Server Technology, Inc.

Planning, Implementing and Application of Cabinet Power Distribution Units

White Paper STI-100-007
2008-April-09

Headquarters
Server Technology, Inc.
1040 Sandhill Drive
Reno, NV 89521
United States of America
+001 (775) 284 2000 Tel
+001 (775) 284 2065 Fax
sales@servertech.com
www.servertech.com
www.servertechblog.com

EMEA
Server Technology Intl.
Sienna Court
The Broadway
Maidenhead
Berkshire
SL6 1NJ
United Kingdom
+44 (0) 1628 509053 Tel
+44 (0) 1628 509100 Fax
salesint@servertech.com

APAC
Server Technology, Inc.
37th Floor, Singapore Land Tower
50 Raffles Place
Singapore, 048623
+65 (0) 6829 7008 Tel
+65 (0) 6234 4574 Fax
salesint@servertech.com

Overview

Over the past decade, our reliance on networks and data centers to provide us with information, to run our businesses, to access and to store data, and facilitate global commerce has grown considerably. As businesses and organizations plan for the next decade and build out their data centers, power consumption and distribution to the devices in the cabinet must be taken into consideration early on in the planning stages. Increased densities, proper installation, power and cord choices all affect the efficiency and reliability of the devices in the cabinet.

Power accounts for one of the largest costs in a data center. By some accounts, energy usage by IT equipment is estimated to be between 60% and 80% of the total IT energy loads (equipment and infrastructure loads).[1] According to a recent USA EPA report, the nation's data centers consume an estimated 61 billion kilowatt-hours (kWh) (1.5 percent of total U.S. electricity consumption)[2] in 2006. This equates to roughly $4.5 billion in total electrical costs. At its current rate, it is estimated that this figure will double in five years. When looking at power, it is not only the power needed by the servers and other equipment in the cabinet, but also the infrastructure to support this equipment (e.g., power delivery systems, PDUs, UPS and cooling systems).

Given these figures, there is considerable attention being paid to finding solutions to monitor, trend, control and ultimately reduce power consumption and the environmental impact of data centers. More importantly, these reductions equate to cost savings.

This paper discusses the different aspects to consider during the planning phase for the implementation of power distribution at the cabinet level. Effective planning will yield considerable benefits and provide tools and information to aid in the management of the data center. Additionally, this will help avoid problems down the road such as insufficient power to fully populate the cabinet, having enough capacity and outlets for redundancy, avoiding the need to run additional power drops in the future and not having enough poles on the distribution panel to provide it.

Number of outlets and type

Before deciding on what type of electrical circuits are needed to deliver power to the rack, it is necessary to determine what the power requirements are. In North America, data centers have a mix of equipment that runs on 120 V and 208 V single-phase power. In most cases 120 V appliances are designed with automatic switching power supplies that will accept an input voltage range between 100 V to 250 V. Equipment manufacturers design products with this consideration to reduce the number of variations so that the same product may be shipped around the world with just a country or region specific power cord. To convert these devices to 208 V power, it may be a simple matter of replacing the input power cord. An additional benefit of switching to 208 V power is that there is typically a 2% to 3% efficiency gain. Running devices at higher voltages reduces current draw resulting in lower losses and heat generation, which reduces the amount of cooling required to dissipate that heat.

The two most common plug and outlet standards in the IT industry are the ones established by the NEMA (National Electrical Manufacturers Association) and the IEC (International Electrotechnical Commission). In North America, NEMA connectors are quite common; however, IEC connectors are gaining popularity as 208 V circuits become the standard in data centers. Internationally IEC connectors are primarily used. Figures 1 and 2 show the common NEMA connectors for 120 V and 208 V, 15 A, 20 A and 30 A, single-phase and 3-phase circuits.

NEMA Nomenclature Explanation (Example L5-20R)	**L**	**5**	**-**	**20**	**R**
L designates twist-lock connections Blank means standard plug					
Voltage Rating 5 rated up to 125 V 6 rated up to 250 V 15 rated up to 250 V, 3-phase, Delta (3 pole, ground) 21 rated up to 120/208 V, 3-phase, Wye (3 pole, neutral, ground)					
Current Rating 15, 20, 50 or 60 following the dash represents the current rating in amps					
P or R represents Plug or Receptacle/Outlet					

