Powering the data center efficiently

best practices

Abstract

In today's rapidly evolving information technology (IT) environment, facility power planning is becoming more critical. This paper will explain trends in power requirements, and it will explain basic considerations of data center power. The paper will describe methods for calculating power requirements and how to use those requirements to select an appropriate power distribution configuration. This paper will also communicate best practices for using certain power products in conjunction with ProLiant servers and rack products.

Introduction

As server density grows, power requirements continue to increase. Today's data centers can push power requirements to as much as 300 watts (W) per square foot. A typical 42U[1] cabinet with servers and storage requires 5 to 8 kilowatts (kW) of load, and today's high-density servers and blades can push this to 12 to 14 kW and even peak as high as 20 kW.

In addition, cable management is becoming a key issue, driven by high-density servers with redundant power supplies. High-density, 1-U servers with redundant power supplies can require as many as 84 power cords in a single rack. As server power consumption continues to rise, power distribution unit (PDU) designs will also track higher current demands of redundantly powered servers.

To support this rapidly evolving environment, facility power distribution planning requirements are becoming more critical. However, determining true power needs is often confusing. A thorough understanding of the following will reduce the confusion:

- Facility planning
- Basic power considerations
- Details of typical power components and their installation
- Power planning
- Cable management solutions

Facility planning

Successful installations and implementations require careful, coordinated planning. Customers should ask themselves the following questions about their data center:

- Have I engaged facility engineering in planning for hardware deployments and future growth?
- Have I adequately planned for growth in power and heat density?
- Is my data center configured to efficiently provide sufficient airflow to the new servers and storage products?

When facilities engineers are involved early and often in the expansion and build out of a data center, risk is minimized and success is maximized. Site preparation has been a standard practice in the mainframe and storage solution environments, now, PC server data centers need to adopt it as standard practice. Although customers own and manage their environments, HP has professional services that can help configure their environment to meet their needs and to operate as efficiently as possible.

---

[1] A U is a standard unit of measure for designating the height in computer enclosures and rack cabinets. A U equals 1.75 inches. For example, a 4U chassis is 7 inches high. A 40U rack cabinet is 70 inches high.

Basic power considerations

When planning an installation, customers must make a few decisions as well as be aware of a few known power supply issues regarding the delivery of power.

Server power supply issues

Input voltage, power factor, inrush current, and leakage current are power supply issues that must be considered to provide safe and efficient power to data center equipment. As a standard feature, all ProLiant power supplies are power-factor-corrected. This eliminates one of these issues and provides for efficient operation.

Input voltage selection

While this is not a choice in much of the world, the selection of the proper input voltage is a very important issue in the Americas and other regions around the world that follow the North American commercial wiring practices. Choosing between low-voltage (LV) operation and high-voltage (HV) operation will have an effect on these areas: power supply output capacity, power conversion efficiency, power supply thermal operation, and finally, power supply reliability.

HP defines LV as 100 to 120 VAC and HV as 200 to 240 VAC. In North America, commercial power is delivered to the server at 120 V or 208 V. All HP servers now have auto-sensing input circuitry that automatically adjusts to the applied input voltage. The only exceptions are those devices that are defined as HV operation only. Make sure to always review a device's input specifications prior to connecting them to the power distribution system.

Some ProLiant servers are equipped with power supplies that have greater capacity when connected to HV input power. When heavily loaded, these products may require HV operation or the addition of another power supply in order to have fully redundant operation. The primary cause for the difference is the limitation of the input devices in the power supply. Components such as the input receptacle, input fuses, and input filter components are many times limited to 10 Amps of maximum input current. For a given output power, a power supply operating at LV effectively has double the input current when compared to HV operation. To achieve maximum output power at LV, the input current would exceed the rating of the input currents.

Power supplies operate more efficiently when operating at HV. The typical server switch-mode power supply has an efficiency rating between 65 percent and 80 percent. Some special-purpose products can hit 90 percent efficiency. Operating at LV causes the power supply to operate at the lower end of this range. At 65 percent efficiency, 35 out of every 100 W pulled from the utility will be wasted by the power supply. In a large data center, this lower efficiency can waste many kilowatt hours and drive electrical bills higher with no benefit. In fact, there are many more negative effects.

Power supply thermal operation is also effected by the choice of input voltage. The input components previously mentioned above actually run hotter when operating with LV input power. This is caused by the almost double input current. The formula for heat generated in a component is $I^2 \times R$, where I is the input current and R is the resistance of the component. Therefore, if the input current is doubled, the heat generated in any given component is going to be four times higher. The life of a component is significantly shortened when it continuously runs at higher temperatures. Therefore, there is an overall impact on the life expectancy of the supply.

Inrush current

Electronic devices containing solid-state power supplies experience an input current during initial start-up that can be several times greater than their operating current. High current surge during start-up, referred to as inrush current, can affect electrical systems by tripping fuses and circuit breakers unnecessarily.

Inrush current is determined using theoretical worst-case measurements. Server specifications designate the amount of inrush current that can be tolerated by the server. For example, the specification for the ProLiant DL380 G2 server is 50 amperes (A or amp) per power supply for a maximum of 2 milliseconds (ms). Typical inrush events that occur in the real world last far less than 2 ms. Inrush current of individual servers is not problematic; however, inrush current is additive and multi-server implementations can easily surpass 50 A, however, for a very short period of time.

