Critical factors in intra-rack power distribution planning for high-density systems

technology brief, 2nd edition

Abstract

As server densities per rack increase, the power densities for those racks increase as well. This paper identifies critical factors that infrastructure designers and system administrators must consider to plan safe and efficient power distribution systems for high-density racks.

Introduction

Power distribution in the data center is changing with the continued increase in server density per rack. The processing power housed in a row of racks a few years ago has been compressed into the space of one rack or less. While the space needed for a given amount of processing has been reduced, the power requirements have been consolidated—and they are growing.

Early methods of distributing power within a rack (intra-rack power) are simply inadequate for meeting the increasing power requirements for racks of high-density equipment. Intra-rack power distribution technology is evolving to provide the levels of power required by racks of high-density servers.

Accurately determining the power distribution needs for a rack of high-density servers requires considering the following factors:

- Voltage level provided by the facility
- Type of power connectors to be used
- Estimated power consumption
- Power redundancy
- System expansion
- Serviceability

Failure to account for these factors can result in unscheduled downtime from equipment damage or overloads, increased costs, and safety issues. This paper explains the significance of these and other factors in planning for high-density computing.

Key factors of power distribution in high-density racks

Power distribution units (PDUs) are highly recommended for high-density systems that place heavy loading demands on the facility alternating current (AC) power bus. PDUs allow IT infrastructure designers to:

- Use segmented circuits to improve serviceability and enable sequential startup
- Achieve balanced loading between dual power buses
- Support redundant operation of equipment with redundant power supplies

A variety of PDU types are available to meet different needs. Key considerations for choosing an appropriate PDU include:

- Level and phase of power to be distributed
- Connector type
- Estimated power consumption of active components in the rack
- Physical space available for PDU mounting
- Special requirements or features such as circuit monitoring

Power level and phase

PDUs are designed to distribute a specific amount of current at either a low-line (100 – 127 VAC) or high-line (200 – 240 VAC) voltage range, depending on which level is supplied to the rack. Most HP server equipment features auto-sensing input circuitry that automatically detects and adjusts to the input voltage range.

High-line voltage is the favored choice for electrical equipment requiring high power consumption. As determined by Ohm's law [voltage (E) x current (I) = power (W)], higher voltage allows power supplies to produce the necessary power with less AC current. For example, a power supply that is required to produce 400 watts of power for a server will need to draw 3.3 amperes if operated with 120 VAC. The same power supply operating at 208VAC only requires 1.9 amperes to produce 400 watts—over a 30 percent reduction of AC current. High-line operation is generally more economical because it allows power supplies to create less heat (thereby reducing the cooling requirements). High-line voltage also can enhance reliability by extending equipment meantime-between-failure (MBF) periods. Note that some power supplies require high-line voltage to operate at maximum potential.

High-line voltage (208 VAC in North America, 200 VAC in Japan, 220 VAC elsewhere) is becoming commonly available in data centers. High-line voltage can be distributed as single or three-phase power. With single-phase power, voltage reaches its maximum and minimum 120 times a second (at 60 Hz). Three-phase power uses three cycles 120 degrees out of phase to provide a more constant voltage to the load (Figure 1). The efficiency of three-phase power allows the use of smaller circuit sizes (smaller-gauge conductors) to distribute the same amount of AC power.

Figure 1. Comparison of single- and three-Phase AC

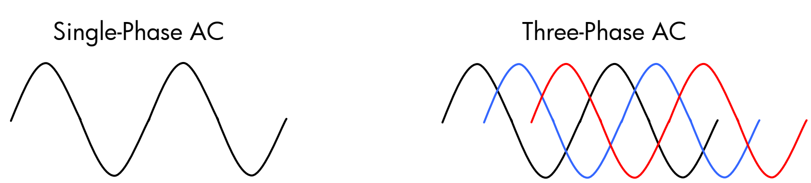

Connector types

Local and national electrical codes specify the types of connectors that should be used when interfacing with facility AC feeds. The location of the facility and the level, amount, and phase of the AC power available determine the type of connector to use. PDUs are designed to distribute power at a specific range and current. PDUs are provided with power cords of appropriate sizes and plug types for the geographic region where they will be used. Table 1 shows profiles of the most common connectors for different regions. Some regions may require connectors not shown in Table 1. Some facilities may require PDU power cords without connectors for hard-wired (terminal) installations.

