

Power Concepts and Solutions for IBM® System x™ Servers

Chris Gillen, IBM

IBM System x - Western Region
Field Technical Sales Specialist



Table of Contents

1	Introduction	5
1.1	Overview	5
1.2	Basic Equations	5
1.3	Understanding Electricity	6
1.4	Redundant Power.....	8
1.5	AC Power.....	8
1.6	Power Hungry Servers	9
2	Delivering Power in the Data Center	11
2.1	Connectors	11
2.2	PDUs.....	12
2.2.1	Introduction	12
2.2.2	PDUs and Racks.....	13
2.2.3	IBM PDU Information.....	14
2.2.4	PDU Summary.....	16
3	Designing the Power Solution	17
3.1	System x Power Requirements.....	17
3.1.1	System x Power Configurator.....	17
3.1.2	System x Power Considerations	18
3.2	Sample Solution Configuration	18
4	BladeCenter Power Considerations.....	21
4.1	BladeCenter Power Domains.....	21
4.2	Designing Power Solutions for BladeCenter	22
4.3	Power Management in BladeCenter.....	23
4.4	BladeCenter H	25
4.4.1	BCH Power Introduction	25
4.4.2	BladeCenter 2900W Power Supply	25
4.4.3	Wiring Examples for BCH	27
4.4.4	FAQ	28
5	UPS	29
5.1.1	BladeCenter and 7500 XHV or 10K XHV UPS	29
5.1.2	BladeCenter and UPS 3000 XHV	31
6	Cooling.....	32
Appendix A.	Connectors	34
Appendix B.	Hardwired UPS Connections.....	35
Appendix C.	Power Transmission	36
Appendix D.	Related Publications and Tools	38
	Related Publications	38
	Related Tools	38

Table of Tables

Table 1	Enterprise PDU Part Number Transitions.....	16
Table 2	PDU Summary.....	16
Table 3	Additional line cord part numbers.....	16
Table 4	Power Requirements for Sample Configuration 1	19
Table 5	Power Requirements for Sample Configuration 2	20
Table 6	BladeCenter Power Requirements 1	22
Table 7	BladeCenter Power Requirements 2	22
Table 8	BladeCenter H Line Cords.....	26
Table 9	UPS Specifications.....	29

Table of Figures

Figure 1 Sprinkler Heads.....	7
Figure 2 Sprinkler Layout.....	7
Figure 3 Useful Voltage.....	8
Figure 4 Three Phase Power.....	9
Figure 5 Power Usage by Component.....	10
Figure 6 CPU and Rack Power Requirements over time.....	10
Figure 7 C13/C14 Connectors.....	11
Figure 8 C19/C20 Connectors.....	11
Figure 9 Single Rack PDU Option.....	12
Figure 10 Redundant Power Distribution.....	13
Figure 11 DPI Front-end PDU Configuration.....	13
Figure 12 Burndy Connector on Front End and Enterprise PDUs.....	15
Figure 13 Enterprise PDU Layout Example.....	15
Figure 14 Sample Configuration.....	18
Figure 15 Power Domains.....	21
Figure 16 BladeCenter with 30A Circuits.....	22
Figure 17 Two BladeCenters with 60A Single Phase Power.....	23
Figure 18 Three BladeCenters with 60A – Three Phase power.....	23
Figure 19 Power Management Policy Settings.....	24
Figure 20 Throttling warning message.....	24
Figure 21 BladeCenter H.....	25
Figure 22 BladeCenter H Power Supply.....	25
Figure 23 BladeCenter H Power Connections.....	26
Figure 24 BCH with 60A Single Phase.....	27
Figure 25 BCH with 60A Three Phase.....	27
Figure 26 BCH with Direct Connections.....	28
Figure 27 BladeCenter and 10K UPS – Sample Configuration.....	30
Figure 28 BladeCenter and 3000XHV – Sample Configuration.....	31
Figure 29 Front to back cooling.....	32
Figure 30 Hot at the top of the rack.....	33
Figure 31 L6-30P.....	34
Figure 32 IEC309 2P+Gnd.....	34
Figure 33 IEC309 3P+Gnd.....	34
Figure 34 IEC309 3P+Gnd Connector.....	34
Figure 35 UPS 7500 Hardwired.....	35
Figure 36 Hardwired Wall Connection (bottom box).....	35
Figure 37 Transmission and Distribution Grid.....	36
Figure 38 Three Phase Distribution.....	37
Figure 39 Transformer.....	37

1 Introduction

1.1 Overview

This guide is intended to present not only the “how” of power delivery for System x Servers, but also the “what” and the “why” behind power delivery configurations for System x solutions. This will be done by providing an overview of how electricity works then providing some reference material and examples to assist in creating valid power configurations. The goal is to present the connection between electrical theory and electrical application.

This guide is US only. All worldwide permutations of power delivery have been removed in the hopes of simplifying the presentation.