Given the number of servers in a data center, power supplies must contain inrush current surge protection to limit the current drawn during startup. If the power supply does not use current surge protection, relays and fuses must be rated higher than any possible surge current. Server power supplies currently offered by HP provide the capability to limit active inrush current. Circuitry in the front end of the power supply limits the amount of current drawn by the power supply during the initial period when alternating current (AC) power is applied.

In addition, HP rack-mount PDU load segments are sized to minimize the effects of short-circuit current and overload current events within the rack. Under normal operating conditions, the load segment breakers should isolate fault currents and prevent them from taking down the entire load connected to the PDU. Breakers capable of withstanding hundreds of amps of inrush current would not provide the protection needed under normal conditions; all the fault current would be passed directly to the AC mains distribution system and could affect more than just the rack involved. However, the breakers will pass relatively large amounts of current for extremely short periods of time.

HP recommends two practices to mitigate the effects of inrush current:

- Install additional PDUs in each rack to distribute and dilute the effect of the inrush current when power is applied to the rack. Ensure that the circuit breakers feeding these PDUs are suitably rated for the duty they will see. In most of the European Union, HP recommends the use of Type D breakers.
- Use an orderly shut-down and start-up procedure that applies power only to a portion of the equipment in the rack when power is restored. Once the initial load is energized, energize the power to the remainder of the load in the rack. This procedure can be accomplished either by unplugging the power cords to a portion of the load or by opening one of the load segment breakers on the PDU.

Leakage current

Earth leakage current is created by filter capacitors located between the primary circuits and the primary grounding (earthing) conductor and subsequently the chassis of the computer. The current is measured from the accessible parts of the equipment back to the phase and neutral conductors. Under normal operating conditions, leakage current does not create a hazard. Because the current is additive when several pieces of equipment are connected together to the same source (for example, an uninterruptible power supply or PDU), the level of leakage current can reach a hazardous potential quickly. If the primary ground or earthing conductor becomes open for any reason, the leakage current and all of its potential will become available on any conductive (metal) surface of the equipment. If an individual comes in contact with the chassis of the equipment and ground, electric shock can occur.

Leakage current can be problematic when the building's installation is required to install residual ground current fault detectors. These devices are set to trip open a circuit breaker when the leakage current exceeds a predetermined level. The United Kingdom is one region where this can be an issue.

For proper operation and safety, the equipment must be properly grounded. In the United States, install the equipment in accordance with National Fire Protection Association (NFPA) 70-1999 National Electric Code (NEC) Article 250 as well as any local and regional building codes. In Canada, install the equipment in accordance with Canadian Standards Association, CSA C22.1, Canadian Electrical Code. In all other countries, install the equipment according to any regional or national electrical wiring codes such as the International Electrotechnical Commission (IEC) 364 parts 1 through 7. Furthermore, ensure that all power distribution devices used in the installation—such as branch wiring and receptacles—are listed or certified grounding-type devices.

Because of the high ground-leakage currents associated with multiple servers connected to the same power source, a reliable grounded (earthed) connection is essential before applying power to the system. HP recommends the use of a PDU that is either permanently wired to the building's branch circuit or features a nondetachable cord that is wired to an industrial style plug. National Electrical Manufacturers Association (NEMA) locking-style plugs or those complying with IEC 60309 are considered suitable for this purpose.

Power factor correction

Before computing and storage devices can use electrical power, the AC provided from the source must be transformed to direct current (DC) by a power supply. The term "power" is the rate at which the electricity does work, such as running a central processing unit (CPU) or turning a cooling fan. The power that the electricity provides (apparent power) is simply the voltage times the current, measured in volt-amperes (VA).

There is a difference between the power supplied to a device and the power actually used by the device because of the capacitive nature at the input of the device to delay current flow. The true power used by the load is measured in watts.

The power factor (PF) of a device is a number between zero and one that represents the ratio between the apparent power in VA and the real power in Watts or, in other words, how efficiently the current is converted to useful work. A power supply that has a PF of 1.0 indicates that the voltage and current peak together, which results in the most efficient loading of the device.

Power supplies for servers usually contain circuitry to "correct" the power factor (that is, to bring input current and voltage into phase). Power-factor correction allows the input current to continuously flow, reduces the peak input current, and reduces the energy loss in the power supply, thus improving its operation efficiency. Power-factor-corrected (PFC) power supplies have a power factor near unity (~1), and thus are highly efficient. The use of energy-efficient PFC devices, including uninterruptible power supplies (UPSs), can lead to significant cost savings.

Delivering power

As ProLiant servers have evolved from their origin as personal computer (PC) servers to data center servers, the long-held assumption that a 120-V, 10-A power cord should be sufficient to power anything in the data center has become invalid.

Also invalid is the assumption that an open receptacle indicates available unlimited power—like blank checks in a checkbook. Just as blank checks do not indicate unlimited money supply, open receptacles do not indicate unlimited power supply. For example, a 24-A PDU with 32 receptacles can power only three ProLiant DL580 G2 servers because each requires 8 A.

To provide more power to a rack efficiently, the paradigm of using 15- to 20-A branches must be broken. In North America, this is the standard branch circuit that an electrician will run off the main distribution board. Some options include 30-A branches, which are becoming standard, and three-phase branches, which are coming very rapidly.

