

Rack Powering Options for High Density

White Paper 29

Revision 7

by Neil Rasmussen

> Executive summary

Alternatives for providing electrical power to high density racks in data centers and network rooms are explained and compared. Issues addressed include quantity of feeds, single-phase vs. three-phase, number and location of circuit breakers, overload, selection of plug types, selection of voltage, redundancy, and loss of redundancy. The need for the rack power system to adapt to changing requirements is identified and quantified. Guidelines are defined for rack power systems that can reliably deliver power to high density loads while adapting to changing needs.

Introduction

Information technology (IT) refreshes in the data center and network room typically occur every 2 to 3 years. As equipment is changed, the power requirement, the voltage requirement, the redundancy requirement, and the connector requirement often change as well. As rack enclosures have become the standard means for housing and organizing computing and communication systems, the power distribution system for the rack enclosure must adapt to these changing requirements.

Power density predictions for racks in data centers have sharply escalated as a result of the high power density of the latest generations of computing equipment. Off-the-shelf IT equipment such as 1-U servers or blade servers can draw 20 kW or more in a fully populated rack. This density cannot be supported in a data center environment where the average rack is fed by a single 120 V 20A power circuit. Twenty of these circuits would be required per-rack to support a 20 kW load in a dual-path electrical environment.

The power requirements of modern computing equipment vary as a function of time depending on the computational load. Until the year 2000, this variation was very small and could be ignored for almost all computing and communication systems. However, the implementation of power management technologies into processors and servers began during the year 2000; today the fraction of computing equipment which has a substantial variation in power consumption in response to the computing load is increasing. This variation can be as high as 200% of the baseline power consumption of the equipment. The power distribution system design for a rack enclosure must comprehend this variation.

This paper is focused on AC rack power distribution. DC power distribution has a very limited role in the modern high density data center, as explained in White Paper 63, [AC vs. DC for Data Centers and Network Rooms](#).

This paper is limited to a discussion of North American voltage and connector standards. The appropriate rack power distribution strategy is considerably different for the 230 V systems, which predominate in most of the world.

Historic means for providing rack power

The most common approach today is to design, engineer, and install power solutions specific to a rack enclosure. Should the requirements for that rack enclosure change, an alternative power solution must be designed, engineered, and installed. While this approach can comprehend any unique power requirement, it involves significant planning, engineering, and rewiring. Rack enclosures are usually fed from a common power distribution panel within the data center or network room. In many instances, this panel cannot be de-energized in order to adapt a rack enclosure(s) power distribution system (i.e. install another breaker). The result is known as “hot work” and not only introduces a very serious safety hazard, but a high degree of risk of creating a fault in the circuit being worked on and / or dislodging / faulting adjacent wiring circuits. Such errors result in undesirable downtime.

Ideally, the rack enclosure power system would be adaptable to any realistically possible combination of equipment, on demand, without the need to perform any work that would be a hazard to safety or that might adversely affect system availability.

Rack powering requirements

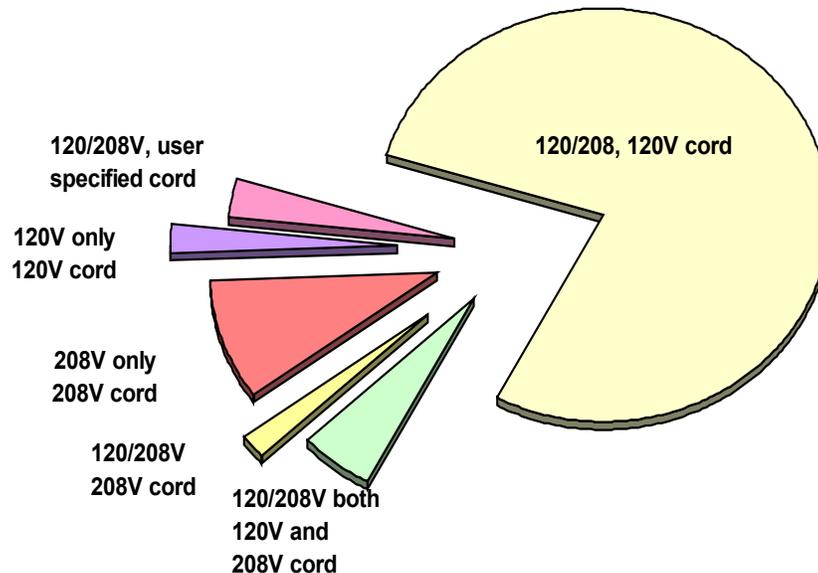
The various dimensions of rack enclosure power requirements are summarized in the following sections. The nature of the requirements is outlined and rational design approaches are summarized.

Voltage requirements

In North America, data centers are provided with both 120 V and 208 V power. The voltage requirement and supplied power cord of the IT equipment break down approximately as follows:

Figure 1

Voltage requirement and supplied power cord for IT equipment in North America



This complex situation suggests the need for the rack power distribution system to provide both 120 V and 208 V. However, it is possible to consider using a system restricted to a single voltage, either 120 V or 208 V. These two possibilities are discussed below.

The selection of 120 V as a single voltage standard for a data center seems most convenient because almost 95% of equipment is provided with a 120 V power cord. Unfortunately, the 5% of equipment that accepts only 208 V is often the most important and mission critical equipment, such as large routers and blade servers. **Therefore, it is not realistic to design a rack power distribution system based solely on 120 V except for very small network rooms.**

The selection of 208 V as a single voltage standard for a data center assures compatibility with over 97% of equipment, including the most critical equipment. However, the supplied 120 V power cords of the majority of the equipment must be replaced by the user with the appropriate 208 V cord, requiring that the user keep appropriate replacement cords on-hand. For some equipment it is also necessary to switch the power supply from 120 V to 208 V operation with a selector switch; the failure to activate this switch on equipment so-equipped can lead to catastrophic failure when powered by 208 V. The 3% of equipment that only operates from 120 V can be excluded from the data center, because in almost all cases these devices are small accessory equipment that has acceptable and readily available substitutes that will operate on 208 V. Nevertheless, providing 120 V in the rack power system can be a very significant convenience, eliminating many plug incompatibilities. For this reason nearly all existing data centers in North America provide both 120 V and 208 V and virtually none have standardized on 208 V exclusively.

