Power to the adaptable data center

Business Unit Manager, Data Center Solutions, Eaton's electrical group

IT managers have more options than ever to tailor the power infrastructure to their unique data center and business requirements.

Every data center is unique, reflecting the needs of the business. How large is the data center? Is it small (<20 racks), medium (20 to 100 racks) or large (>100 racks)? What is the power density per rack? Is it low (1-3 kW), moderate (3-10 kW) or high (>10 kW)? What level of reliability is required? Tier I (basic), Tier II (with redundancy), Tier III (concurrently maintainable) or Tier IV (fault-tolerant)?

Whatever the answers, everything could change tomorrow. Data centers change continually to serve shifting priorities, applications and transaction volumes –and to reflect new technologies, regulations and market forces. Virtually every large data center undergoes moves, additions and changes (MACs) once a month or more, according to a Network Computing survey of 474 IT managers (Figure 1). Frequent change affects 85 percent of mid-sized data centers and 67 percent of small data centers, according to the report. Is your power infrastructure up to these realities? 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.

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.

There has been some talk lately about establishing an industry standard for power system design. But when you consider the tremendous variation among organizations and data center architectures, there really is no one right way to optimize power management for everybody. There are many considerations to bear in mind, and there are technology options to create your best power system.

What type of power drops should feed the data center?

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 120V 15A power drop offers 1.4 kW of available power, enough to support four 1U servers, or two 2U servers. On the high end, a 208V three-phase, 60A power drop provides 17.3 kW of power, enough to support 49 1U servers, 31 2U servers or four blade servers (see Table 1).

Table 1. Power Requirements for Servers

Power drop	Available power (kW)	1U servers supported	2U servers supported	Blade chassis supported
120V 15A	1.4	4	2	–
120V 20A	1.9	5	3	–
208V 20A	3.3	10	6	–
208V 30A	5.0	14	9	1
208V 3ph 30A	8.6	25	16	2
208V 3ph 60A	17.3	49	31	4

Source: Eaton

Figure 1

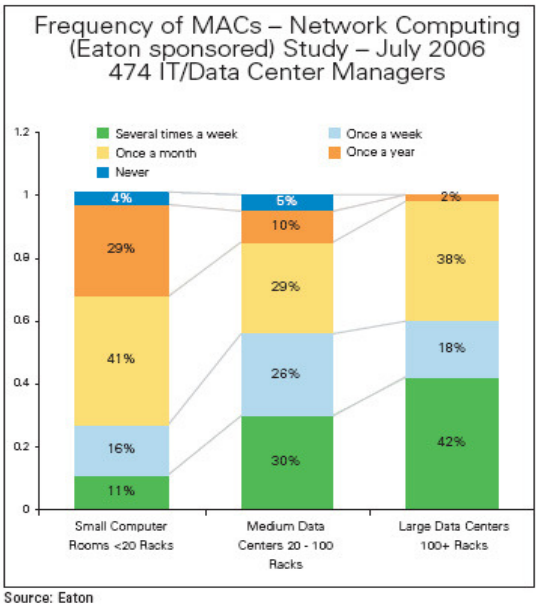

Source: Eaton

Looking at it a different way, to support a fully populated 42U rack with dual-corded 1U servers, you would need (see Table 2):

- Ten 208V 20A feeds/power strips, or
- Six 208V 30A feeds/power strips, or
- Four 208V 3ph 30A feeds/power strips, or
- Two 208V 3ph 60A 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 any 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.

Table 2. Number of power drops / rack power strips for a fully populated, 42U rack with servers

Drop / Power Strip Type	Available Power per Drop / Power Strip	1U	2U	Blade chassis
120V 15A	1.4 kW	11 + 11	9 + 9	–
120V 20 A	1.9 kW	8 + 8	7 + 7	–
208V 20A	3.3 kW	5 + 5	4 + 4	8 + 8
208V 30A	5.0 kW	3 + 3	3 + 3	5 + 5
208V 3ph 30A	8.6 kW	2 + 2	2 + 2	3 + 3
208V 3ph 60A	17.3 kW	1 + 1	1 + 1	2 + 2
208V 3ph 125A	36.0 kW	1 + 1	1 + 1	1 + 1
Servers supported		42	21	6 systems

Source: Eaton

How should power be distributed to racks?