Table 1 – NEMA style power plugs and outlets

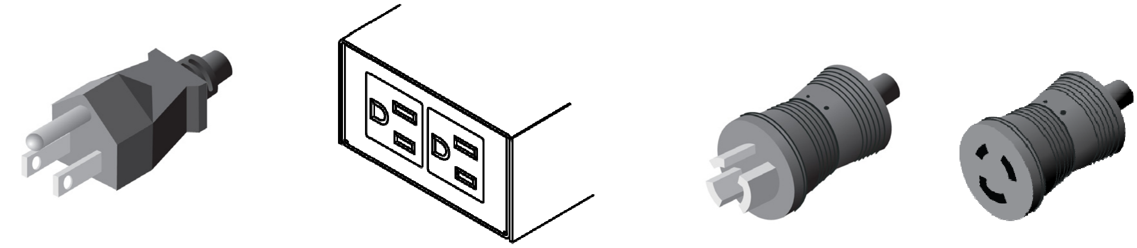

NEMA 5-15P
120 V, 15 A Plug

NEMA 5-15R
120 V, 15 A Outlets

NEMA L5-20P
Twist-lock, 120 V,
20 A Plug

NEMA L5-20R
Twist-lock, 120 V,
20 A Outlet

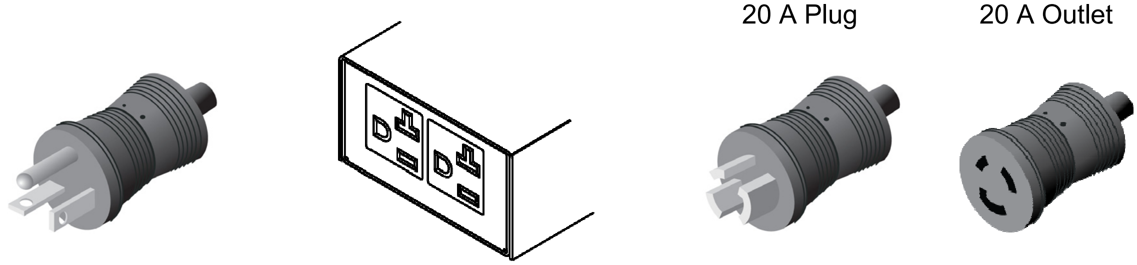

NEMA 5-20P
120 V, 20 A Plug

NEMA 5-20R
120 V, 20 A Outlets
(Convenience type accepts
NEMA 5-15P or 5-20P)

NEMA L5-30P
Twist-lock 120 V,
30 A Plug

NEMA L5-30R
Twist-lock, 120 V,
30 A Outlet

Figure 1 – Common 120 V NEMA style connectors for North America

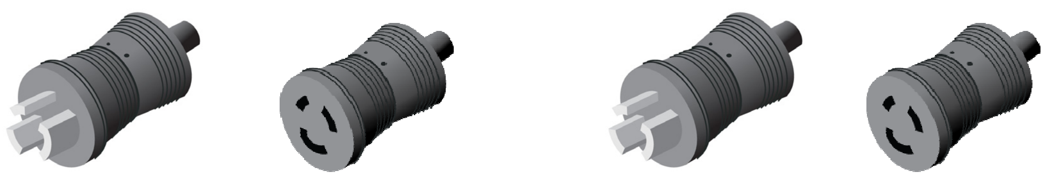

NEMA L6-20P
Twist-Lock, 208 V,
20 A Plug

NEMA L6-20R
Twist-Lock, 208 V,
20 A Outlet

NEMA L6-30P
Twist-Lock, 208 V,
30 A Plug

NEMA L6-30R
Twist-Lock, 208 V,
30 A Outlet

Figure 2 – Common 208 V NEMA style connectors for North America

IT equipment may be supplied with a localized or regionalized power cord that may not be compatible with the cabinet level power distribution unit used to distribute power to the cabinet. Fortunately, most IT equipment manufacturers have standardized on IEC connectors on their products such as the Dell 2161 DS KVM Switch, which is supplied with an IEC 60320 C14 inlet. This allows the power cord to be changed out. See Figure 3.

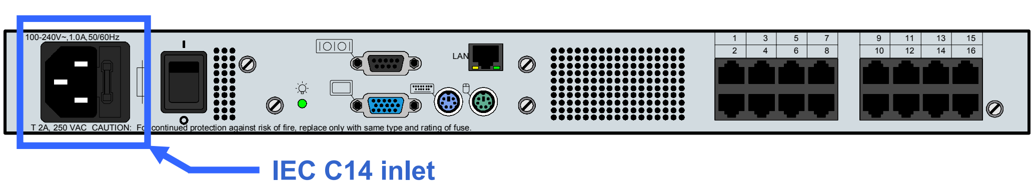

Figure 3 – Dell 2161 DS KVM switch

So if a device with an automatic switching power supply is supplied with a 120 V NEMA 5-15P power cord, it may be replaced with a 208 V power cord utilizing an IEC connection. The most common IEC plugs and outlets are the IEC 60320 C13 (outlet) / C14 (inlet) and the IEC 60320 C19 (outlet) / C20 (inlet) as shown in Figure 4. The IEC connectors are becoming the industry standard globally with many PDU manufacturers offering these types of outlets on their products. For 208 V applications in North America, utilizing IEC connectors for power distribution is more practical and provides greater outlet density than designing PDUs with twist-lock NEMA plugs.

IEC plug and outlet nomenclature

- Odd numbers are used to designate Outlets/Receptacles (C13 and C19)
- Even numbers indicate Inlets/Plugs (C14 and C20)
- Outlet/Receptacle and Inlet/Plug pairs are grouped in numerical sequence.
  For example:
  A C19 Outlet/Receptacle accepts a C20 Inlet/Plug
  A C13 Outlet/Receptacle accepts a C14 Inlet/Plug

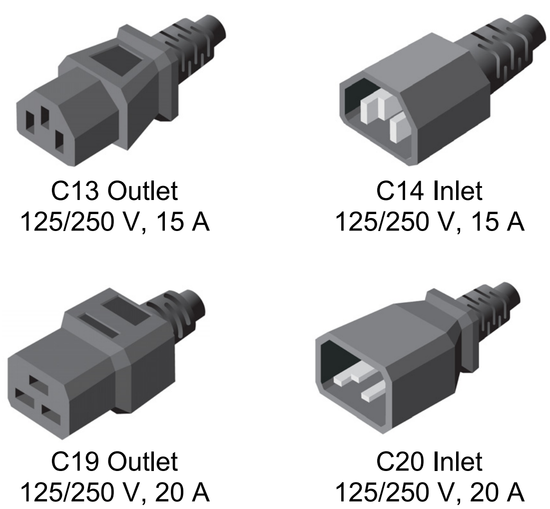

Figure 4 – Common IEC 60320 plugs and outlets

Internationally (and less common in the United States), the IEC 60309 pin-and-sleeve connectors are used on the power input feed of PDUs and various IT equipment. The plug is designed with several features so that it cannot be improperly inserted into the wrong outlet type. The plug's diameter is based on the load rating, the higher the current, the larger the diameter. An alignment notch prevents the plug from being improperly inserted into a socket. Additionally, an oversized ground pin prevents it from being inserted into a line (live) terminal. Figure 5 is an example of a 16 A, 250 V pin-and-sleeve plug. Common IEC 60309 250 V connectors are 16 A USA and International; 30 A (32 A) and 60 A (63 A) USA (International).