The rack environment is single-phase. There is an **insignificant** quantity of rack mounted IT equipment manufactured that requires three-phase power (Some brands of blade servers being notable examples). Occasionally, a pre-configured OEM rack enclosure is wired using an internal power distribution unit (PDU) that takes in three-phase power and provides three branches of single-phase power to the single-phase IT loads. It is important to note that

these IT loads are actually single-phase. Despite the absence of three-phase loads, there is a good case to be made that three-phase power should be distributed to racks as will be shown later in this paper. One key benefit of providing 3-phase to the rack is that an inherent property of three-phase distribution is the ability to power 120 V and 208 V loads from the same branch circuit.

Power requirements

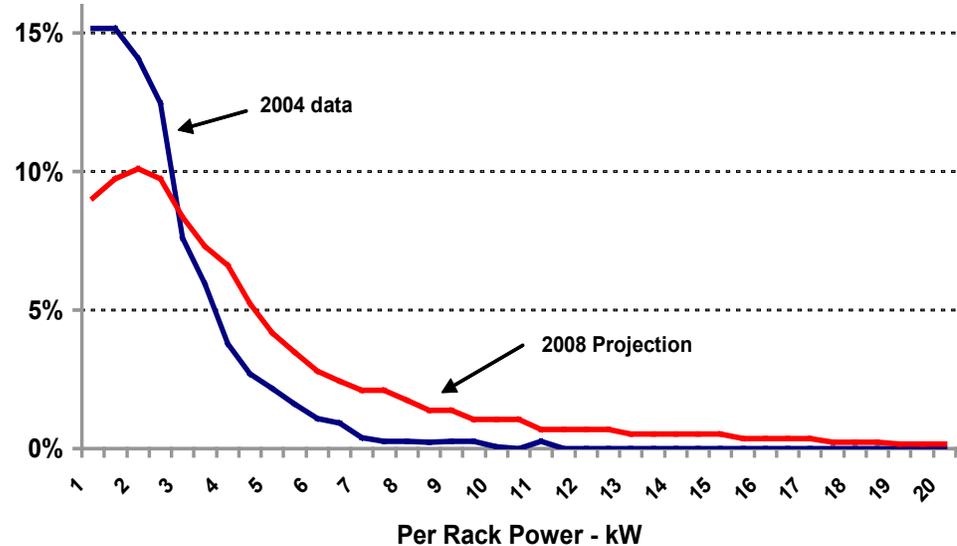
Power densities within the rack enclosure can vary greatly dependent upon the equipment installed. In the extreme low load case, a rack enclosure may only have passive patch panels or a few internetworking switches with a power draw of <100 W. In the extreme high load case, a rack enclosure may be completely filled with high-density servers for a total load of 20 kW or more.

In addition to supplying the total rack power requirement, the rack power system must also be able to provide the required power to an individual device. Sending multiple branch circuits to a rack may appear to provide the total power requirement, but the power requirement of an individual large load may exceed the capability of any of the branches. For example, sending any number of 20 A branch circuits to a rack where a single piece of equipment requires 30 A is insufficient. Another example is a blade chassis with a 30 A plug that may be initially populated with only a few blades and use 5 A on a 30 A circuit. Some users may think they can put multiple blade chassis' on a single 30 A circuit, but as they populate the chassis they overload the circuit. In cases like these, it is recommended that only one load device be attached to each branch circuit.

The appropriate design value for average rack power level is a subject of considerable controversy. A Schneider Electric survey of usage patterns in corporate data centers, network rooms, and communication rooms for the year 2004 identified the distribution of per-rack power consumption shown in **Figure 2**. This graph shows the frequency of occurrence of racks configured to different power levels. The frequency of occurrence goes down with increasing power level. 95% of racks draw power below 6.5 kW.

A projection of usage in the year 2008 (based on technology / client trends) is also shown in **Figure 2**. This indicates that the average power per rack is increasing over time. It is possible to configure IT equipment today that would exceed a 20 kW per-rack power requirement if fully populated into a rack enclosure. While possible to achieve, we did not find this occurrence to be frequent in real-world installations. The data collected indicates that the average power density per rack enclosure will rise significantly. However, power densities >6 kW will still remain a small fraction of the installed base.

Figure 2
Per-rack power consumption frequency distribution



An analysis of the underlying data of the distribution of rack power indicates the following:

- The very low loads are mainly rack enclosures with wiring patch panels, switches, and hubs
- Loads in the 1 kW range are mainly sparsely populated rack enclosures
- Loads in the 2-3 kW range are mainly rack enclosures that are populated with typical equipment but with significant unfilled rack space
- Loads in the 5 kW range are partially loaded with 1U servers, or contain a mix of technologies
- Loads in the 7 kW + range are extremely rare but, according to customers, are going to become more common with the recent density increases resulting from server technology advancements

The mean value for rack power consumption in corporate computing environments is around 1.5 kW. However, organizations were found to have variable mean power consumption within their overall rack enclosure environment. These organizations take different approaches as to how much equipment they pack into a rack enclosure. Some leave large amounts of space within the rack enclosure unused while others may pack equipment as tightly as possible. Therefore, the market-wide mean rack power consumption is not necessarily a good predictor of average rack power consumption within an organization.

The electrical circuits between the last over-current protector and the equipment are called “branch circuits”. It is important to understand that electrical codes specify that the branch circuit rating of a circuit is dictated by the load and cannot be arbitrarily chosen by the user. In the USA, virtually all branch circuits in rack enclosures are rated for 20 A¹.

Table 1

Branch circuit power limitations

| Voltage | Max branch current rating | Maximum kW capacity per branch | Total rack power for 1/2/3/4 branch circuits |
|---------|---------------------------|--------------------------------|--|
| 120 V | 20 A | 2.4 kW | 2.4 / 4.8 / 7.2 / 9.6 |
| 208 V | 20 A | 4.2 kW | 4.2 / 8.3 / 12.5 / 16.6 |

The maximum power available to the rack enclosure depends on the number and type of branch circuits provided within the rack enclosure. Clearly the number of branch circuits will need to be greater than one to support the power density of current and future IT technology.