Hardware

More power requires larger hardware—plugs and wires have to change from the PC line cord. Figure 1 shows two plugs common in HP products, and Figure 2 shows the larger wiring of these plugs.

Figure 1. Common plugs.

NEMA L6-30P,
208 V, 30 A, 3 w

NEMA L15-30P,
208 V, 30 A 3 ph, 4 w

Figure 2. Large wiring in common plugs.

Larger copper cabling
SJ/SO 3 x 10awg
SJ/SO 4 x 10awg

The industry uses two primary connector standards: NEMA and IEC.

North America uses the NEMA standard. The L in front indicates a locking or twist-lock connection. The absence of L indicates a standard push-in plug. The 5 or 6 following that is the voltage rating of the plug. A 5 indicates a rating up to 125 V; a 6 indicates a rating up to 250 V. The 15, 20, or 30 is the rated current in amperes. P and R stand for plug and receptacle, respectively. In the case of the top common plug shown in Figure 1, the designation NEMA L6-30P indicates a locking plug rated for up to 250 V and 30 A.

The other primary standard is the IEC standard. The IEC is a standards body in Geneva, Switzerland that defines the most common connectors: the IEC 320 general-purpose household connectors and the IEC 309 industrial-grade connectors. The C13/C14 connectors are the 10-A power supply connectors used on 90 percent of HP equipment. With all IEC 320 connectors, females (receptacles) have odd numbers; males (plugs) have even numbers. The C19/C20 connectors are rated for 16 A. Larger or higher power devices may require C19/C20 connectors if the input can exceed the 10-A range. One example is the ProLiant DL580 G2 server.

Inside a rack, only two connectors are used for distributing power: Those are the NEMA 5-15 for low-voltage, 120-V, applications, and the IEC 320, C13/C14.

IEC 309 pin-and-sleeve connectors are commonly used on international PDUs, servers, and storage devices, and they are beginning to be used in North America. Figure 3 shows three types of IEC 309 pin-and-sleeve connectors, all rated at 250 V, that are becoming more popular in the United States:

- 16-A single-phase
- 32-A single-phase
- 32-A three-phase

Figure 3. IEC 309 pin-and-sleeve connectors.

Power density

From an electrical supply standpoint, deploying full racks of servers can be problematic. The total rack load is not the problem. Rather, the issue lies with how power is distributed—the relationship between and restrictions on the line cords and distribution outlet. For example, the 250-V rating on a plug is just its maximum voltage, not the rating of the circuit behind it. Table 1 shows power limitations associated with each common type of branch circuit in North America.

Table 1. Commercial branch circuits in North America.

Branch voltage	Branch circuit size	Max load 80% per NEC	Maximum available power
120 V	20 A	16 A	1920 VA
120 V	30 A	24 A	2880 VA
208 V	20 A	16 A	3328 VA
208 V	30 A	24 A	4992 VA
208 V	20 A 3∅	27 A	5760 VA
208 V	30 A 3∅	42 A	8650 VA

In North America, the standard distribution to the data center is 208 V. Distributions of 220/221 V (sometimes also referred to as 230/240 V) do not occur in commercial power distribution in North America. Since everything runs off of 208 V, it is not necessary to install transformers to convert 208 V to 220 V.

As shown in Table 1, a 120-V, 30-A circuit is limited to less than 3,000 VA, or 3,000 watts of power. The same size wire at 208 V, 30 A can deliver almost 5,000 watts, or 5,000 VA. However, even 5,000 VA is becoming a limitation with racks that are very heavily loaded. Today's power densities dictate leaving the 120-V infrastructure, using multiple 208-V, 30-A feeds and looking forward to using three-phase power to meet the density demands.

Distribution panel efficiency

The distribution panel is key—how power is pulled out of the panel determines how efficiently that power is used. In most cases, more power is left in the panel than is used on the floor.

The problem is the number of poles or attach points in the panel. A standard power distribution panel for a data center will provide approximately 150,000 VA with 84 poles. A 208-V distribution requires two poles, which allows 42 two-pole positions out of the distribution panel. As a result, power appears plentiful. In actuality, distribution limitations leave power in the panel.

For example, a cabinet full of 21 ProLiant DL380 G2 servers requires 8,560 VA to operate. A 24-A PDU at 208 V is limited to about 5,000 VA. Consequently, the load requires at least two 24-A PDUs. Redundancy requires four PDUs. Since each PDU requires two poles off the distribution panel to get 208 V, each cabinet requires four breakers and eight poles. With these requirements, the 84-pole panel will provide redundant power for 10 cabinets and 210 servers using today's distribution method. A 10-cabinet implementation will use only 86,000 VA of the 150,000 VA possible from the panel. Since there are not enough poles to pull additional power, more than 40 percent of the total available power is left at the panel.

Increasing power capacity

Installing a larger current-density PDU is not a solution for the problem. The lack of cost-effective plug standards above 30 A would require that larger single- and two-phase PDUs be hard wired into the data center by an electrician.

The only potential solution today is to move to three-phase power distribution. The NEMA L15-30P is a very common plug used for three-phase power.

Three-phase power

All electricity is generated and distributed by utilities as three-phase AC, usually at a voltage in excess of 150,000 V. By the time the power reaches a data center, its voltage has been stepped down to 480 V or less.

