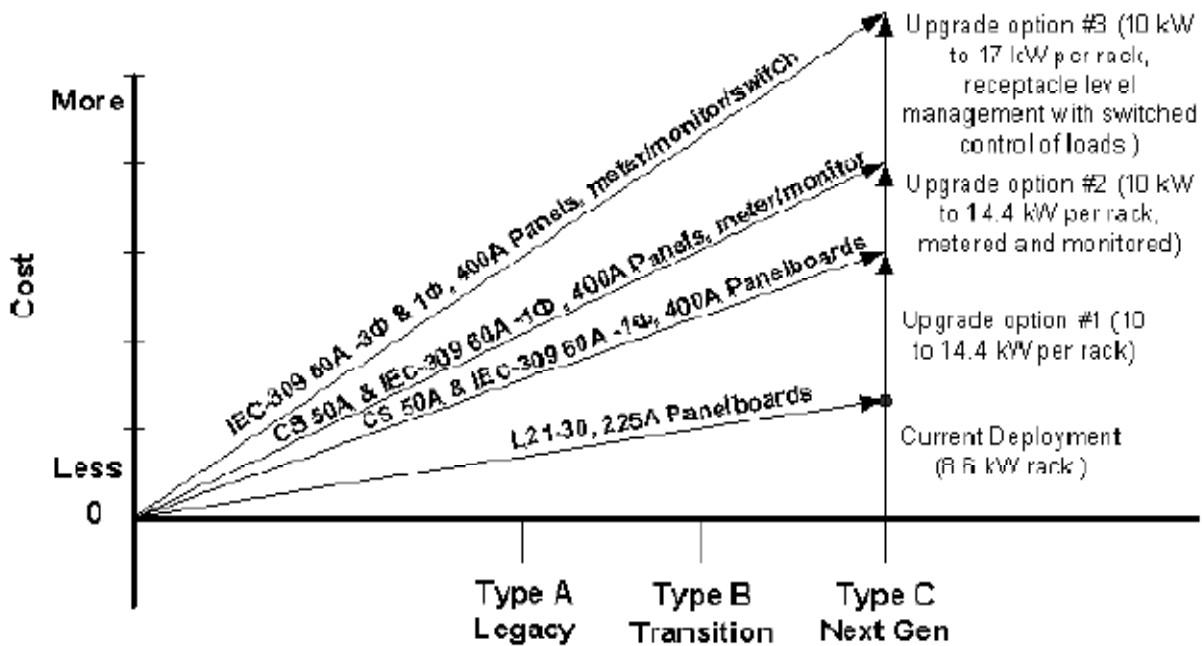


Figure 3. The Vector Approach—Type C data center

Type D: The next-generation, high-density data center

Typical attributes

- High-volume data centers with heavy use of blade servers or virtualization
- High percentage of very high-density systems
- Common to have 10–17 kVA in a single enclosure
- Some enclosures soon to consume as much as 25–30 kVA
- “Greenfield” new-build sites designed by electrical engineering consultants

Key issues

- Very high cost of electrical infrastructure components
- Significant challenges from both a thermal and power perspective
- Very high criticality to each enclosure

Power planning considerations

Organizations in this category reap the benefits of virtualization, consolidation and other strategies that reduce the number of servers, but there is a trade-off. Power consumption and heat dissipation at the enclosure level are exponentially higher than with lower density data centers. The connectivity to the upstream power architecture is more complex. Professional consultants are inevitably required to plan the power infrastructure for such sites.

The path forward

- **Deploy ultra-high-density, three-phase enclosure PDUs.** For example, 60A models are available that provide up to 17 kVA to an enclosure. A 125A model can deliver 33 kVA.
- **Hardwire enclosure PDUs instead of using plugs and receptacles.** This approach reduces connections, improves reliability and eliminates the cost of plugs and receptacles under the raised floor. The cost to deploy an enclosure can be reduced by as much as 20–40 percent. However, racks with PDUs need to be in place before the electrical distribution system can be installed, and moving racks in the future will require an electrician.
- **Deploy highest density loads directly from the distribution panelboard,** due to the high power requirement. This means individual circuit runs from the panelboard to the high-power device, eliminating another pluggable power strip between source and load. The more simple the architecture, the fewer potential points of failure.
- **Consider using distributed power systems** or high-density overhead or under-floor electrical busway systems in lieu of panelboards. These systems allow higher-current circuits (600A to 4000A) to be fed directly to the data center. Replaceable receptacle boxes can be placed close to the rack to power the loads.
- **Investigate all possible airflow management practices,** such as hot or cold aisle containment systems, chimney cabinet configurations, sealed cable entries, blanking panels and air isolation curtains to more effectively manage hot/cool air.
- **Consider scaling back the density of the deployment** by limiting the number of high-power servers in each rack or spreading the load out over more floor space. The high cost of powering and cooling very high power density loads could outweigh the benefits.