Combining the data from **Table 1** and data from **Figure 2**, the following conclusions can be drawn:

- A single 120 V branch circuit can supply the load requirement for the most common rack enclosures today, but this will not be true in the future.
- Two 120 V branch circuits, or one 208 V branch circuit, can supply the load requirement for approximately 80% of rack enclosures today, but only 40% of racks in the future.
- Three 120 V branch circuits in North America can supply the load requirement for over 92% of rack enclosures today, and still over 85% in the future.

¹ Note that 30 A circuits are also sent to the rack but these circuits are feeder circuits and not branch circuits, because they require additional circuit breakers in the rack as described in later sections. Very few rack devices can directly utilize a 30 A branch circuit; these are usually blades, routers or free-standing servers or storage devices.

Note that the inability to provide sufficient branch circuits to a rack enclosure does not prevent operation of the system. If the rack has insufficient power distribution capacity, the power drawn by a rack enclosure can be reduced by removing equipment from it and moving it to another rack enclosure. However, the consequence is that a reduction in space utilization occurs. For an occasional rack enclosure this is not a serious problem. The costs and benefits of spreading loads within the data center are discussed in other white papers.

An adaptable rack enclosure power system would be able to provide enough power to supply the maximum anticipated load to any rack enclosure at any time without re-engineering the power system. In North America three 120 V branch circuits per rack enclosure or two 208 V branch circuits is a practical design baseline, with the ability to add additional circuits easily as needed.

Redundancy requirements

Providing redundancy and/or fault tolerance in the power system can increase the availability of a computing system. In high availability environments, a common way to provide redundancy is to supply two independent power paths to each piece of computing equipment; the equipment in turn accepts the two power feeds via independent paralleled power supplies that are sized such that the equipment will continue to operate with only one power path.

This system provides the following key advantages:

- If a power supply fails the system continues to operate
- If one power feed fails due to equipment malfunction the system continues to operate
- If one power feed fails due to user error the system continues to operate
- If the power supply fails in a way which faults the power feed and trips breakers, the equipment sharing the breaker is not affected.
- If one power feed needs to be shut down for maintenance or upgrade, the system continues to operate

For this approach to be effective, the following requirements must be met:

1. The protected equipment must support dual power feeds and operate with one feed faulted
2. The loading of breakers within each power path must always be less than 50% of trip rating during normal conditions, so that the increase in load that will accompany failure of the alternate path does not cause breakers to trip. This also helps prevent tripping of the alternate path due to low line voltage conditions.

Meeting these two requirements can be very difficult. Some computing equipment is only available with a single power cord. There is also equipment manufactured with three power cords, where any two are needed for proper operation. These types of equipment cannot operate with the loss of one power feed. In these cases an Automatic Transfer Switch (ATS) can be used which generates a single feed from two inputs. Such an ATS may be deployed centrally or it may be deployed in a distributed manner by placing small rack mount ATS in the rack enclosure with the protected equipment. For more information see [White Paper 48, *Comparing Availability of Various Rack Power Redundancy Configurations*](#).

An adaptable rack enclosure power system would be able to support a single or dual path environment or a hybrid of both single and dual equipment. In addition, it is necessary to provide current monitoring to ensure that all circuits are loaded below 50% capacity in order to prevent breaker tripping during a loss of one power path.

Redundancy requirements

One of the most misunderstood concepts of power distribution is the over-current protection of branch circuits. Each branch circuit in a rack enclosure must be supplied by an independent circuit breaker and a typical rack enclosure will need multiple branch circuits. **Figure 3** illustrates the common methods of power distribution to the rack, showing different branch circuit configurations. In **Figure 3a**, a single branch circuit supplies a rack. For a 20 Amp system, this arrangement is limited to 2.4 kW max capacity for 120 V or 4.2 kW max capacity for 208 V. To achieve higher rack power multiple branch circuits are required. There are two options for providing multiple branch circuits to a rack enclosure, and these are shown in **Figure 3b** and **3c**.

Figure 3

Illustration of methods of delivering branch circuits to the rack, showing alternate ways of supplying multiple branch circuits to a rack

3a.