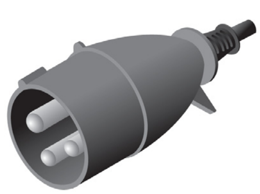

Figure 5 – IEC 60309 pin-and-sleeve plug

Power Cord Assemblies and Jumper Cables

There are two types of power cords. The first type is a Cord Assembly (e.g. NEMA L6-20P plug to an IEC C19 outlet) that connects a device or PDU directly to utility power coming from a centrally located distribution panel. The second type is a Jumper Cable (e.g. IEC C14 to a C13) that connects an IT device to a cabinet power distribution unit (PDU).

Most IT equipment is available with detachable cords, however some power supplies are designed with hardwired 30 A twist-lock power cords. Avoid these configurations if possible since this will limit the available options for PDUs with remote power management and the use of PDUs. With a few exceptions, most IT equipment are connected via Jumper Cables to a PDU rather than through a Cord Assembly directly to the distribution panel.

To help the cabinet to cool more efficiently, Jumper Cables should be routed properly to allow for optimum airflow and its lengths sized accordingly. An IT device may be supplied with a detachable power cord with a regionalized plug. This cord assembly may be replaced with a jumper cable since the device will be plugged into a PDU that will have either IEC C13 or C19, or NEMA 5-20R outlets.

Electrical circuits

After determining which IT equipment supports 208/230 V, which most do, the next step is to establish what the power requirements are for each device. The power ratings found on the product nameplates are notoriously high when compared to its actual power consumption. The reason for this is that manufacturers design power supplies to be used across multiple product lines and also in anticipation of additional features and upgrades in the future. Best practices are to check the equipment manufacturers' website to see if they have real-world power loads and capacities available. See Appendix A for a listing of power calculator tools available from Cisco, Dell, HP, IBM and Sun.

For example, an Hewlett Packard's Proliant® DL380 server power supply has a maximum output rating of 1000 W at 208 V. Using HP's Active Answers Power Calculator, a dual processor server running at 3.73 GHz with 8 GB RAM and two 60 GB SATA drives, only requires 567 W. The actual power requirements are much lower than the power supply ratings. If a server does not have this information available, a factor of 0.6 to 0.7 can usually be applied to the nameplate rating as a means of calculating typical power usage.

Data center managers need to be careful in anticipating how much power to supply to the rack since power demands are continually increasing. According to a recent report by Jonathan Koomey, staff scientist at Lawrence Berkeley National Laboratory, electricity usage in the data center is growing at an annual rate of 14% in the U.S. and 16% for the world.[1] Although each generation of processors consume less power per computational unit, growth in power consumption is attributed to increases in the number of new servers deployed.

Facilities managers need to decide whether to run single-phase or three-phase power circuits to the cabinets. Consider this: one 3-phase, 208 V, 30 A circuit provides as much capacity as three 120 V, 30 A single-phase circuits. See Table 2. This reduces the number of cables running to the cabinet (one 4-wire or 5-wire whip versus three 3-wire whips), thus improving airflow; and costs (cabling, labor and operating). To establish the overall power requirements of the rack, managers need to determine: the total power draw of the equipment, a safety factor, extra headroom to allow for future upgrades of equipment and expansion, and considerations for redundancy.

Table 2 lists the different load capacities for 120 V, 208 V and 230 V, single and 3-phase circuits.

North American Power Circuits

Voltage (V)	Current (A)	Phases	Max Power (kW)	De-rated Power[1] (kW)
120	20	1	2.4	1.9
120	30	1	3.6	2.9
208	20	1[2]	4.2	3.3
208	30	1[2]	6.2	5.0
208	20	3	7.2	5.8
208	30	3	10.8	8.6
208	60	3	21.6	17.3

1. De-rated by 20% in accordance to NEC (USA)
2. Two 120 V phases with 120° separation (1/3 cycle offset)

International Power Circuits

Voltage (V)	Current (A)	Phases	Max Power (kW)
230	16	1	3.7
230	32	1	7.4
400	16	3	11.1
400	32	3	22.2

Table 2 – Load capacity

De-rating

When loading a circuit it is generally not recommended operating at the limits of the circuit rating for the PDU since this may trip internal circuit breakers or blow fuses. To avoid this, users should apply a de-rated factor. In the USA, the National Electric Code (NEC) specifies that all circuits operate at 80% or less of maximum capacity.

3-Phase Power

3-phase power is delivered to the data center facility via the power grid. In the past, this was distributed to the cabinet level as single phase power since densities and power requirements did not demand it. A single phase 120 V, 30 A drop will provide up to 3.6 kW of power or 6.2 kW for a 208 V, 30 A drop. As densities and the power requirements of servers increase, the options are to run new drops to the cabinets or consider implementing 3-phase circuits. There are several advantages to 3-phase power.