Table 1. AC power connectors (rack-to-facility interface)

Geographic region or country	**Low-line**	**High-line**
North America & Japan	NEMA 5-15 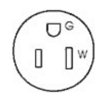	NEMA 6-15 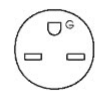
	NEMA 5-20 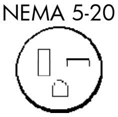	NEMA L6-20 
	NEMA L5-30 	IEC 309 2 phase [1]  IEC 309 3 phase [1] 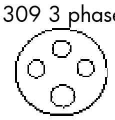
Europe	n/a	Shuko 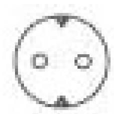
United Kingdom	n/a	BS1363 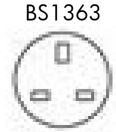

[1] Becoming more common in other regions.

For more information on AC power connector types refer to the HP technology brief "Selecting power cords and jumper cables for use with HP ProLiant servers" at the following URL:
http://h18004.www1.hp.com/products/servers/technology/whitepapers/datacenter.html

Power consumption

The type (low-line or high-line) and amount of power (in VA) required by the equipment in a rack largely determines the type of PDU(s) to use. HP provides online power calculators for making accurate power consumption estimates for most ProLiant servers. Each calculator (Figure 2) allows the IT infrastructure designer to predetermine the power requirements for a server chassis according to configuration parameters entered into the calculator. With accurate estimates of power requirements for individual server chassis, infrastructure designers can easily calculate power consumption for a fully-configured rack.

A catalog of HP Power Calculators is available for downloading at the following URL:
http://h71028.www7.hp.com/enterprise/cache/80364-0-0-0-121.aspx?jumpid=reg_R1002_USEN

Figure 2. HP Online Power Calculator (partial screen)