Three-phase current is represented by three separate single phases. The voltage of each phase is represented by a sinusoidal wave that alternates between positive and negative values (see Figure 4) at a frequency of 60 cycles per second (cps), or 50 cps in many European countries. The current in each phase has a similar waveform at the same frequency.

Figure 4. Representation of sinusoidal voltage waves of three-phase current and corresponding transformer phase windings.

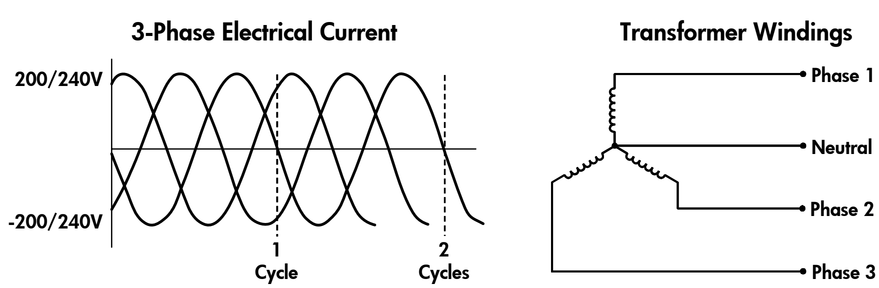

Most computer equipment operates on single-phase current. Single-phase loads, such as computer equipment, are connected to one of the transformer's phase windings and the neutral connection. Equipment requiring more power, including data center environment support systems, runs on three-phase current. Three-phase loads, such as air-conditioning equipment, are connected to all three transformer windings.

Larger systems are moving to higher amperage or three-phase power. Many enterprise-class machines presently use three-phase power, and most data centers are already wired for three-phase power. A common and economical method of supplying power to high-density data centers is to use a 208-V, three-phase system, known as high-line AC power. The system operates with 208 V between any two transformer windings, giving 120 V to neutral (Figure 5).

Figure 5. High-line AC power

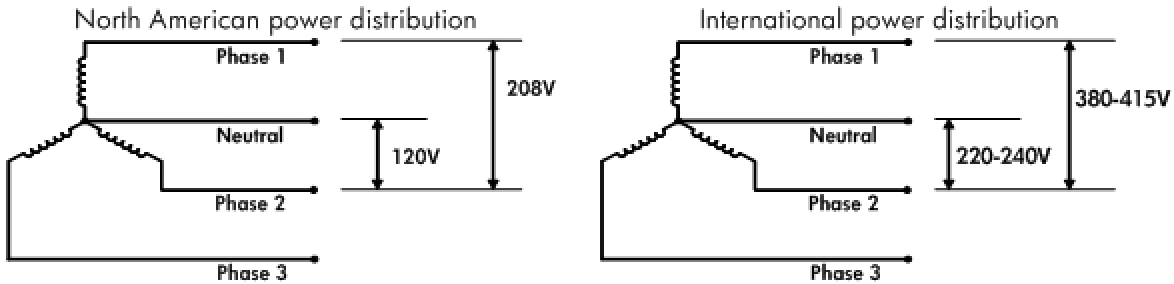

Three-phase power is typically more efficient than single-phase power since it provides more than 150 percent of maximum available power provided by single-phase power. In addition, three-phase power provides more power using smaller circuit sizes, which translates to smaller cables for customers. Since power density goes from 5,000 VA to 8,600 VA, 24 A to 42 A, fewer overall panel positions are used for rack implementations. Three-phase power is also more stable than single-phase power since it has three voltage waves and is at peak voltage more of the time. Table 2 compares single-phase, two-phase, and three-phase power distributions and illustrates the increased efficiency of three-phase power.

Table 2. Panel utilization for racks of ProLiant DL380 G2 servers.

PDU model	24A 1-phase PDU, 120 V	24A 2-phase PDU, 208 V	24A 3-phase PDU, 208 V
VA per PDU	2880	4992	8650
PDUs per rack of 21 ProLiant DL380 G2 servers	8	4	2
Breaker poles required per rack	8 (1 poles per PDU)	8 (2 poles per PDU)	6 (3 poles per PDU)
Racks per 84-pole distribution panel	10	10	14
Total VA used from panel	89,460	86,730	121,422

While the best setup for a data center application depends upon the particular facility requirements, HP recommends that rack-optimized equipment be installed for operating at high-line voltage for the following reasons:

- A single cabinet can provide greater capacity. For the same size circuit, almost twice the power can be delivered to a rack at high-line power than at low-line power. For example, a 120-V, 30-A branch circuit can deliver 2,880 VA to a rack, while a 208-V, 30-A branch can deliver 4,992 VA to a rack.
- The amount of infrastructure (for example, the number of PDUs and the number of circuits pulled from the panel) can be halved.
- Some products require 200-V to 240-V input power to operate at their full rated capacity.
- Power supply efficiency is higher at higher voltage operation, which wastes less energy. Power factor is also higher with higher voltage input.
- Power supplies run cooler at higher input voltages and therefore will last longer and improve overall availability.