In North America, 208 V 3-phase power delivers 1.732 times more power than a single-phase 208 V circuit. While internationally, a 400 V 3-phase circuit delivers three times more power than a 230 V single-phase circuit. A well designed PDU is equipped with 3 load monitors on each phase that will help operators properly balance each phase of the circuit, thereby reducing the complications resulting from an unbalanced circuit. With more available power, 3-phase circuits will make it easier to add equipment and have available capacity for future upgrades. Additionally, the higher capacity per drop reduces the number of drops needed per cabinet which will help improve airflow.

North American 3-phase 20 A circuits only delivers 34.6 A at 208 V, a little more than a single-phase 208 V 30 A circuit. Therefore a 30 A 3-phase circuit is more desirable since it can deliver up to 51.96 A at 208 V with similar costs to implement.

Data centers in North America that have 3-phase power have an option of a Delta or Wye supply. If all the equipment in the cabinet operates at 208 V, the Delta configuration (3 poles and 1 ground) is typically used. 208 V power is delivered by connecting two transformer windings together (phase to phase), see Figure 6. To reduce harmonics and limit the additive effect of reverse currents, power draw should be balanced across all three phases.

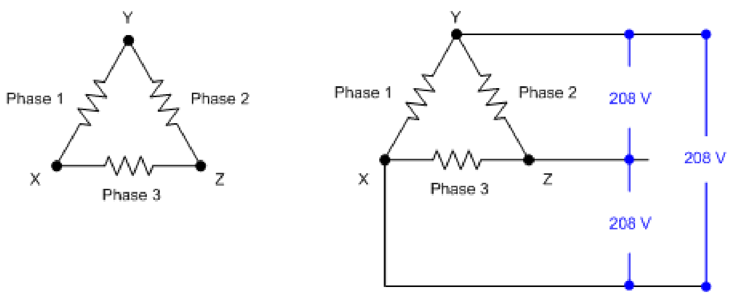

Figure 6 – North American 3-phase Delta power

If a rack has a mix of 208 V and 120 V devices, then a Wye configuration (3 pole, 1 neutral and 1 ground) is used because the PDU can be configured to supply both voltages, Figure 7. To supply 120 V, an outlet is configured with one phase and neutral. To deliver 208 V, an outlet is configured phase to phase.

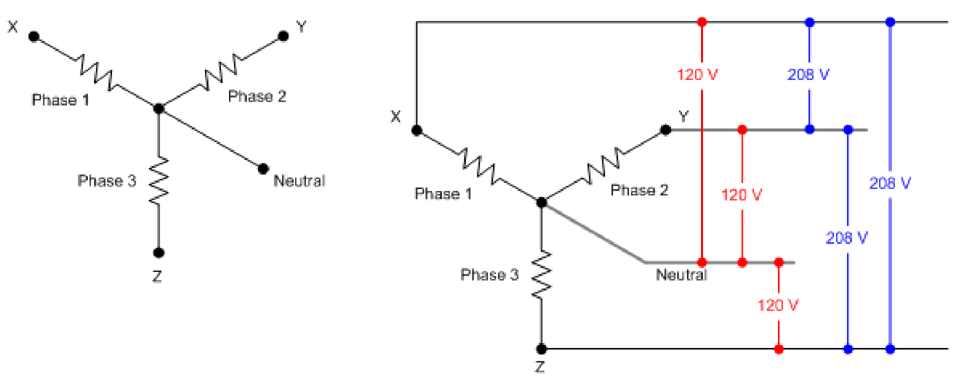

Figure 7 – North American 3-phase Wye power

Outside North America, a 3-phase circuit provides 230 V from phase to neutral so a Wye configuration is needed. Connecting any one of the 3 phases to neutral will provide 230 V. See Figure 8.

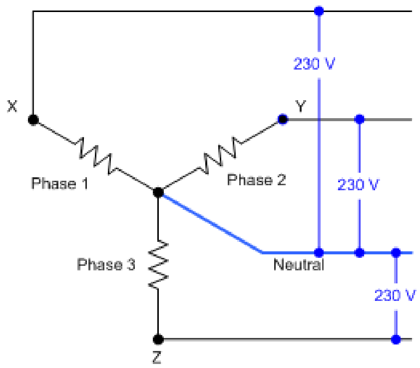

Figure 8 – International 400 V, 3-Phase Wye power

Redundancy

To provide redundancy, most IT equipment is designed with two or more power supplies. Ideally, each power supply draws power from a different cabinet PDU. Each cabinet PDU is fed by a separate power source. Typically most multi-supply devices, under normal operating conditions, have each power source deliver roughly 50% of the load to the server, so that if one fails or power is lost the other can still source 100% of the load.