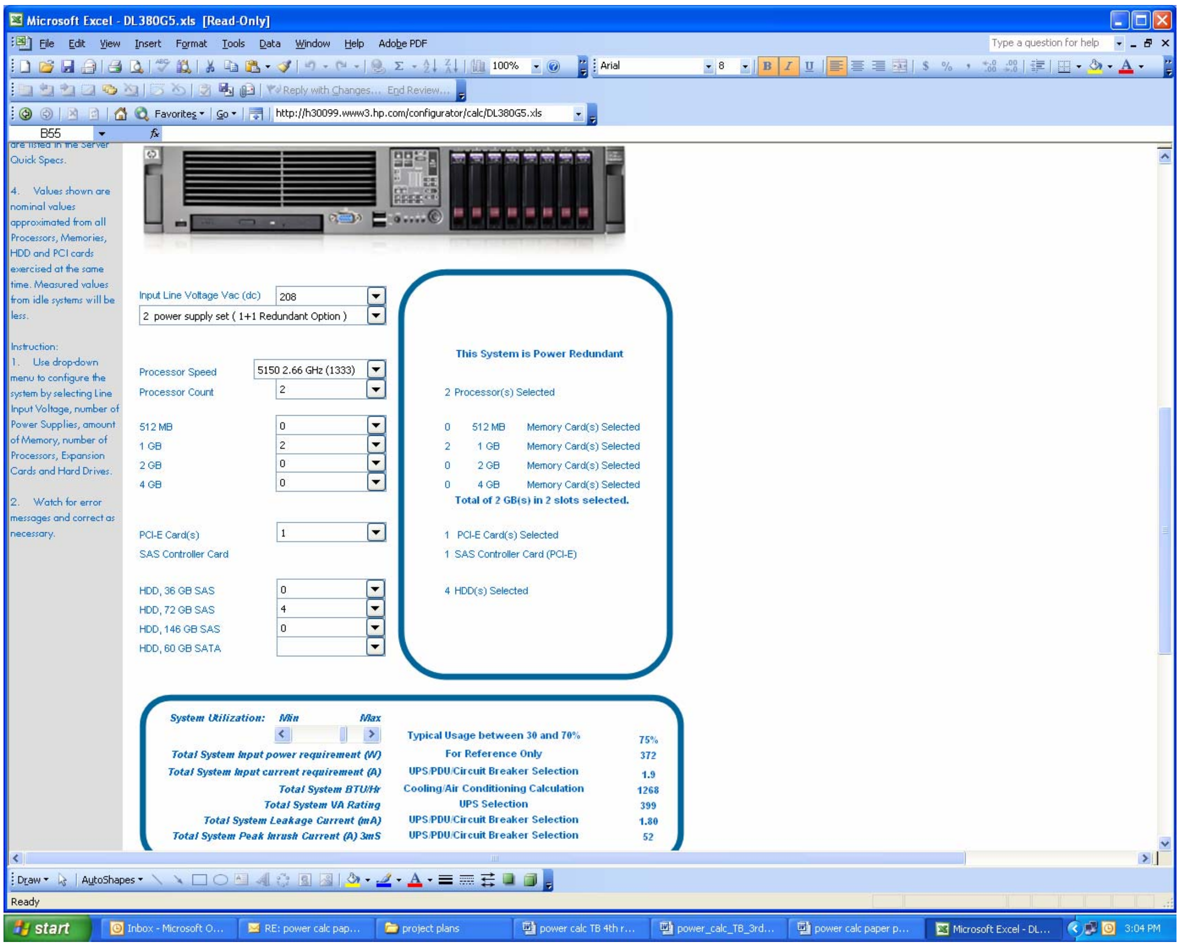

The HP Online Power Calculator also provides estimates of current draw, watts consumed, and inrush current. These are factors to consider when choosing PDUs. The HP Online Power Calculator also provides a BTUs/hr calculation for planning the cooling requirements of an installation.

Space issues

The 1U and blade server form factors allow for high-density installations that can occupy 100 percent of the front U space in a rack. This leaves only rear U or side ("zero U") space for mounting power distribution components. HP has designed components to address space issues that are becoming more acute in today's high-density systems. HP modular power distribution components include a control core module that offers flexible mounting options as illustrated in Figure 3.

Figure 3. Modular PDU mounting methods

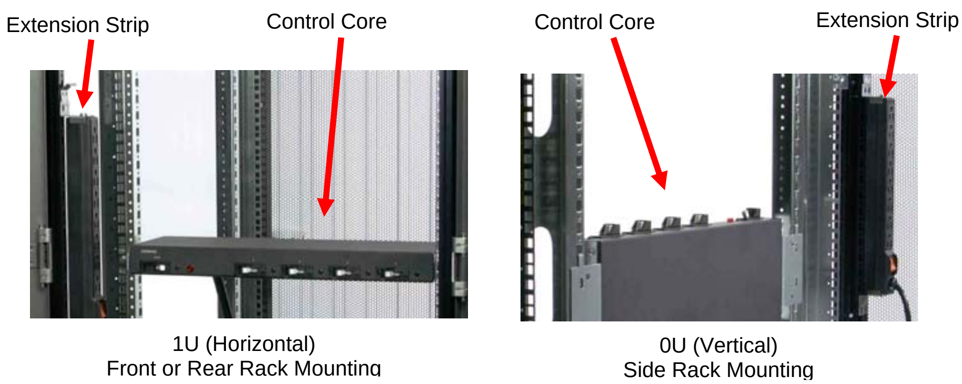

HP Monitored PDUs are self-contained in an extension strip form factor and designed to mount vertically in the same manner as the extension strips shown in Figure 3.

Note that some server systems may occupy most of a standard-sized rack space. They leave little room for power components or cables. These systems may require a rack extension kit (Figure 4) to accommodate cables and PDUs and to allow access to field-serviceable components. Two optional rack extensions kits, a 42U kit and a 36U kit, are available for the HP 10000 series equipment racks.