Typical components

A well-designed electrical system for the data center ensures adequate and consistent power to the computer hardware and reduces the risk of failures at every point in the system. Adequate and consistent power can be provided with a combination of components, including:

- Racks
- Power distribution units
- Circuit breakers

Racks

Racks range from 70 to 87 inches (1.8 to 2.2 meters) in height, which means they can hold from six 8U servers to forty-two 1U servers, not including storage or UPSs. A typical 42U cabinet with dual processor (2P) servers and storage requires over 12 kW of power. IA64 servers and higher capacity drives may increase the power requirement to 15 kW per rack.

Providing adequate and consistent power to racks of equipment can be challenging. For example, cabling for a rack full of equipment can be troublesome. HP has developed racks and rack options that help to alleviate many of the challenges associated with powering racks of servers. For example, many racks now provide extra room for cables and space for mounting 0U products like the 0U power distribution unit.

Power distribution units

In traditional data centers, power goes from the transformers to a number of sub-panels that contain circuit breakers. A wire from each circuit breaker provides power to each outlet where the equipment is connected. As long as the power provided is sufficient for the equipment in place, this system is adequate. However, if the outlets, wire, and breakers need to be upgraded for higher amperage or three-phase power, the process can be very expensive and time consuming. As a result, HP recommends using PDUs in installations where a number of server units can place serious loading demands on the AC power bus.

HP modular PDUs integrate the outlets, wire, and breakers in a convenient location on each rack of equipment. HP modular PDUs range from 16 A to 40 A, with up to 32 outlet receptacles. For mission-critical environments with redundant power systems, HP also offers a fault-tolerant dual input PDU that automatically switches over to a secondary input source if the first source fails.

With so many components in every rack, cable management is becoming increasingly important. The HP modular PDU provides an excellent cable management solution. The PDU is split into two pieces:

- A control core connects to the power bus (or to a UPS), brings the high-amperage conductor into the rack, and splits it up across the breakers. The control core includes a 15-A circuit breaker for each of the four C19-type extension bar outputs and is available in low- and high-voltage versions.
- Extension bars bring the receptacles to the back of the rack where they can be easily accessed. The modular PDU comes with three types of extension bars to accommodate a variety of distribution requirements. The extension bars include both single-bar and double-bar mounting brackets for attachment to vertical rack supports.

With the extension bar in the back, shorter power cables can be used. Most HP servers ship with either an 8- or 12-foot power cord because not every customer buys HP modular PDUs for their units. For those who do, a cable option provides 4-foot IEC-to-IEC jump cords that go from the servers to the PDU. The 4-foot cables are long enough to go from the cable management arm and plug into the back of the PDU with no extra cable to bundle. They also allow pulling the server out on its rail in the locked position.

In most of today's HP servers, the redundant power supplies are both located on the right-hand side of the server. As a result it is necessary to get all the power distribution to the right-hand side of the rack. HP developed a new set of mounting brackets that allows installation of two power strips on the right-hand side of the cabinet. This makes better use of the vertical setup in the cabinet. The new brackets actually allow installing five of the double power strips on one side (Figure 6) as opposed to four single power strips with the old brackets. This provides a total of 80 outlets and is the setup required for racking 40 ProLiant DL360 G3 servers.

Figure 6. Double-mounted power strips.

The control core may be rack-mounted in a 0U or 1U configuration. A 0U configuration (Figure 7) may be preferable in a high-density installation requiring the maximum amount of vertical space for servers and other active components. The 0U configuration saves vertical rack space for equipment requiring more accessibility and provides easy access to power connections.

A 1U configuration (Figure 8) may be preferable in installations where operator or maintenance accessibility to all components is key. The 1U configuration provides easy access to all switches and circuit breakers, easy viewing of circuit status LEDs, and easy access for service replacement or upgrade.

Figure 7. Modular PDU in 0U configuration.

Figure 8. Modular PDU in 1U configuration.

Circuit breakers

A circuit breaker is a device that interrupts the path when necessary to protect other devices attached to the circuit—for example, in case of a power surge. HP recommends using circuit breakers with nominal current rating of at least a 32 A. In addition, for large data center applications, HP recommends use of Type D breakers for international customers and UL409 breakers for North American customers. A 32-A Type D breaker is rated at 640 A for up to 5 seconds and for more than 1,000 A at 10 ms.

Power planning

Facility cooling and power distribution planning requirements are becoming more critical as servers and data center components require more and more power. Determining actual power needs can be confusing.

The amount of power a facility will need can be calculated using one of three methods:

- Using nameplate ratings – This method worked in yesterday's environment, but it is a costly method resulting in wasted infrastructure dollars.
- Using actual power measurements – This is the most accurate approach, but numbers are difficult to generate and collect when in a planning mode.
- Using ProLiant Power Calculators – HP recommends this as the best practice for advanced planning since it produces more realistic numbers.

Nameplate ratings

Regulatory agencies, such as UL and TÜV, specify the amount of power required for each product. This power specification is called a nameplate rating. Nameplate ratings are the maximum amount of power used by a device.

A product may not always use the amount of power specified by its nameplate rating. For example, a typical server that is not fully populated with processors or disk drives will use 50 percent to 60 percent less power than its nameplate rating.

If a facility is planned based on nameplate ratings, the amount of power infrastructure provided could significantly exceed the actual power needed to operate the equipment. The result could be unnecessary expenditure.

For example, an actual installation using ProLiant DL360 servers was planned on the basis of the server nameplate rating of 292 W per server. However, actual measurements of a ProLiant DL360 server with two processors and two 18-gigabyte (GB) disk drives averaged 192 W. When that discrepancy is multiplied for a rack of 42 servers, the total amount of unused power is 8,064 W—a huge unnecessary expense.