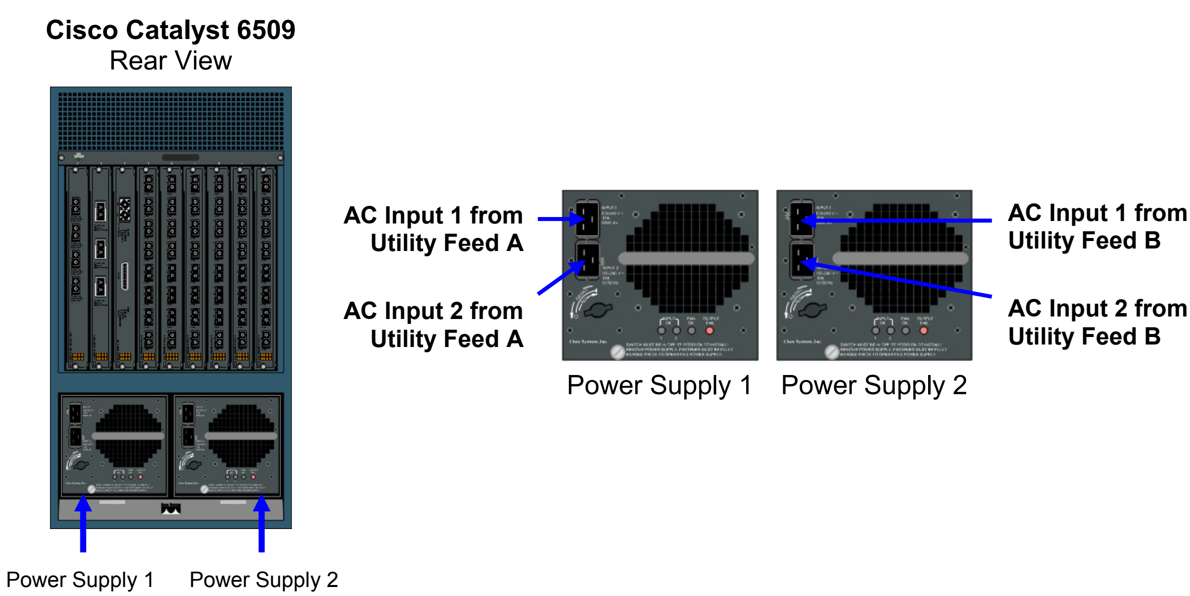

Figure 9 – The Cisco® Catalyst® Switch may be configured with two power supplies. Each supply is capable of powering the switch if the other fails. Additionally, each supply has two input power feeds that may be fed by separate and independent circuits.

One critical aspect of redundancy that should not be overlooked is how each cabinet PDU is loaded. Each circuit should be designed so that it is capable of handling the entire load of the rack; therefore, in a Tier 3 or 4, N + 1 system (as defined by the Uptime Institute), each PDU under normal operating conditions should not be loaded to more than 40% of the circuit capacity. Why 40%? In the USA, National Electric Code (NEC) requires that a circuit be loaded to no more than 80% of capacity. Since each rack is typically powered by two cabinet PDUs (A and B input power feeds), each PDU should be loaded no more than 40% because if one circuit is lost, the other PDU can source the entire load and remain at 80% of its capacity.

The following example illustrates how two PDUs are properly loaded under normal operating conditions but still fail when one power input feed is lost.

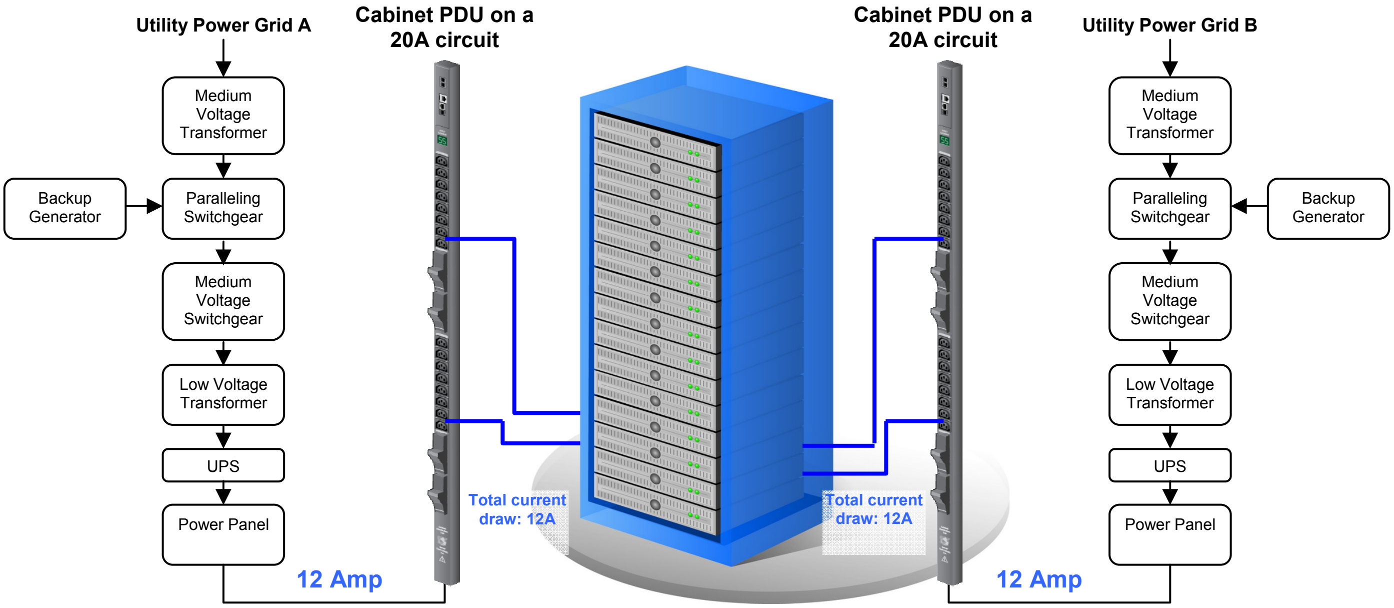

Figure 10 – Normal operation but improperly loaded

Each cabinet PDU is rated and capable of supplying 20 amps (Fig. 10). However, in the USA, the NEC requires that the circuit be loaded no more than 80% of its capacity or 16 A (derated) for this 20 A PDU. In this example each A and B circuit is feeding one input of a rack of dual power supply servers. The total load for each circuit is 12 A which is less than the derated value however when Circuit-A is lost (Fig. 11), the full load is drawn from Circuit-B resulting in a total current draw of 24 A. This will cause Circuit-B to blow a fuse or trip a circuit breaker leading to the loss of the whole rack of servers.