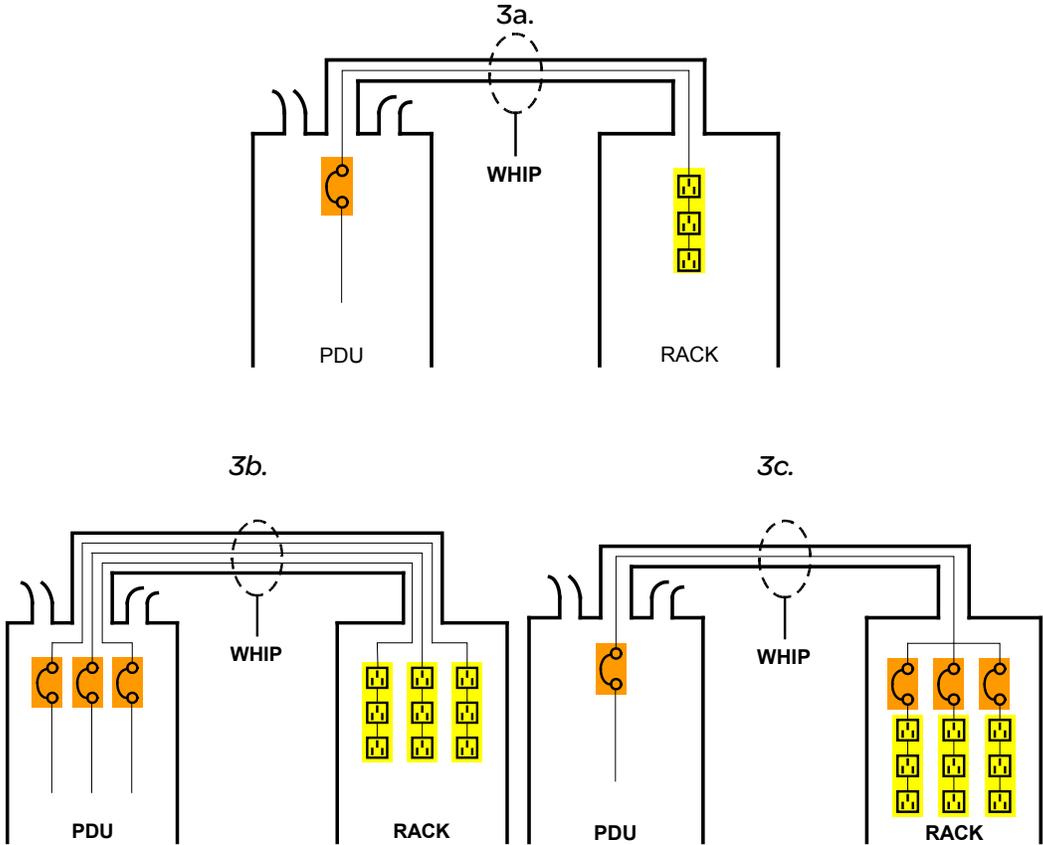
Single branch circuit to the rack

3b.

Multiple branches: generate required branch circuits at the PDU

3c.

Multiple branches: generate required circuits within the rack



The arrangements of **Figure 3b** and **3c** are capable of providing the same power, but with a different wiring and circuit breaker arrangement. Note that in **Figure 3b** the conduit or “whip”² contains multiple branch circuits. These alternatives give rise to significant advantages and disadvantages as summarized in **Table 2**.

² A whip is a power feed to a rack. A whip could potentially have multiple branch circuits running to the same rack. For example, a 3-phase whip is fed from a 3-pole breaker or 3 single-pole breakers.

Table 2

Comparison of the two options for providing multiple branch circuits to a rack enclosure

| | Number of feeds | Location of breakers | Advantages | Disadvantages |
|---|--|---|--|--|
| Generate required branch circuits at the PDU | One feed per required branch circuit | PDU only | Fewer total breakers No fault coordination issues Each branch usable to full capacity Feed to rack may be pluggable | More circuit breaker positions required in the PDU |
| Generate required branch circuits within the rack enclosure | Fewer feeds per voltage required within the rack enclosure | Feed breakers in PDU Branch breakers located within rack enclosure | Fewer feeds to rack Fewer breakers in the PDU | More breakers to monitor for overload. Fault coordination. Branches may not be usable to full capacity. More breakers in series reduces reliability Feed may need to be hardwired. |

The summary of **Table 2** suggests there is a significant advantage to avoiding the need to generate branch circuits in the rack enclosure.

Most users who send 30 A or larger circuits to racks do not understand that these circuits are feeder circuits and **not** branch circuits. When 30A whips are used, the typical 20A branch circuits required to supply receptacles **must** be provided by branch-rated breakers located **within the racks**.

An adaptable rack power system would eliminate the need for hard wiring, cascaded breakers, and breaker coordination analysis, which would suggest a preferred approach of multiple branch circuits per rack. Ideally, multiple branches would be provided using a single multi-conductor cable to the rack in order to simplify wiring to the rack.

Connector requirements

Over 99.99% of all AC powered equipment used in rack enclosures are connected via a power cord and plug (a negligible fraction is hard-wired). While most have plugs, these plug types can vary greatly from one product to another.

In North America, the approximate statistical breakdown of frequency of occurrence for plug types is provided in **Table 3**:

Table 3

Frequency of occurrence of different plug types on IT equipment in North America

| Frequency | Plug type | Description | Where used |
|-----------|-------------------------------------|-------------------|--|
| 80% | NEMA 5-15P | 120 V 15A | Small servers, hubs, departmental switches, monitors, power blocks |
| 13% | IEC 60320 (previously IEC-320) C-13 | 120 or 208 V 15A | Enterprise servers |
| 3% | NEMA L6-20P | 208 V 20A Locking | Enterprise servers, routers |
| 2% | IEC-60320 C-19 | 120 or 208 V 20A | Some blade servers |
| 1% | NEMA 5-20P | 120 V 20A | Large departmental servers |
| 0.9% | NEMA L6-30P | 208 V 30A Locking | Enterprise servers, routers |
| 0.1% | NEMA 6-15P | 208 V 15A | Large departmental servers |
| 0.01% | NEMA 15-30P | 208 V 30A 3-phase | Some blade servers |

In the rest of the world, plugs vary by country, which makes this situation more complex. However, unlike North America, most countries use a single service voltage and therefore power cords are rated for 230 V 16 A. This means that a single plug type serves most applications from small hubs to enterprise servers. Therefore, outside of North America, a single country-specific plug type handles 99% of applications. For more information on the different types of North American plug types see White Paper 20, [The Different Types of AC Power Connectors in North America](#).

It is important to note that 99% of all IT equipment in the data center utilizes a detachable power cord. This means that the plug type can be changed. OEM suppliers take advantage of this when creating complete rack systems and often will install cords with IEC 60320 (previously IEC-320) style power plugs on all the equipment in the rack enclosure, along with an IEC60320 style outlet strip. This has the benefit that a single configuration can be used worldwide.

The number of receptacles required in a rack enclosure varies dramatically with the installed equipment. A rack enclosure may contain only a single load as a minimum. Conversely it may be populated with (42) thin servers with dual power cords for a total requirement of 84 receptacles.

An adaptable rack enclosure power system would be able to provide power receptacles for all the various plug types, which might be encountered, as well as two feeds, each containing (42). To accomplish this, one must provide a large quantity and assortment of receptacles in every rack enclosure, or provide a number of easily changeable outlet strip options to meet evolving requirements.

Harmonic requirements

Historically, computing equipment generated harmonic currents on AC power lines, which led to the need to incorporate specialized features into power systems such as oversized neutral wiring and K-rated transformers. During the 1990's, regulations placed on the design of computing equipment, combined with the gradual retirement of older equipment, resulted in the elimination of this as a problem by the year 2000. Today, lighting and power distribution transformers are the predominant source of harmonic currents in data centers. Harmonic-rated wiring and transformers are not needed in the rack enclosure power environment. See White Paper 26, [Current Harmonics and Neutral Overloads](#) for more information on this subject.