The conventional approach brings power in from the utility or an uninterruptible power system (UPS) to a transformer that "steps down" power to the desired voltage. Power then goes through the main breaker to a panel board, then to power strips in racks. The complexity of this arrangement – particularly multiple connections from panel board to power strips – makes it expensive, difficult to install, hard to monitor and prone to failure. Power distribution units (PDUs) provide a more reliable and streamlined solution (see Figure 3). A PDU eliminates the need for separate transformer, main breaker and panel boards. You have one factory-tested, pre-packaged solution, with options for integrated metering and monitoring.

Figure 2 Power distribution inside IT equipment

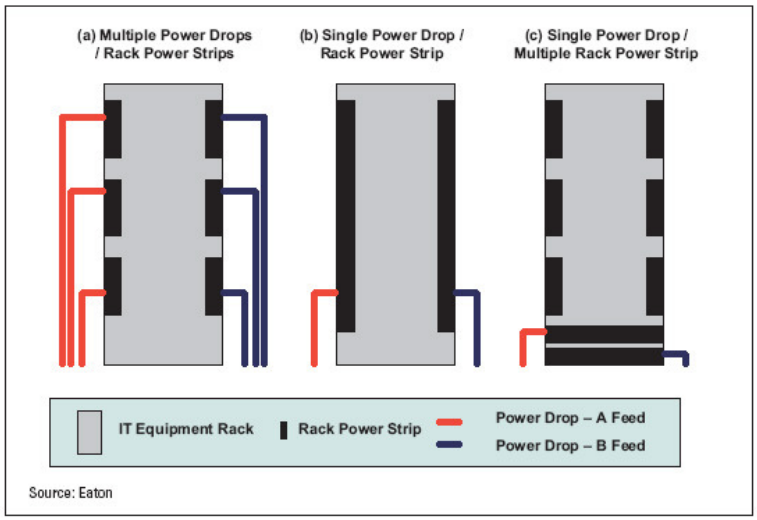

Source: Eaton

The following distribution options are typical:

- Directly from PDU distribution panels to rack power strips, using ULlisted, factory-tested power cables or conduits with branch distribution breakers. This scenario is suitable for both raised and non-raised floor applications.
- From the PDU with subfeed breakers that feed a remote power panel (RPP). This option, suitable for raised and non-raised floor data centers, is highly flexible, reduces cabling, and offers faster and easier installation.

How will you monitor power consumption and quality?

Is your IT equipment receiving computer-grade power? If you add a new piece of equipment, will you overload that branch circuit? Which rack has enough power to accommodate a new piece of equipment? How can you optimize the available power distribution in your data center without more capital investment? These questions are answered with 7x24 power quality metering, which can be conducted at several points in the power distribution chain (see Figure 4).

At the PDU or RPP level, a meter can monitor hundreds of branch circuits feeding tens of racks, with one IP address. At the subdistribution level, a meter can monitor tens of power strips in a few racks with one IP address. At the rack power strip, a meter can monitor a few branch circuits with one IP address. The right strategy for your needs will be a trade-off between the number of IP addresses you're willing to allocate, the number of devices you want to monitor, and degree of detail required.

Figure 4. Real time power distribution monitoring options

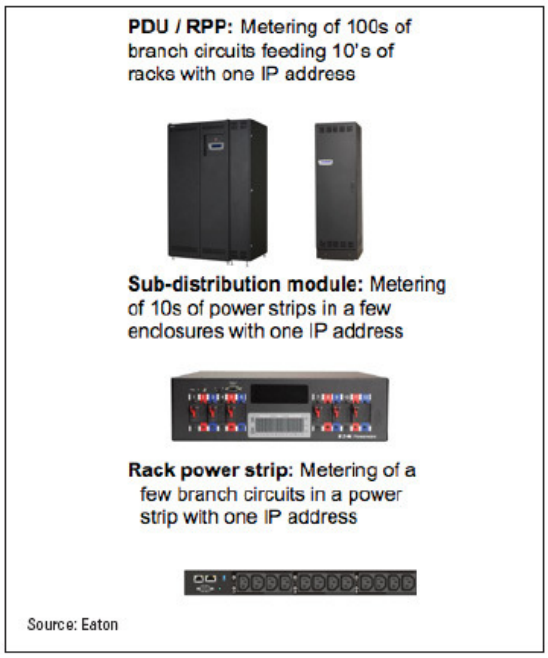

How will you provide clean, reliable backup power?