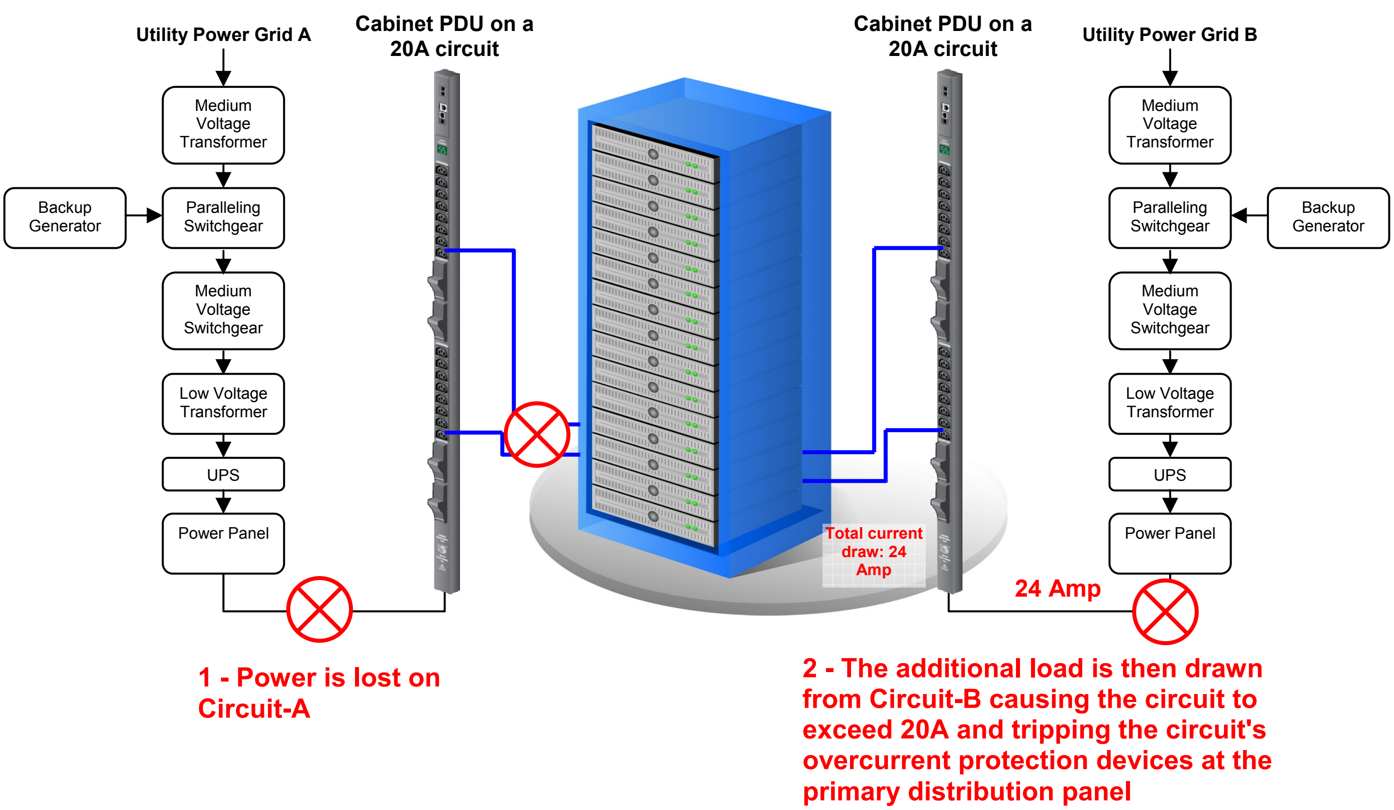

Figure 11 – A fault in Circuit-A will result in the loss of Circuit-B

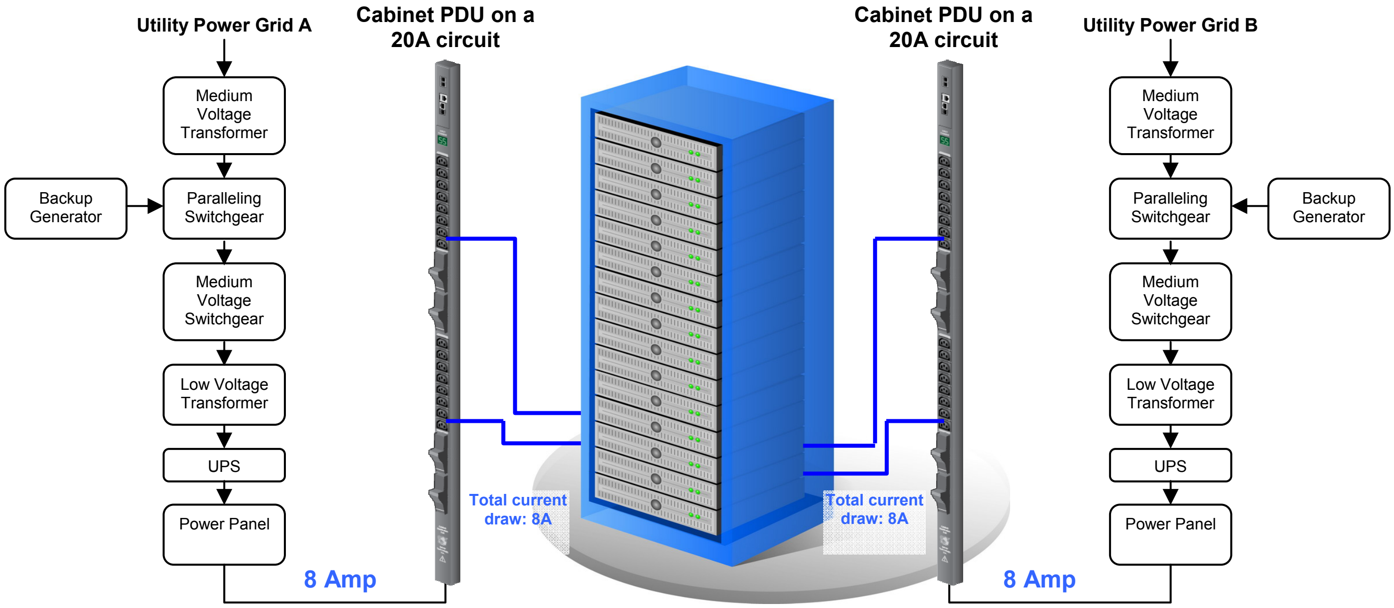

Figure 12 – Normal operation and properly loaded

In this example (Fig. 12), each 20 A circuit is loaded for a total of 8 A. When Circuit-A is lost (Fig. 13), the total load on Circuit-B is 16 A which will not cause a secondary failure. To provide power redundancy for critical systems, the total power capacity at the rack level will need to be at least two times the amount of power consumed by the devices.

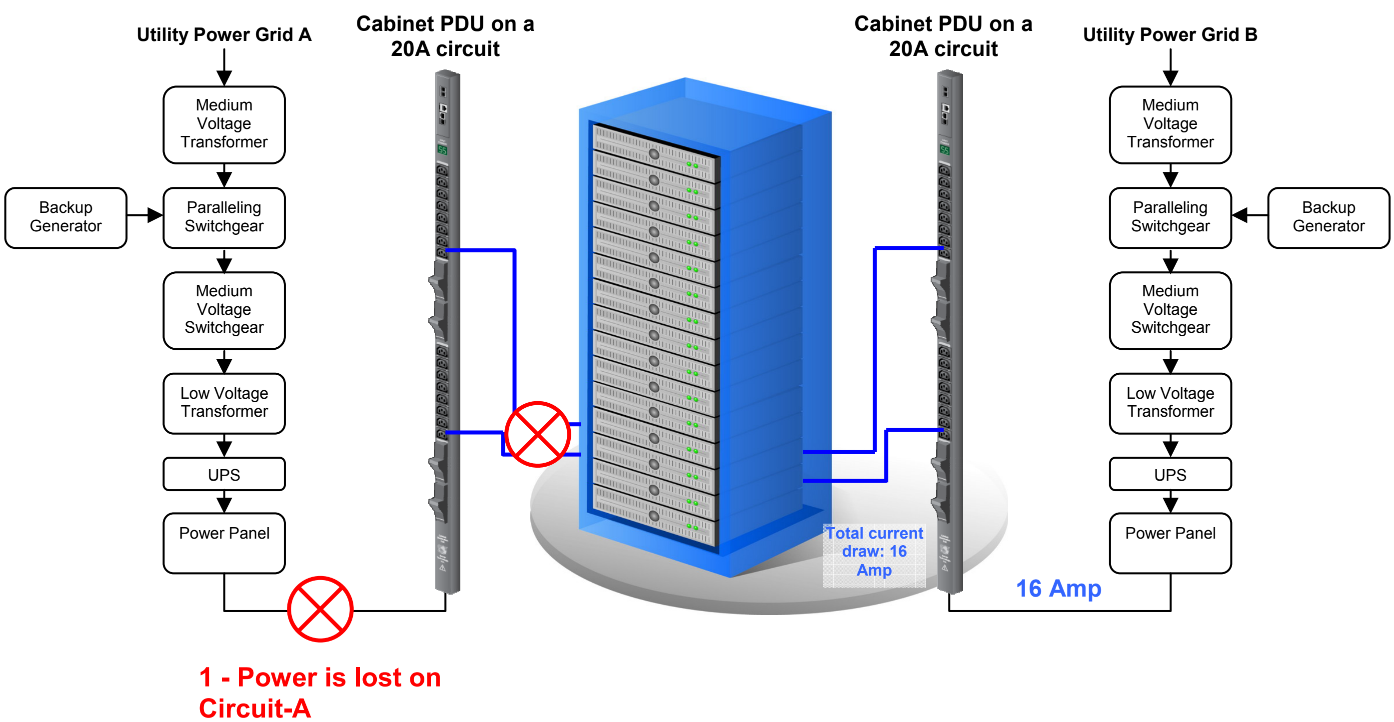

Figure 13 – Circuit-B is capable of handling the entire power load and all the equipment in the rack remains operational

Branch Circuit Protection

PDUs designed for medium to high density applications for use in North America shall have branch circuit protection (as a requirement of the National Electric Code) and more importantly have selective coordination. PDUs fed by 30 A or higher circuits are required to have branch circuit protection. In an overcurrent situation, a well engineered PDU design will isolate the faulted branch circuit while maintaining uninterrupted power to the rest of the branches, thereby reducing the number of outlets that are affected. Figure 14 shows an example of a PDU with fused branch circuit protection. Fuses in Branch A are blown when the total current for Branch A exceeds the fuse's current rating for a defined period of time. This removes power to Branch A while leaving Branch B unaffected. In this example, selective coordination is achieved with all the connected devices on Branch B still operating by not tripping the upstream circuit breaker at the power distribution panel board. On a PDU with circuit breakers, selective coordination is difficult to achieve without performing a coordination study.