De-rating requirements

Operating with circuits at the rated current limit is not desirable since the circuit breakers are on the edge of the trip rating and may trip. Therefore the user should not expect to utilize the full current or power rating of the system, but should apply a de-rating factor. The National Electric Code specifies a 20% de-rating factor be applied to rack mounted PDUs. This is a practical and recommended de-rating factor. However, some users conservatively specify lower de-rating factors, such as 30% or even 40%. The power capacities of the distribution architectures described in this paper are full rated values unless otherwise stated, and therefore the actual values should be computed by applying the desired de-rating factor.

Cabling requirements

Cables to deliver power to the rack enclosures are an essential part of the rack enclosure power system. A common practice today is to use under floor power cabling. The under floor power cabling method presents a number of barriers to adaptability, which are described in White Paper 19, [Re-examining the Suitability of the Raised Floor for Data Center Applications](#).

In an adaptable rack enclosure power system, the cabling provided to each rack would provide all the branch circuits and voltages that might ever be required. No changes to the cabling would be required due to equipment changes in the rack enclosure. It would also be easy and safe to provide the appropriate power feeds to additional rack enclosures in the future.

Current monitoring requirements

Rack enclosure power systems are subject to constant load changes due to the installation and removal of equipment and to the dynamic power draw variation in the installed equipment. These circumstances lead to a requirement to monitor power flowing in branch circuits in order to prevent failures or hazard due to overloads. This subject is described in detail in White Paper 43, [Dynamic Power Variations in Data Centers and Network Rooms](#).

Consistency requirements

Due to the large number of power circuits in the typical data center, there is a significant advantage to minimizing the different types of power distribution provided (branch circuit ratings, poles per whip, circuit breaker types and location, etc). Ideally a uniform single type of power feed would be provided to every rack to maximize flexibility and reduce human error.

Human error is a constant threat in a data center and has been the cause of much downtime. Standardizing on a common power distribution circuit that fulfills the need 97% of the time is

Selecting the appropriate power distribution system

just one method of reducing the risk of human error. With standardized whips users are less apt to become confused, parts are minimized and learning curves are accelerated all which lower the risk of a costly mistake.

Despite the number of requirements, there are still many different combinations of circuits that can be used to power rack enclosures, each providing different total power capacity and differing in key features. There are at least 25 practical but different ways to provide power to rack enclosures in the range of 2.4 kW to 36 kW per rack. The details of these alternatives are provided in **Appendix A**.

Through a systematic investigation of these alternatives, it is possible to determine that these alternatives are not equivalent in their costs and benefits, and that some options are clearly preferred. The analysis of **Appendix A**, when considered with the requirements defined in the previous section, suggests that there are four essential preferred forms of rack power distribution between the PDUs and the racks, which are used in multiples per rack to achieve a desired power density. The four preferred forms are:

1. 208 V 20 A whips
2. 120 / 208 V 20 A 3-phase whips
3. 208 V 50 A 3-phase whips
4. Dedicated whips

Figure 4 illustrates these four preferred forms by showing the branch circuit configurations. The characteristics and advantages of these four essential preferred types of power distribution are provided in **Table 4**. In **Table 4**, the shaded attributes represent the best performance for that characteristic. The figure shows the clear advantage of 3-phase distribution to the rack.

The analysis of this paper suggests that distribution whips that include single or 3-phase 30 A whips are not preferred except in the case of dedicated whips. The 30 A whip value is not optimal because the most common branch circuit size that must be generated in a rack is 20 A, which leads to two undesirable problems with a 30 A whip size:

1. The coordination between a 20 A branch circuit breaker and a 30 A feeder breaker is difficult to achieve increasing the likelihood of cascade breaker tripping.
2. Two 20 A branches are required to fully utilize a single 30 A whip phase feed, and if one of these branches is fully utilized the other can only be halfway utilized. This is an inefficient use of circuit breakers; furthermore, the feeder breaker may trip before the branch breakers trip.

When a density higher than that provided by one or more 20 A or 20 A 3-phase whips is required, a 50 A 3-phase whip is the preferred solution over a 30 A 3-phase whip.

Figure 4
Illustration of the four preferred forms of power distribution to the rack