The UPS provides a critical line of defense against power outages, transients and anomalies that could otherwise compromise data center availability. Several deployment approaches are available (see Figure 5):

- Centralized, where a large-capacity UPS powers the entire data center. This works well when growth can be accurately forecast. The basic architecture simplifies monitoring and maintenance, but a failure of the UPS can bring down the entire computer room.
- Zoned, where the computer room is divided into zones, each powered by a UPS. This approach is more scalable, and a UPS failure would only affect a single zone, not the whole data center.
- Distributed, where a UPS serves one or a few racks. This approach adds complexity but limits the impact of a UPS failure.

In practice, only four to eight UPSs can be connected in parallel for capacity or redundancy in any deployment choice, so it's important to size the UPS appropriately both for present and anticipated load requirements, for three to five years out.

Figure 5. UPSs can be deployed in centralized, zoned or distributed arrangements

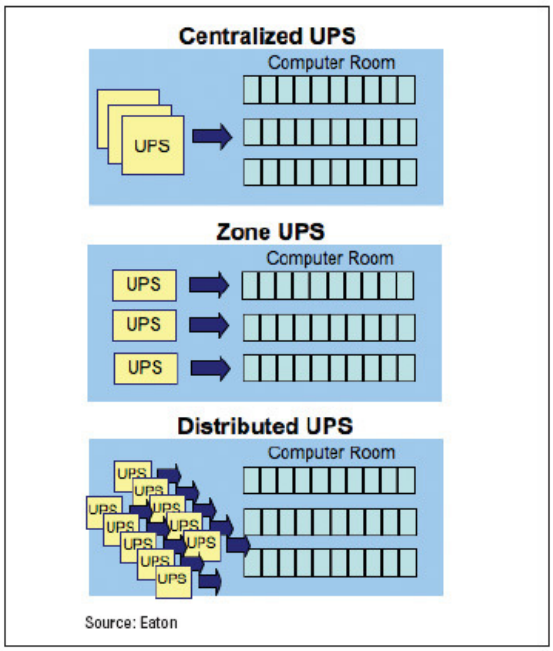

Where should you locate the UPS?

You would probably install it in a separate equipment room when the UPS is large (>200 kVA) or uses flooded (wet cell) batteries, where computer room real estate is at a premium, or if maintenance technicians should not enter the computer room.

You might prefer to keep the UPS in the computer room if it is important that the UPS receive the same security and conditioned environment as other IT assets, to be managed like every other device in the data centers. The sleek design of modern UPSs blends well in today's computer rooms (see Figure 6).

How will you establish redundancy at the load level?

Most modern IT equipment can be procured with dual internal power supplies that are fed from separate power sources for maximum reliability. If one power source fails, or one IT equipment power supply fails, the system doesn't have to crash. This capability is required for Tier III or IV reliability.

To achieve Tier III reliability, one power path should be supported by an active power system with UPS protection. For Tier IV reliability, you would need multiple active and independent power paths with UPS protection (see Figure 7).