Figure 4. Rack extension kit

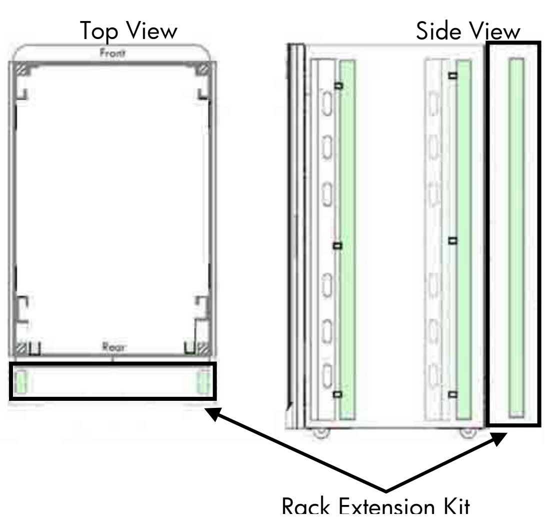

Segmented power distribution

In high-density systems where a single rack can support hundreds of users at a time, the power infrastructure should be designed so that a single system failure will have minimal effect on overall operations. Segmenting servers into groups for power distribution reduces the chance for unscheduled downtime resulting from overload or maintenance conditions that might otherwise shut down everything else in the rack. Segmenting servers into two or more groups also facilitates proper cable management and serviceability.

Segmenting can reduce the amount of inrush current associated with servers by allowing IT administrators to sequentially power up server groups. Circuits can be segmented by using a modular system of control cores and extension strips (Figure 5A) or by using monitored PDUs that individually support multiple load segments (Figure 5B). In either case, a circuit breaker is associated with each load segment (LS).