4a.
208 V 20A whip

4b.
208 V 20A 3-phase whip

4c.
208 V 50A 3-phase whip

4d.
Dedicated whips

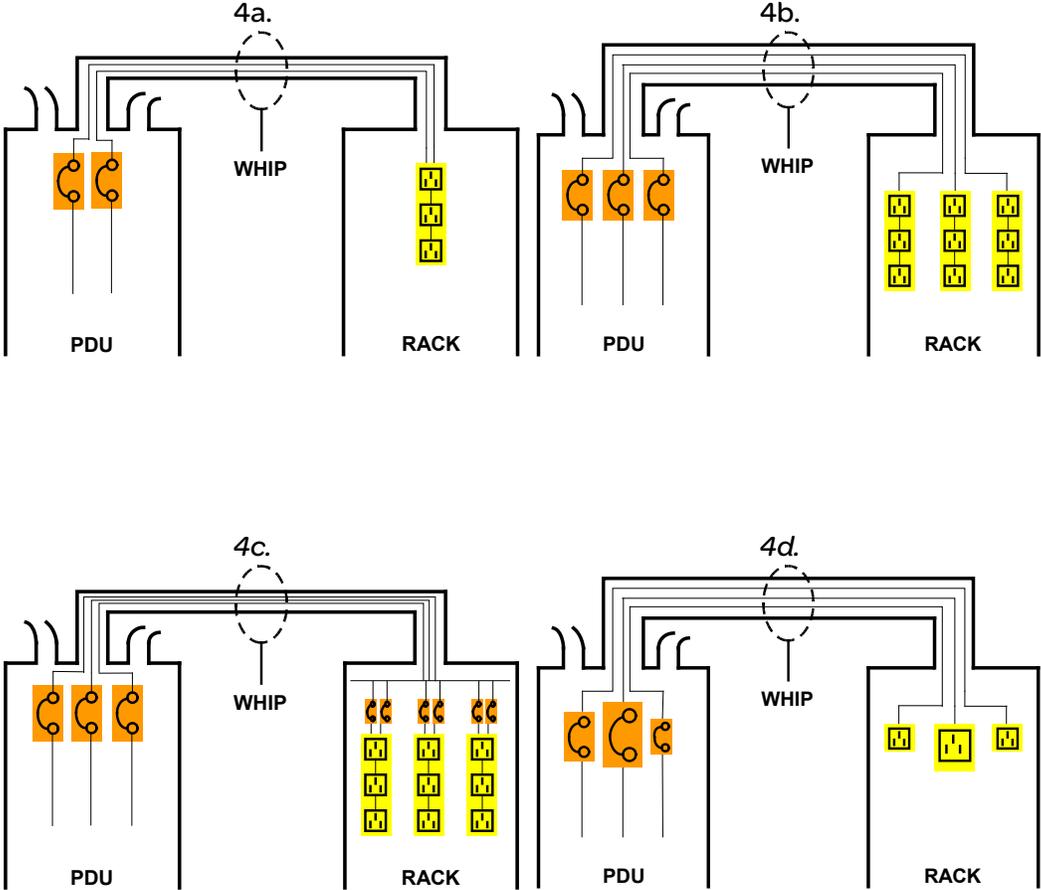


Table 4

Characteristics of the four preferred types of rack power distribution

| Characteristic | 208 V 20A | 120/208 V 20A 3-phase | 208 V 50A 3-phase | Dedicated 208 V | Comment |
|--|----------------------------|-----------------------------|-------------------------|---|---|
| kW levels obtainable from 1, 2, 3 or 4 whips | 4.2 8.3 12.5 16.6 | 7.2 14.4 21.6 28.8 | 18 36 54 72 | 4.2; 6.2 8.3; 12.5 12.5; 18.7 16.6; 25 | Dedicated whips assume 20A or 30A single phase whips. Note that 3-phase dedicated whips are also possible but extremely rare. |
| kW per whip | 4.2 | 7.2 | 18 | 6.2 | |
| kW per PDU breaker panel pole position | 2.1 | 2.4 | 6.0 | 3.1 | |
| Max single load kW | 4.2 | 4.2 | 10.4 | 6.2 | There are currently almost no production IT rack products that require over 4.2kW per plug, but this could change. |
| # of 1kW IT devices per whip | 4 | 6 | 18 | 1 | Note that a dedicated feed only supports a single device. |
| # of breakers in series with load | 1 | 1 | 2 | 1 | |
| 120V support | N | Y | N | N | The ability to handle 120V circuits is an advantage since most equipment comes with a 120V power cord. |
| Fault coordination | E | E | G | E | Placing additional branch breakers in the racks, as required for the 30A designs, compromises fault coordination. |
| Cost per kW @ 2kW per rack | \$350 | \$310 | \$475 | variable | Includes source breaker, whip, and rack PDU. Note that a dedicated whip is required for each IT device. |
| Cost per kW @ 10kW per rack | \$275 | \$124 | \$95 | variable | Includes source breaker, whip, and rack PDU. |

Note: Shading indicates best performance for the characteristic

Note: whip costs do not include any installation costs

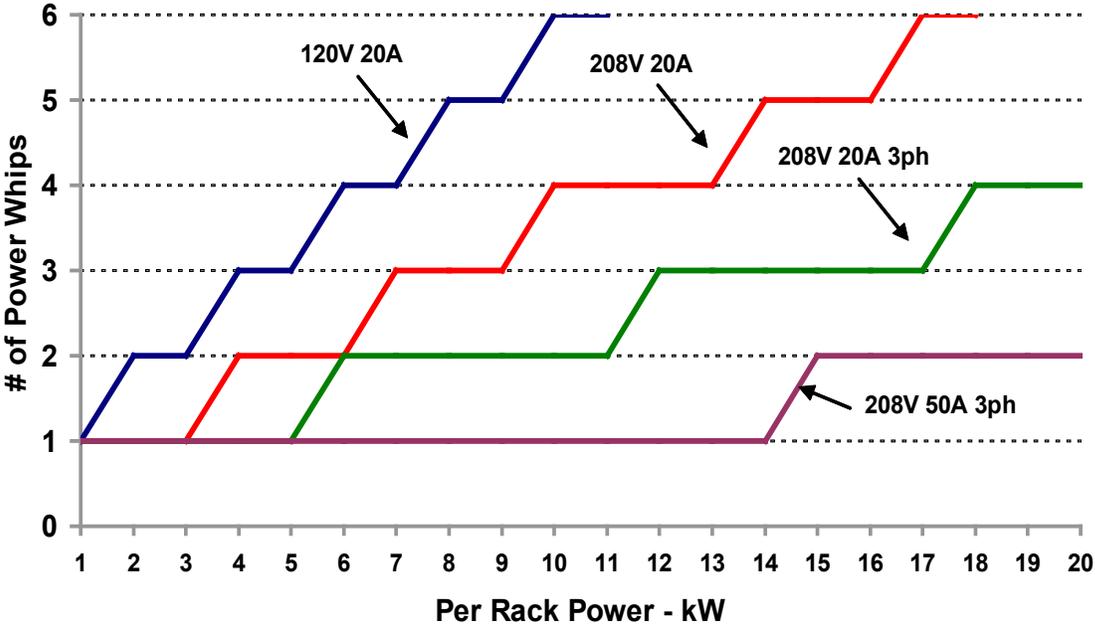
Fault coordination – legend: E = excellent G = good P = poor

The preferred forms may be used in multiples to achieve higher total rack power ratings. **Figure 5** shows the number of whips required to achieve increasing total rack power ratings. The four curves on the chart represent the whip count for various types of power distribution. The 120 V 20 A rack power distribution circuit is not a preferred approach but is in the chart for reference. Note that the number of whips required increased linearly with rack power rating as expected. For power ratings greater than 11 kW per rack the number of whips becomes very large for all systems except the 208 V 50 A 3ph system.