Online calculator

HP provides power calculators to simplify facility planning. The HP power calculators provide information based on actual system measurements, not nameplate ratings. The measurements are taken using production servers with real processors, real hard drives, real memory, and actual operating systems running internally developed exercise utilities. All major system components (processors, memory, and disk drives) are exercised at 100 percent duty cycle. Calculator results may be higher than actual usage, but they are much closer than the nameplate ratings and provide extra headroom.

Figure 9 shows a sample screen shot of the power calculators, which can be found at http://h18001.www1.hp.com/partners/microsoft/utilities/power.html.

Figure 9. Screen shot of ProLiant DL380 G3 server power calculator.

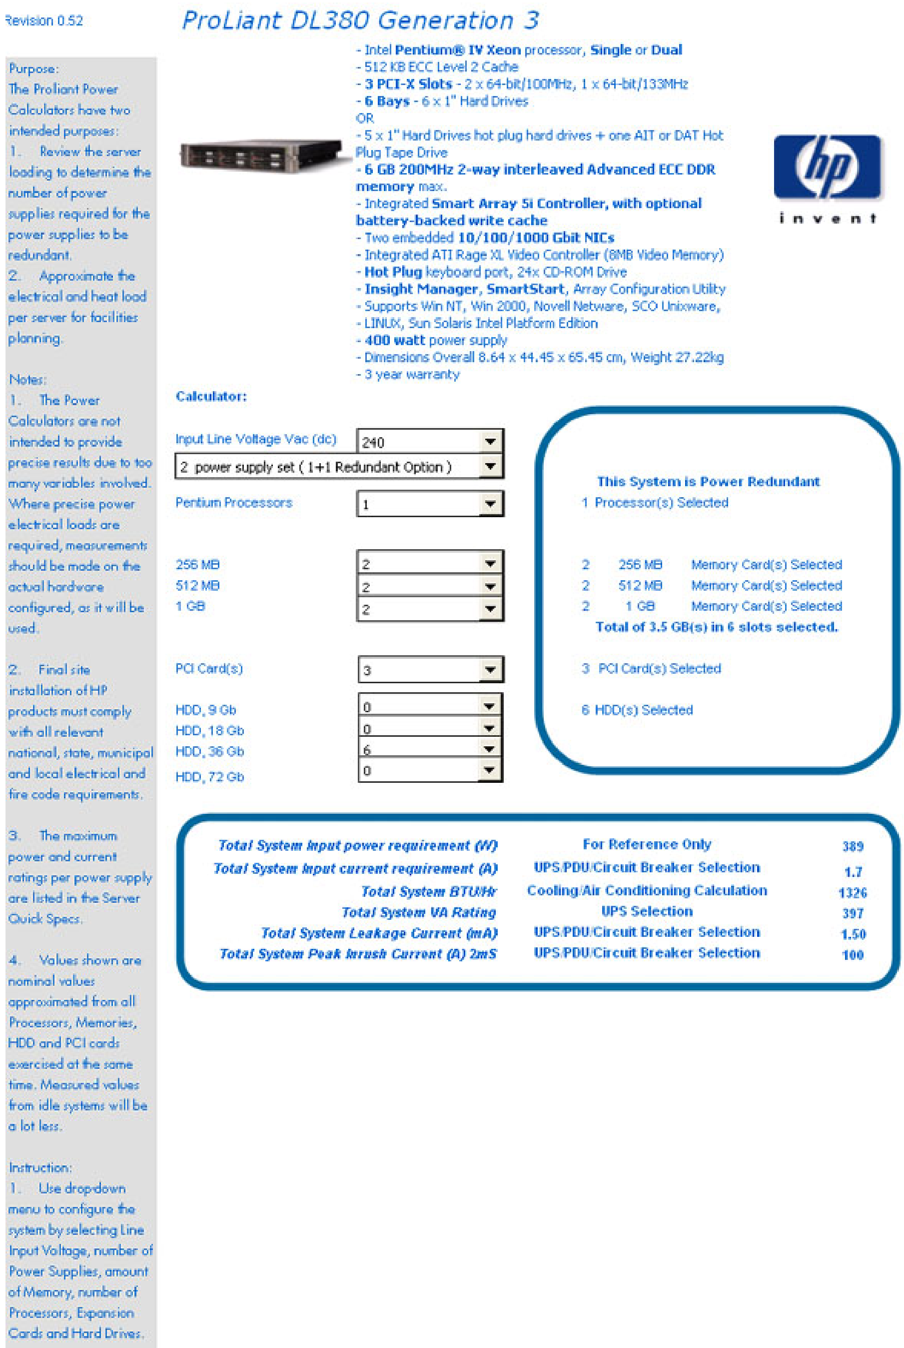

Factoring for future growth

Two to three years ago, the number of servers in a rack was four to six with an average load per rack of 1,500 to 3,000 W. Today, each rack averages eight to twelve servers with an average load per rack of 5,000 to 6,000 W. Implementations of loads up to 7,000 to 8,000 W are not unusual. HP has built the equipment to supply up to 18,000 W per cabinet.