Figure 5. Segmented power distribution

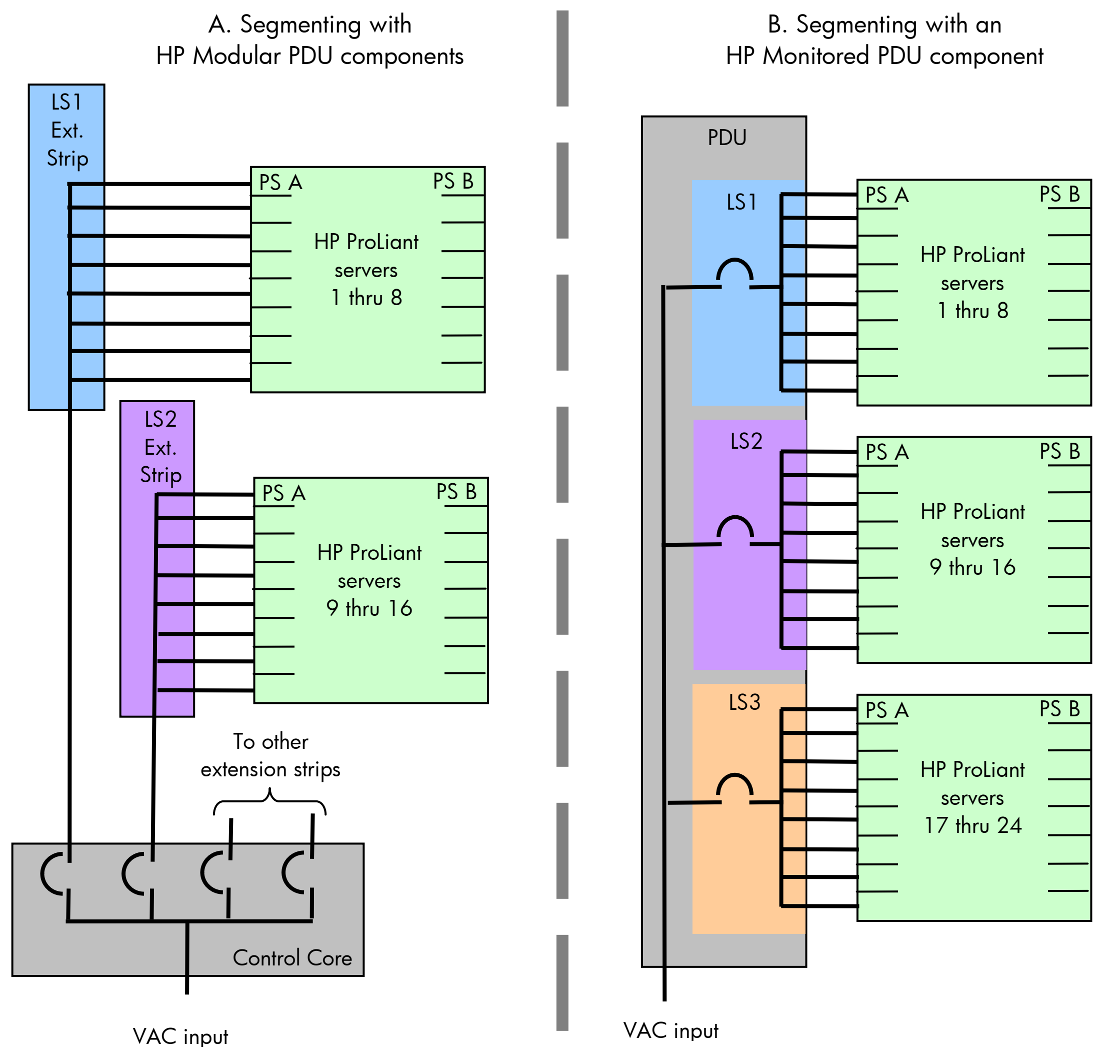

Redundant power distribution

Redundant power systems require two power distribution circuits, one for each power bus. Redundant operation requires configuring servers with redundant power supplies, each capable of supporting the full load of a server, if necessary. For redundant operation, the PDU or PDUs for each bus should be configured to handle less than half of their rated capacity for normal operation so that they can handle the full load of the system if one bus fails. The example shown in Figure 6 is based on an HP ProLiant server with a current draw of 1.9A operating with 208 VAC common to North America and Japan. The PDU for each bus feed must handle the full rack load (in this case, 32A) if one bus fails. In Figure 6, multi-segment 40A PDUs satisfy this requirement with adequate headroom.

Figure 6. Redundant power distribution

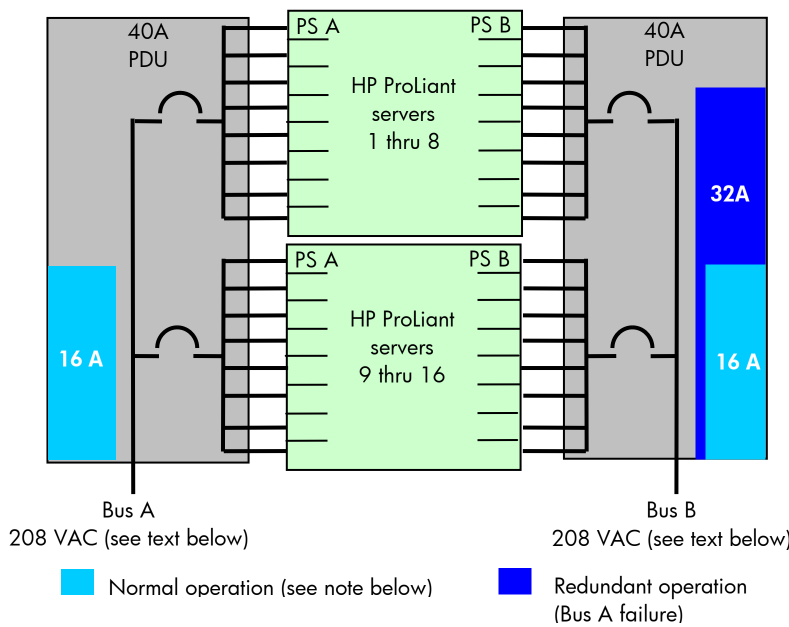

However, if the system in Figure 6 were operating in a 230 VAC environment, 32A PDUs could be used. The current draw of the same ProLiant server operating 230 VAC will be 1.7A each, resulting in a PDU current load of 13.6 A for normal operation and 27.3 A for all16 servers in redundant operation.

NOTE:
Depending on the server system employed, the actual load on redundant buses during normal operation may not be evenly divided.

Power circuit monitoring

Enterprise data systems can be maintained efficiently with minimum manpower by employing components that continuously monitor current loading on power circuits. Monitored systems make it easier to balance power distribution circuits, to gauge expansion capabilities, and to avoid unscheduled downtime resulting from circuit overloading or equipment damage.

Monitored PDUs from HP include an integrated display panel that provides:

- Real-time measurement of current load on each load segment
- Out-of-range voltage indication (10 to 20 percent under-voltage or over-voltage condition)
- Load warning alarms for each load segment (load detection of 80 to 102 percent rated capacity)
- Overload warning for each load segment (load detection of 102 percent or greater than rated capacity)

For maximum convenience, HP monitored PDUs can accept an optional PDU Management Module. This module enables IT administrators to monitor one or a pair of HP Monitored PDUs over a network. The HP PDU Management Module can be accessed through the built-in web interface or the text-based serial/Telnet interface. The HP PDU Management Module can also be integrated with SNMP management software such as HP Systems Insight Manger (SIM) that can receive and log SNMP traps (alerts) to manage alarm conditions.