Note that for 2N dual power path systems the number of whips in **Figure 5** must be multiplied by 2.

Figure 5

Number of rack power whips or feeds required as a function of the rack power rating



Note: this chart assumes whip derating to 80% of rated capacity

Dedicated whips

One of the preferred forms of power distribution is dedicated whips. However, dedicated whips are preferred only when the load equipment in a rack has a small number of power cords that each draw a large amount of power, particularly if the equipment uses an unusual plug type or power configuration. Certain blade servers and SAN storage units meet these criteria. For example, a large blade server power subsystem that has a 3-phase NEMA L15-30 power plug. The advantage of dedicated whips is that they never require an additional circuit breaker in the rack, increasing reliability and saving expense. The big disadvantage of dedicated whips is the lack of flexibility when equipment is changed in the future. Dedicated whips are compared to standard 208 V 50 A whips in **Table 5**.

Table 5

Comparison of the two options for providing very high density power

| | Number of whips | Method of change | Advantages | Disadvantages |
|--|-------------------------------------|---|--|---|
| Standard 208V 50A 3-phase whips | One whip per every 18kW | Easy: Plug in different breaker / outlet strips in the rack | Fewer breakers in the PDU Various outlets provided for small or unplanned equipment | Some breakers located in the rack |
| Dedicated whips | One whip per power cord in the rack | Difficult: Power down and run new wires to the PDU | No fault coordination issues Fewer total breakers Can handle any strange plug or circuit size Lowest cost | Must plan in advance and know every power cord that might exist in the rack No provision for small ancillary IT equipment Feed needs to be hardwired. |

In general, dedicated whips should only be used when the power requirement or plug configuration cannot be provided by a standard 208 V 50 A 3-phase circuit, or when cost is much more important than the ability to reconfigure the rack later.

Distribution selection strategy

From this analysis it is possible to draw conclusions regarding a preferred arrangement for branch circuits, which is as follows:

1. Use a single 120/208 V 20 A 3-phase power whip to supply common medium density racks up to approximately 6 kW / rack; supply this by default to every rack.
2. Use two 120/208 V 20 A 3-phase whips to supply higher density racks up to approximately 12 kW per rack.
3. For densely packed 1U server or blade server applications, use either one or two 208 V 50 A 3-phase whips.
4. For certain extremely high density loads that have input current requirements per power cord of over 20 A, run separate dedicated 20 A or 30 A single or three phase whips per device, depending on the device plug type (consult application notes for specific recommendations for specific brands and models of blade servers).

Adaptable power architecture for rack enclosures

The growing recognition of the issues described in this paper have led data center and network room designers and operators to develop their own ingenious solutions to address the need for adaptable rack enclosure power systems. Nevertheless, an integrated and cost-effective approach from the point of view of the equipment providers has been lacking. A fully integrated approach would include a modular system that implements all aspects of power distribution from the AC mains connection of the facility, through the UPS, through the power panels, all the way down to the mechanics of connection of the plugs in the rack enclosures.

The first integrated and adaptable rack enclosure power system was introduced in 2001 and is shown in **Figure 6**. InfraStruxure® system includes components that are engineered to meet the requirements of an adaptable rack enclosure power system. The system includes prefabricated multi-branch power distribution whips, Quick-change multi-voltage metered outlet strips with various receptacle configurations, pre-engineered circuit breaker coordination, single and dual-path power feed support configurable at the rack enclosure or row level, point of use DC capability, and rapid conduit-free installation. This system is provided as an integrated system that can be configured to order from stock components.



Figure 6

An example of an adaptable rack power system

In addition to the capability of the adaptable rack enclosure power system to respond quickly and economically to change, there are cycle time and cost advantages associated with the initial installation of the system, including a dramatic simplification to the up-front engineering and installation work associated with data center design. Furthermore, the ability to adapt the rack enclosure power system can allow the system to be “right sized” to the actual load requirement and grow with expanding needs. The economic benefits of rightsizing can be well over 50% of the lifecycle cost of a data center or network room and are discussed in more detail in White Paper 37, [Avoiding Costs from Oversizing Data Center and Network Room Infrastructure](#).

Conclusion

Individual rack enclosure power consumption in the data center or network room varies widely and is expected to grow in the next few years. Rack enclosure equipment is replaced 5 or more times during the life of a data center in a piecemeal manner. This situation requires a rack enclosure power distribution system that can cope with the changing requirements. Key requirements of an effective rack power distribution system were described, which suggest a practical rack enclosure power architecture that can meet the requirements for an adaptable rack enclosure power system. The recommended approach standardizes on four key ways to distribute power, along with a strategy for selection of the best approach for a given installation. When this approach is implemented, the result is a power distribution system which reduces human error, adapts to changing requirements, minimizes the need for advance planning and meets the requirements of high density IT equipment.



About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world’s largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 19 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC InfraStruXure system.

Prior to founding APC in 1981, Neil received his bachelors and masters degrees from MIT in electrical engineering, where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.



[AC vs. DC for Data Centers and Network Rooms](#)

White Paper 63



[Current Harmonics and Neutral Overloads](#)

White Paper 26



[Comparing Availability of Various Rack Power Redundancy Configurations](#)

White Paper 48



[The Different Types of AC Power Connectors in North America](#)

White Paper 20



[Re-examining the Suitability of the Raised Floor for Data Center Applications](#)

White Paper 19



[Dynamic Power Variations in Data Centers and Network Rooms](#)

White Paper 43



[Avoiding Costs from Oversizing Data Center and Network Room Infrastructure](#)

White Paper 37



[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

Appendix A: detailed analysis of an exhaustive set of rack power delivery options

Various practical circuit types, including multiples, which can provide rack power, are listed in **Table A1**. The table is ranked by increasing rack power capacity.