Generation-to-generation density growth is something that also needs to be examined. Replacing a fully loaded ProLiant 6400 server with a ProLiant DL580 G2 server would cause an increase in power density of almost 140 percent. When examining spec sheets, customers should carefully compare the power density of the product currently used to the power density of the product considered for purchase.

Growth in power requirements is averaging 20 percent to 30 percent generation to generation. When planning data center implementations, customers should be sure to factor power at 100 percent of present use plus 20 percent.

Providing redundant power

The specific meaning of 1+1, N+1, and N+N power often raises questions. The first digit is the number of power supplies needed support the full configuration of the server. The second number is the number of online spares. To be redundant, most ProLiant servers require a 1+1 environment—that is, a single power supply can power the server and a second power supply provides redundancy.

The ProLiant ML570 server comes with three power supplies and supports N+1 redundancy. When using low voltage power, a minimum of two power supplies are required to support the full configuration. A third power supply provides redundancy so that one of the three power supplies can fail and the other two will support the load.

HP has only one product today that supports N+N redundancy: the ProLiant BL p-Class server blade. A single-phase ProLiant BL p-Class power enclosure will accommodate up to four power supplies. Two power supplies support the full capacity of the system with two online spares. A failure can occur in any combination of two power supplies in the system, and the other two power supplies will still provide full-capacity power.

Redundancy must be considered when determining how to distribute power through the PDUs. Since the redundant power supplies load share, a 1+1 system should be configured so that under normal operation, each power supply (and the PDUs and distribution line) takes 50 percent of the load. Then, if a failure occurs, the full load can be supplied by one power supply and line cord.

Configuration example: normal 1+1 redundant power operation

Figure 10 shows how a normal, maximum capacity, 1+1 redundant power system should be configured. At a server level, each power supply supports 50 percent of the load. At the rack level, each PDU supports 50 percent of the rack load.

Figure 10. Normal 1+1 redundant power operation.

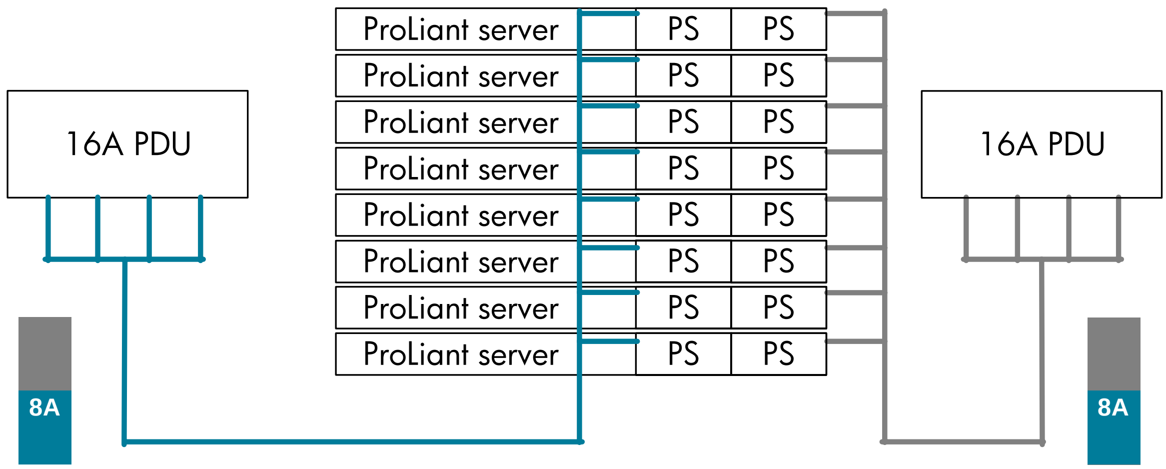

At a server level, each power supply supports 50 percent of the load. At the rack level, each PDU supports 50 percent of the rack load. In the event of a failure, one of the feeds drops out and the entire load transfers to the other side (Figure 11), with no downtime.

Figure 11. Feed failure in a 1+1 power configuration.

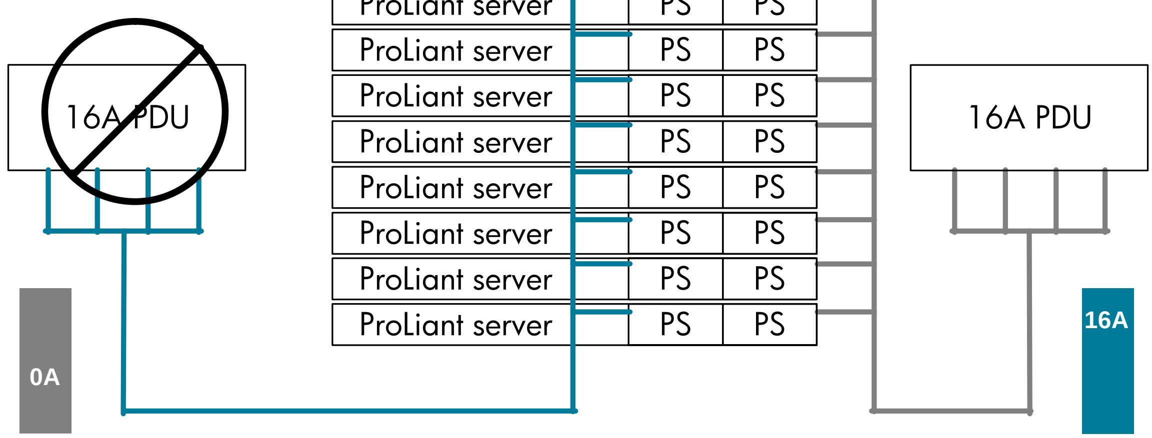

Configuration example: heavily loaded PDUs

It is becoming a common practice for customers to heavily load their PDUs. If an electrician's measurements indicate that the PDU is drawing less than its capacity, they add products to that rack and increase the load from the PDU. Figure 12 illustrates a configuration of this type.