Figure 6. Modern UPSs are designed to blend with IT equipment

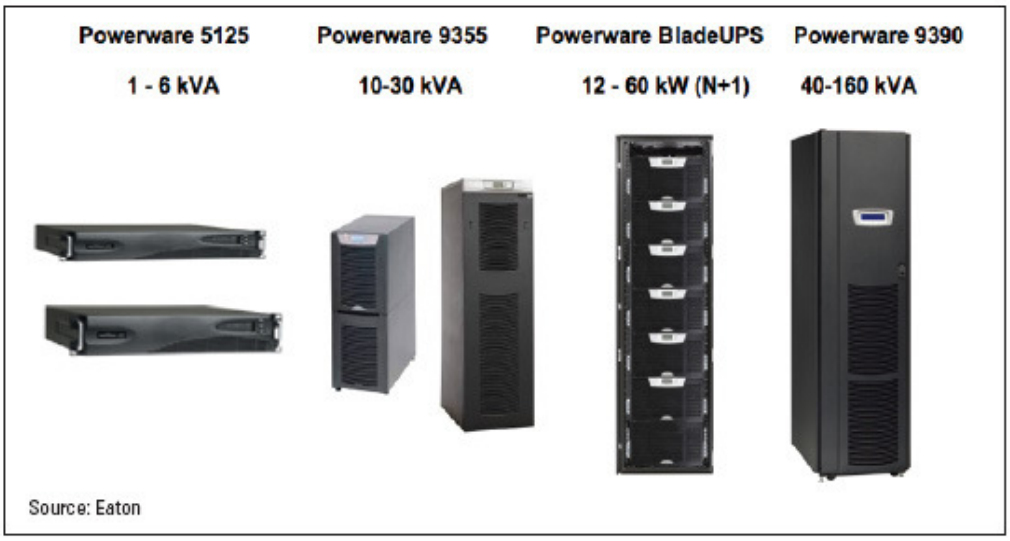

Figure 7. Power protection for IT equipment with redundant power supplies

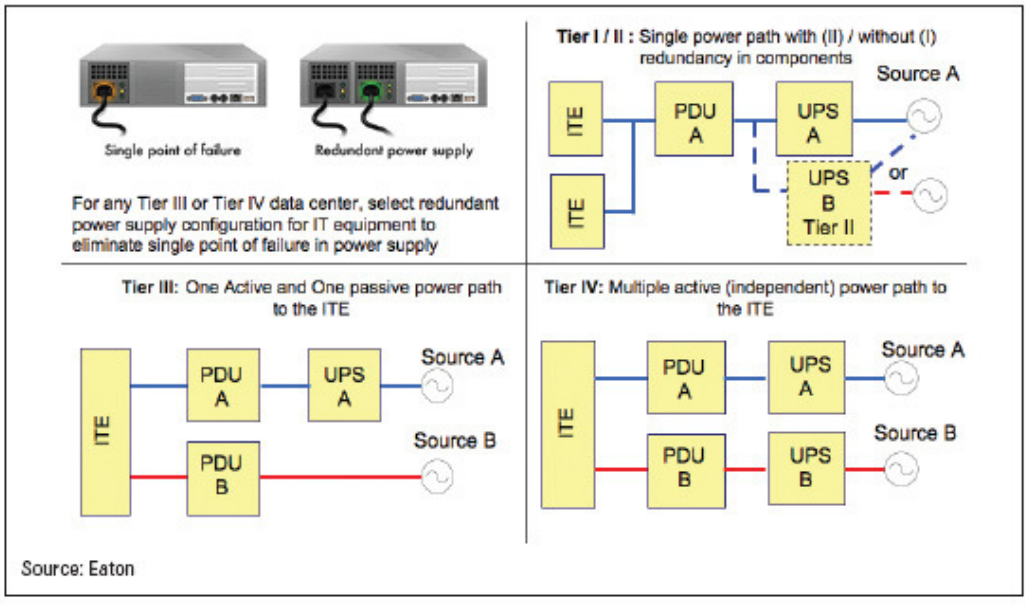

Conclusion

IT managers have many power distribution and protection options that should be aligned with current and future business needs and local utility conditions. The best strategy for your data center will not be dictated by any one-size-fits- all industry-wide standard. It depends on many factors, such as power consumption, data center configuration, reliability, cable management, metering and flexibility for future expansion. When planning an upgrade, expansion or new build of a power infrastructure, don't let anyone sell you on a so called "industry standard" one way. Standardization within your data center will drive simplicity and ease of management, but the right standards for your data center will probably be quite different from your competitors' or neighboring businesses. Demand that the vendor's engineers optimize the design for your unique requirements.

For more information or to discuss your data centre power infrastructure requirements contact:

Wright Line at **800-225-7348** or **www.wrightline.com**

At Wright Line we deliver customized, scalable, reliable solutions for today's technology intensive environments.