Other considerations

An effective power distribution solution for a high-density rack involves much more than meeting the power requirement needs of rack-mounted components. Factors including cable routing, air flow, and cooling can also become critical in high-density systems. For more information on these and other related subjects, refer to the HP technology papers web site at the following URL: http://h18004.www1.hp.com/products/servers/technology/whitepapers/index.html

Sample power distribution solutions

This section includes examples for implementing power distribution methods for various rack configurations. These configurations are examples only. Numerous configurations may be appropriate depending on factors mentioned earlier such as:

- Server power consumption
- Server configuration (number of processors, hard drives, expansion cards, etc.)
- Server utilization
- Redundant/non-redundant requirements

NOTE:
These sample configurations do not take into account battery backup (UPS) considerations.

Rack configurations for distributing power to HP BladeSystem c-Class servers

This section describes power distribution configurations based on the HP BladeSystem c-Class system using the HP BladeSystem c7000 Enclosure and modular or monitored PDUs (Figure 7). The HP BladeSystem c7000 Enclosure uses auto-ranging power supplies and supports redundant power supply configurations. Up to four BladeSystem c7000 enclosures can be installed in a standard 42U rack.

Figure 7. HP ProLiant c-Class system and power distribution components

Table 2 shows the suggested PDU configurations when using one to four HP BladeSystem c7000 enclosures in a standard rack.

Table 2. Power distribution configurations for using one to four HP BladeSystem c7000 enclosures in a standard rack

No. of c7000 enclosures	Input Feed	Recommended HP PDU with AC input plug type	No. of PDUs Reqd. [1]	1Φ output current per PDU	Total 1Φ output current available [3]	Total output power available
1	1Φ 208VAC	Modular, 252663-D75 w/CS8265 50A non-NEMA	2	40	80 A	16.6 kW
	1Φ 220VAC	Modular, 252663-B34 IEC309 60A 3-wire blue	2	40	80 A	17.6 kW
2	1Φ 208VAC	Modular, 252663-D75 w/CS8265 50A non-NEMA	4	40	160 A	33.3 kW
	1Φ 220VAC	Modular, 252663-B34 IEC309 60A 3-wire blue	4	40	160 A	35.2 kW
	3Φ Δ 208VAC	Monitored, AF916A w/IEC309 60A 4-wire blue	2	84	168 A	35.0 kW
	3Φ Y 220VAC	Monitored, AF917A w/IEC309 32A 5-wire red	2	96	192 A	42.2 kW
3	1Φ 208VAC	Modular, 252663-D75 w/CS8265 50A non-NEMA	6	40	240 A	50.0 kW
	1Φ 220VAC	Modular, 252663-B34 w/IEC309 60A 3-wire blue	6	40	240 A	52.8 kW
	3Φ Δ 208VAC	Monitored, AF916A [2] w/IEC309 60A 4-wire blue	2	84	168 A	35.0 kW
	3Φ Y 220VAC	Monitored, AF917A [2] w/IEC309 32A 5-wire red	2	96	192 A	42.2 kW
	3Φ Y 220VAC	Monitored, AF917A [2] w/IEC309 32A 5-wire red	4	96	384 A	84.5 kW
4	1Φ 208VAC	Modular, 252663-D75 w/CS8265 50A non-NEMA	8	40	320 A	67.0 kW
	1Φ 220VAC	Modular, 252663-B34 w/IEC309 60A 3-wire blue	8	40	320 A	70.4 kW
	3Φ Δ 208VAC	Monitored, AF916A [2] w/IEC309 60A 4-wire blue	2	84	168 A	35.0 kW
	3Φ Y 220VAC	Monitored, AF917A [2] w/IEC309 32A 5-wire red	4	96	384 A	84.5 kW

[1] The number of required PDUs coincides with number of facility AC feeds. For PDU choices, access the online PDU matrix through the URL provided in the "For more information" section.

[2] Configuration requires optional 600mm rack extension kit (P/N AF058A).

[3] Redundant operation requires double this amount of Total 1Φ output current.

Products and part numbers discussed in this document are subject to change without notice. For the latest product and part numbers, access the URL provided in the "For more information" section.