Figure 12. 1+1 redundant configuration with heavily loaded PDUs.

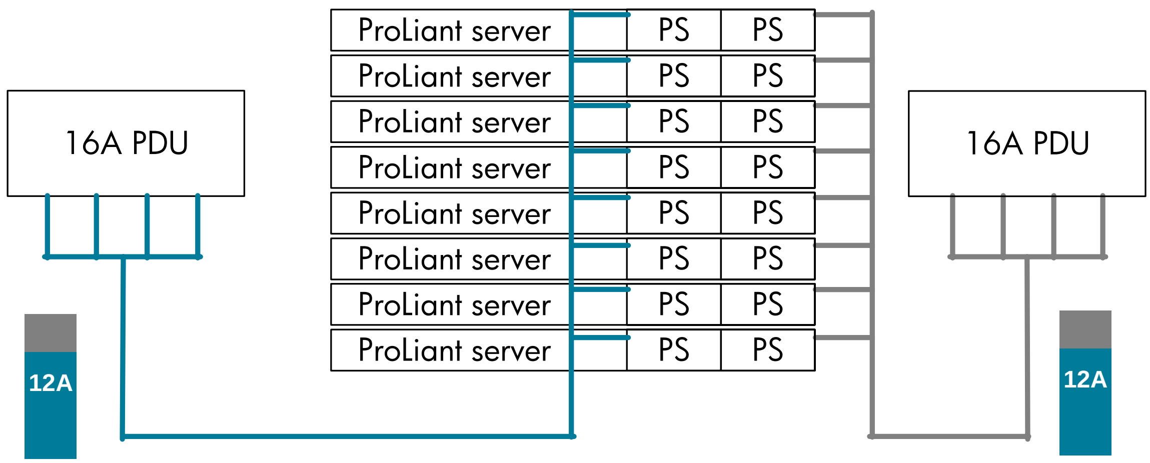

Overloading the PDUs in this way undermines redundancy. In the event of a failure, the entire load will shift to the remaining PDU, and that will result in an overload situation and ultimate failure of the entire rack (Figure 13). Loading PDUs in this way is very dangerous. In a normal operating environment, each PDU should never be loaded to more than 50 percent of its rated capacity.

Figure 13. Feed failure with heavily loaded PDUs.

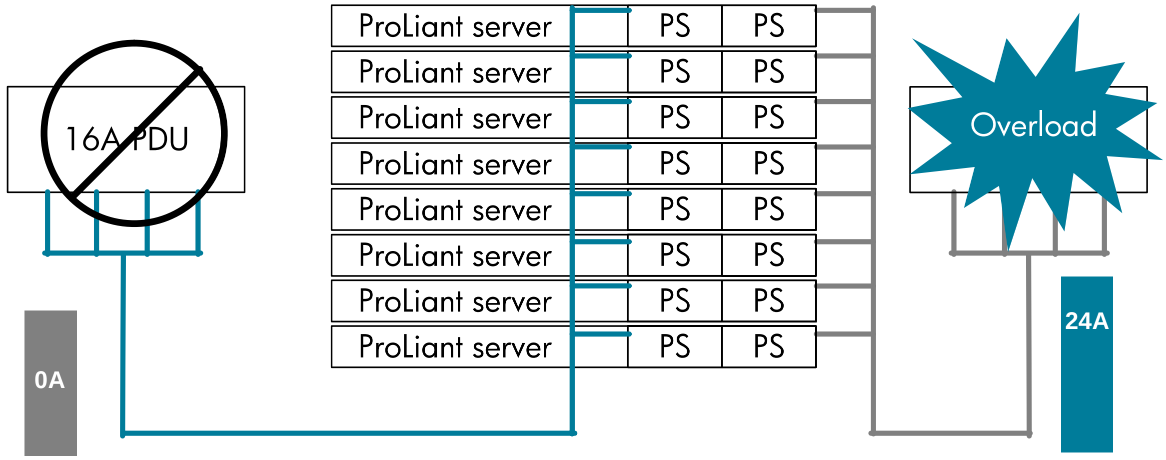

Conclusion

Power distribution planning is critical in today's ever-changing IT environment. Enterprises must provide adequate and consistent power to data centers to meet the 24x7 demands of users. While determining how much power is required and how to provide that power can be confusing, HP has developed tools and best practices to reduce the confusion. Understanding basic power considerations and the capabilities and installation of typical power components provides a basis for planning a facilities power distribution. HP power calculators help to automate the process.

For more information

Power calculators for HP hardware are available at
http://h18001.www1.hp.com/partners/microsoft/utilities/power.html

Call to action

To help us better understand and meet your needs for ISS technology information, please send questions and further comments about this paper to: TechCom@HP.com.

TC040401BP, 04/2004