Table A1

Characteristics of practical rack power circuit configurations

| total kW | circuit type | # whips | kW / whip | # poles | kW / pole | max single load kW | # breakers | 120V support | 208V support | coordination |
|----------|------------------|---------|-----------|---------|-----------|--------------------|------------|--------------|--------------|--------------|
| 2.4 | 120V 20A | 1 | 2.4 | 1 | 2.4 | 2.4 | 1 | Y | | E |
| 3.6 | 120V 30A | 1 | 3.6 | 1 | 3.6 | 2.4 | 3 | Y | | P |
| 4.2 | 208 20A | 1 | 4.2 | 2 | 2.1 | 4.2 | 2 | | Y | E |
| 4.8 | 120V 20A | 2 | 2.4 | 2 | 2.4 | 2.4 | 2 | Y | | E |
| 6.2 | 208 30A | 1 | 6.2 | 2 | 3.1 | 6.2 | 6 | | Y | P |
| 7.2 | 120V 20A | 3 | 2.4 | 3 | 2.4 | 2.4 | 3 | Y | | E |
| 7.2 | 120V 30A | 2 | 3.6 | 2 | 3.6 | 2.4 | 6 | Y | | P |
| 7.2 | 120/208 20A 3 ph | 1 | 7.2 | 3 | 2.4 | 4.2 | 3 | Y | Y | E |
| 8.4 | 208 20A | 2 | 4.2 | 4 | 2.1 | 4.2 | 4 | | Y | E |
| 10.8 | 120/208 30A 3 ph | 1 | 10.8 | 3 | 3.6 | 6.2 | 9 | Y | Y | P |
| 10.8 | 208 30A 3 ph | 1 | 10.8 | 3 | 3.6 | 6.2 | 9 | | Y | P |
| 12.6 | 208 20A | 3 | 4.2 | 6 | 2.1 | 4.2 | 6 | | Y | E |
| 12.4 | 208 30A | 2 | 6.2 | 4 | 3.1 | 6.2 | 12 | | Y | P |
| 14.4 | 120/208 20A 3 ph | 2 | 7.2 | 6 | 2.4 | 4.2 | 6 | Y | Y | E |
| 16.8 | 208 20A | 4 | 4.2 | 8 | 2.1 | 4.2 | 8 | | Y | E |
| 18.0 | 208 50A 3 ph | 1 | 18 | 3 | 6.0 | 10.4 | 9 | | Y | G |
| 18.6 | 208 30A | 3 | 6.2 | 6 | 3.1 | 6.2 | 18 | | Y | P |
| 21.0 | 208 20A | 5 | 4.2 | 10 | 2.1 | 4.2 | 10 | | Y | E |
| 21.6 | 120/208 30A 3 ph | 2 | 10.8 | 6 | 3.6 | 6.2 | 18 | Y | Y | P |
| 21.6 | 120/208 20A 3 ph | 3 | 7.2 | 9 | 2.4 | 4.2 | 9 | Y | Y | E |
| 21.6 | 208 60A 3 ph | 1 | 21.6 | 3 | 7.2 | 12.5 | 12 | | Y | G |
| 25.2 | 208 20A | 6 | 4.2 | 12 | 2.1 | 4.2 | 12 | Y | Y | E |
| 24.8 | 208 30A | 4 | 6.2 | 8 | 3.1 | 6.2 | 24 | | Y | E |
| 28.8 | 120/208 20A 3 ph | 4 | 7.2 | 12 | 2.4 | 4.2 | 12 | Y | Y | E |
| 36.0 | 208 50A 3 ph | 2 | 18 | 6 | 6.0 | 10.4 | 18 | | Y | G |

Fault coordination – legend: E = excellent G = good P = poor

The options described in **Table A1** are limited to a practical maximum of 6 whips, and includes only multiples where all of the whips are of the same circuit configuration. Dedicated whips may be some mixed combination of the whips in the above list per rack.

While this list is exhaustive, certain combinations can be eliminated if the objective is to provide power to high density rack enclosures. A power system for high density rack enclosures must provide 208 V capability since there are many high density IT loads that do not accept 120 V. Furthermore, a power system for high density loads should also be able to power single loads of at least 3 kW since there are many high density loads that require up to 3 kW. In addition, a power system for high density loads should not exhibit a poor breaker coordination ratio. By excluding options that do not meet these criteria, the list of **Table A1** is reduced to the list of **Table A2**.

Table A2

Characteristics of rack power circuit configurations suitable for high density loads

| total kW | circuit type | # whips | kW / whip | # poles | kW / pole | max single load kW | # breakers | 120V support | 208V support | coordination |
|----------|------------------|---------|-----------|---------|-----------|--------------------|------------|--------------|--------------|--------------|
| 7.2 | 120/208 20A 3 ph | 1 | 7.2 | 3 | 2.4 | 4.2 | 3 | Y | Y | E |
| 8.4 | 208 20A | 2 | 4.2 | 4 | 2.1 | 4.2 | 4 | | Y | E |
| 12.6 | 208 20A | 3 | 4.2 | 6 | 2.1 | 4.2 | 6 | | Y | E |
| 14.4 | 120/208 20A 3 ph | 2 | 7.2 | 6 | 2.4 | 4.2 | 6 | Y | Y | E |
| 16.8 | 208 20A | 4 | 4.2 | 8 | 2.1 | 4.2 | 8 | | Y | E |
| 18.0 | 208 50A 3 ph | 1 | 18 | 3 | 6.0 | 10.4 | 9 | | Y | G |
| 21.0 | 208 20A | 5 | 4.2 | 10 | 2.1 | 4.2 | 10 | | Y | E |
| 21.6 | 120/208 20A 3 ph | 3 | 7.2 | 9 | 2.4 | 4.2 | 9 | Y | Y | E |
| 21.6 | 208 60A 3 ph | 1 | 21.6 | 3 | 7.2 | 12.5 | 12 | | Y | G |
| 25.2 | 208 20A | 6 | 4.2 | 12 | 2.1 | 4.2 | 12 | Y | Y | E |
| 28.8 | 120/208 20A 3 ph | 4 | 7.2 | 12 | 2.4 | 4.2 | 12 | Y | Y | E |
| 36.0 | 208 50A 3 ph | 2 | 18 | 6 | 6.0 | 10.4 | 18 | | Y | G |

Fault coordination – legend: E = excellent G = good P = poor