Rack configurations for distributing power to HP ProLiant DL servers

This section describes power distribution configurations for a high-density server system based on HP ProLiant DL 360 G5 servers. The ProLiant DL 360 G5 is a 1U server that uses an auto-ranging power supply and supports power redundancy when used in the 1 + 1 power supply configuration option. These configurations use monitored PDUs and fixed cord extension bars (Figure 8).

Figure 8. HP ProLiant DL360 G5 system and power distribution components

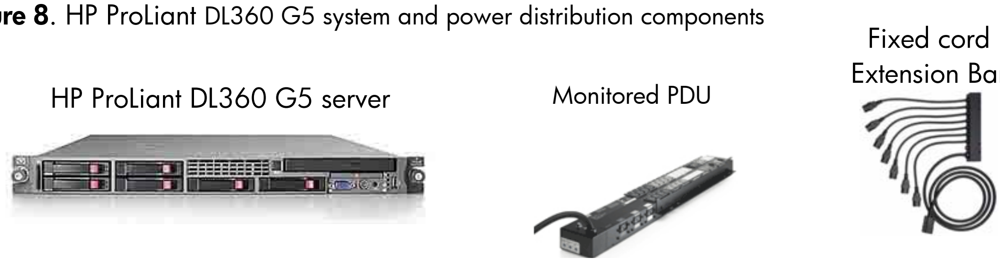

Table 3 shows the difference in power requirements of the ProLiant DL360 G5 server when it is used in two different scenarios. In one scenario, the server is configured and utilized to about half its potential. In the other scenario, the server is nearly at maximum configuration and estimated to run at a relatively high (70 percent) state of use. Both configurations employ redundant power (two supplies per system) and use high-line voltage. The maximum configuration operates off three-phase power. Although a standard 42U rack can physically accommodate 42 ProLiant DL360 G5 servers, it is advisable to leave space for server expansion, power distribution components such as PDUs and UPSs, keyboard/video/mouse (KVM) equipment, and/or networking components.

Table 3. HP ProLiant DL360 G5 server power distribution configurations

No. of servers in rack	System configuration (each server)	Total VA required	Total current draw	No./type of outlets required	Recommended PDU
30	One 2.0-GHz processor, 2 GB memory, one PCI-E card, two 36 GB hard drives, system utilization estimated at 40%	8.5 kVA	36 A	60/C13	1Φ 208VAC: Monitored PDU AF506A (x2) or monitored PDU AF505A 1Φ 220VAC: Monitored PDU AF508A (x2)
30	Two 3.0-GHz processors, 16 GB memory, two PCI-E cards, four 146-GB SAS hard drives, system utilization estimated at 70%	14.5 kVA	60 A	60/C13	3Φ 208VAC: Monitored PDU AF916A and fixed cord extension bar 351655-B21 (x5) 3Φ 220VAC: Monitored PDU AF509A (x2)

Products and part numbers discussed in this document are subject to change without notice. Refer to the URL provided in the “For more information” section for the latest available products and part numbers

Total VA and current requirements were determined using the HP Power Calculator Utility available at the following URL:
http://h30099.www3.hp.com/configurator/calc/Power%20Calculator%20Catalog.xls

Conclusion

High-density server systems that draw significant amounts of facility power require careful consideration to ensure that safety, availability, and manageability are not compromised. The variety of power distribution products available from HP allows an almost unlimited number of power distribution configurations to satisfy those goals.

NOTE:
Products and part numbers discussed in this document are subject to change without notice. Refer to the URL provided in the "For more information" section for the latest available products and part numbers.

For more information

For additional information, refer to the resources listed below.

Resource	Hyperlink
HP power distribution products	http://h18004.www1.hp.com/products/servers/proliantstorage/power-protection/pdu.html
HP PDU comparison chart	http://h18004.www1.hp.com/products/servers/proliantstorage/power-protection/power-distribution/powermonitoring/index.html
HP Industry Standard Servers technology papers	http://h18004.www1.hp.com/products/servers/technology/whitepapers/index.html.
HP remote management technology	http://h18013.www1.hp.com/products/servers/management/remotemgmt.html?jumpid=hpr_R1002_USEN

Call to action

Send comments about this paper to TechCom@HP.com.

TC070503TB, May 2007