Figure 4. The Vector Approach—Type D data center

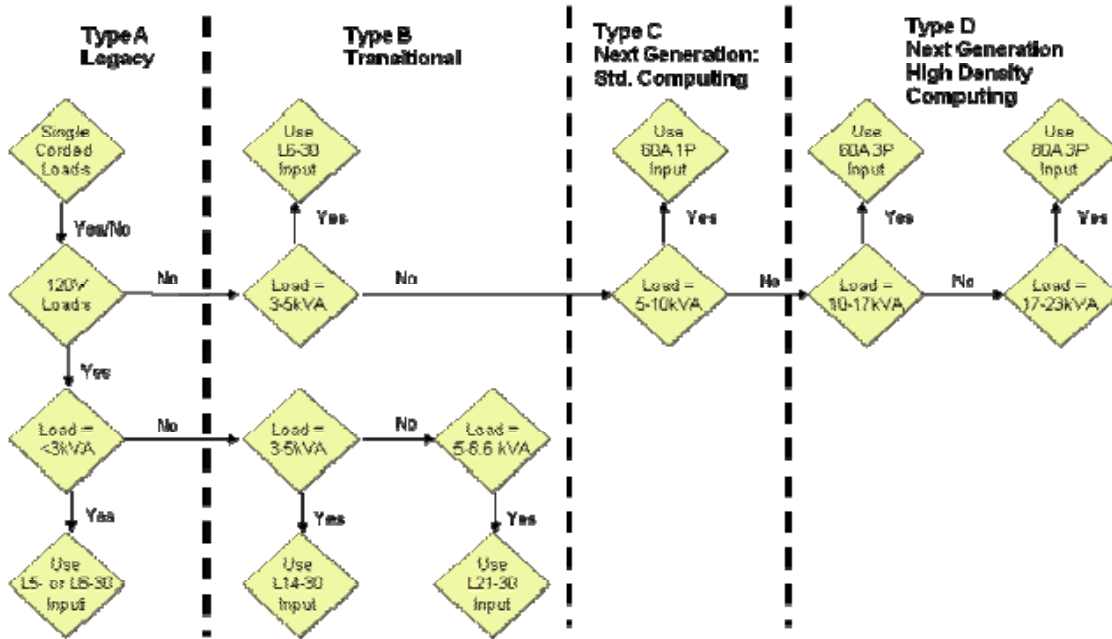


Figure 5. Typical flow analysis for transitioning between data center types.

Build flexibility and visibility into enclosure-level power distribution

How and where power components are implemented in your data center dictates how flexible and scalable the power chain will be, especially as the data center changes and grows. This is true no matter where your organization is on the evolutionary scale.

The good news is that there are more options than ever to tailor the power system for your unique data center requirements—and for the velocity of change. As you plan to upgrade power systems or build a new facility, you need a power infrastructure that is as adaptable as the IT infrastructure must be.

Data Center Resources provides a complete suite of power distribution products to help IT managers meet escalating power requirements:

- **Rack PDU's** distribute from 4–36 kW of power in high-density rack environments—or anywhere power must be distributed to multiple pieces of equipment. We are unique in the industry for providing tiered choices of PDU products along two dimensions—tiered both in power capacity and in functionality. Choose the combination of features and power rating you need to best suit each application.
- **Enclosure Distribution Modules** provide up to 36 kW of plug-and-play primary power distribution from a three-phase UPS or utility source to secondary power distribution devices or directly to IT equipment.
- **Power monitoring and management** software delivers the detailed and aggregated information needed to balance loads, reduce energy costs, prevent tripped circuits and proactively plan for change.



- DCR provides one-stop shopping for ancillary products and accessories, such as **power cables** for most any application.

Closing thoughts

No CIO or facilities manager wants to be the star player in one of the stories described earlier. The risk is real. Even newer data centers built on prevailing best practices have their Achilles' heels. The power distribution infrastructure could have some hidden pitfalls and roadblocks that will be expensive bottlenecks to future growth.

The challenges are intensifying as high-density computing and virtualization become more prevalent. In these environments, IT applications, their processing loads—and the power resources they require— can shift at will, on a moment by moment basis, stressing the power system in ways never before imagined.

Since there are finite limits to the capacity of many power components—and electrical architectures must adhere to the National Electrical Code—it is essential to plan a power infrastructure with a big-picture and longer-term perspective, one that will not require significant upgrades or wholesale replacement later.

It is time to redefine the way power distribution systems are designed, to account for both the magnitude and *direction* of the data center—a Vector Approach.

For more information

To learn more about how to optimize power distribution for your data center in transition, contact Data Center Resources at <http://www.datacenterresources.com> or 1-866-740-2121.