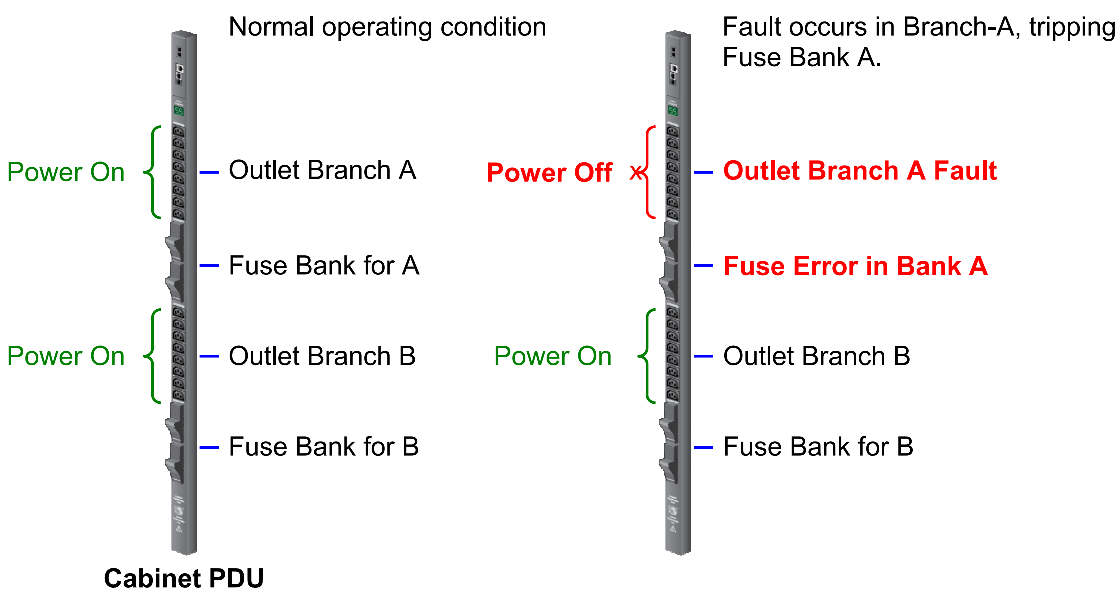

Figure 14 – Branch circuit protection

Addressing Single Corded Devices and Redundancy

In the data center today, there are still installations taking place with single corded devices due to legacy and cost issues. Considerations should be taken for these devices. Automatic transfer switches (ATS) and fail-safe transfer switches (FSTS) take in two power input feeds and distribute power to single corded devices. If there is a power interruption to one input power feed, the ATS or FSTS transfers the load from one input feed to the other.

Some features to consider when looking for an ATS or FSTS are:

- The ability of the product to distribute power from both input power feeds (known as infeed sharing). This reduces the load that needs to be transferred in a blackout or brownout situation. This design feature helps prolong the life of the products because only half of the total load is transferred resulting in less wear on the switching components.

- The ability of the ATS/FSTS to transfer loads from two power supplies that are not phase synchronized. Transfer switches that do not have this design feature will arc as the load is being transferred. This damages the relay contacts within the ATS/FSTS. In extreme circumstances, the relays fuse and the ATS/FSTS no longer functions.

Remote Power Management

The smart or intelligent PDUs are a powerful tool for administrators to securely manage their networks by providing the capability to monitor and control cabinet power remotely for a data center or branch office. PDUs may be remotely configured and managed while providing power and environmental monitoring reducing the costs associated with having someone onsite 24/7 or having to travel to a location during off hours. Remote management allows administrators to reboot a single outlet or a group of outlets for servers with multiple power input feeds and bring them up simultaneously to avoid inrush current issues. Additionally, PDUs with SNMP support provide alerts when thresholds are exceeded or fault conditions occur along with email alerts as well.

Another aspect to consider is out-of-band management. Does the PDU offer a way to access the device if the network is down to reboot a router or switch? A PDU equipped with a console port will allow administrators to access the PDU through a command line interface (CLI) or external modem.

With energy costs rising and a greater push to reduce power consumption, companies are looking at means to monitor usage. One option is to monitor power usage through the PDU and correlate this to CPU utilization and other metrics as a means to identify under utilized equipment, available capacity and cost savings resulting from consolidation or virtualization.

Summary

Proper planning and implementation of a rack level power distribution strategy provides for a more robust network and greater uptime. It starts by evaluating the power needs of the devices and upgrading 120 V devices to 208 V or converting units with automatic switching power supplies to 208 V. The next step is understanding and designing the appropriate types of power drops with enough capacity for redundancy for both single-corded and devices with multiple power supplies. Anticipating future power demands will reduce the need to run additional power drops, so consider providing 3-phase circuits. Equally important is how to remotely manage these devices. Smart and Switched (IP addressable) PDUs have the ability to provide load status, environmental readings, and individual outlet control to reboot locked up servers remotely, either in-band or out-of-band. These features are ever more clearly needed when that pager goes off at 2 A.M.

Appendix A:
Power Calculation Tools

Cisco
http://tools.cisco.com/cpc/

Dell
http://www.dell.com/calc

Hewlett Packard
http://www.hp.com/go/bladesystem/powercalculator

IBM
http://www.ibm.com/systems/bladecenter/powerconfig

Sun
Power Calculators are available but not in one centralized location. At www.sun.com use their search feature and query for "Power Calculator"

Sources:

1. Estimating Total Power Consumption by Servers in the U.S. and the World, Jonathan G. Koomey, Ph.D., Staff Scientist, Lawrence Berkeley National Laboratory and Stanford University. Sponsored by Advanced Micro Devices, February 15, 2007.
2. Report to Congress on Server and Data Center Energy Efficiency, U.S. EPA, August 2, 2007.