This document is a supplement the IBM® publications provided with various products and is intended to be a personal productivity tool. It is not intended to be comprehensive, and is provided for guidance only, on an 'as is basis' without warranty of any kind. Please be aware that its contents have not been certified by IBM.

1.2 Basic Equations

The motion of electric charges (for example, electrons in a conducting material) constitutes an electric current. Electrical current is measured in *amperes*, or *amps* (**A**).

A number of materials, most notably metals such as copper and silver, contain many free electrons which move in response to electric fields. These materials are capable of carrying an electric current. Each free electron in a metal wire carrying current is accelerated by the electric field until it loses its velocity as a result of a collision within the metal. It then accelerates again until it suffers a subsequent collision. The energy required to force the successive accelerations and move an electron is called the electric *potential difference* between two points. This potential difference is termed a *volt* (**V**).

The resistance that each free electron encounters as a result of collisions when moving in a conductor depends on the resistivity of that conductor. This resistance is affected both by the natural properties of the conductor and as well as the shape of the conductor. The resistance value (**R**) is measured in ohms (Ω) and the equation is expressed as:

Resistance (R) is equal to a resistivity constant (ρ) times length (L)
divided by the cross section area (A)

$$R = \rho \frac{L}{A}$$

Longer wires have more resistance, and wires with small cross sectional areas have more resistance.

The higher the resistance, the higher the force necessary to push the current through the circuit. Since the wires inside a toaster offer a high resistance to the flow of electrons, a flow of current makes them get red hot.

There is a direct relationship between current, voltage and resistance which is specified by an equation known as Ohm's law:

$$\text{Voltage (V) is equal to amperage (I) times resistance (R): } V = I * R$$

The basic point to remember is that as the resistance of a circuit goes up, the amount of current that flows in the circuit will go down if voltage remains the same.

The amount of work that the electricity can perform is expressed as power. The unit of power is the watt (**W**). The relationship between power, voltage and current is known as Joule's law and is expressed as:

Power (W) is equal to amperage (I) times voltage (V): $W=I * V$

And by extension: $W = I^2 * R$

Since the coils in a toaster have a high R value, they generate a lot of power as heat and some as red light.

Quiz 1A: If server draws 500W how many Amps will it draw at 208V?

2.4A at 208V ($2.4*208 = 500$)

Quiz 1B: How many amps will that server draw at 110V?

4.6A at 110V ($4.6*110 = 500$)

The power equation is the important one. With computers, W, the power required by our servers, is essentially a fixed value in the above equation *per configuration*. That means that we need to manage the current and the voltage values.

In most cases, the real limitation is amps since a wire (or PDU) will give off heat (W) based on the values of $I^2 * R$. Consider a server that pulls 300W of power. At 208V, it will pull 1.5A and at 120V it will pull 2.5A. The circuit/PDUs will then generate either $2.25*R$ or $6.25*R$ in heat, so obviously the $2.25*R$ is a more manageable value.

1.3 Understanding Electricity

It is useful to make an analogy to water when trying to understand electrical theory. In a comparison with water flowing through a pipe, the electrons would be the water and the wire is the pipe.

Amperage can be compared to the volume of water flowing through a pipe. When connected to a circuit, an ammeter will measure the actual amount of current flowing through the circuit, just as the water meter for a house measures how much water the sink, washing machine, shower and sprinkler system uses. When relatively few electrons flow through a circuit, the amperage is low. When many electrons flow, the amperage is high.

Voltage can be viewed as the equivalent of water pressure. Water pressure is measured in units such as pounds per square inch (psi) and as previously mentioned; electrical pressure is measured in volts. When a voltmeter is connected to a circuit, it is measuring the electrical pressure. A battery or a generator is the equivalent to a water pump.

The final part of Ohm's law is the resistance of the conductor. To use the water analogy again, the large wire would be like a fire hose (low resistance); the small wire would be like a garden hose (high resistance).

Let's take a closer look at a garden sprinkler system to understand how current and voltage affect our power designs.

First off, consider that the sprinkler heads are the equivalent of servers. A 1U server might be a dripper (below left), and a high power, 4 socket server such as an x3850 would be a larger, impact sprinkler head (below right).



65PJADJ-TNT

Figure 1 Sprinkler Heads

Figure 2 Sprinkler Layout shows us a sample sprinkler layout.




Figure 2 Sprinkler Layout

At the main valve there is 50psi. The system is laid out using 3/4" PVC. In this example, the maximum amount of water that can be delivered through the main valve is 16 gps. In real life, this value can be determined by running a flow test. Knowing the maximum flow, we can determine how many sprinkler heads we can connect per zone (circuit). For example, we can hookup 5 sprinkler heads that deliver 3 gps before we need to create another sprinkler zone. Once a particular zone is maxed out, we need to create a new one.

If the system is laid out using 1" PVC, then it might deliver 24 gps and we would be able to add 3 more sprinklers to this zone. This would essentially be the equivalent of having larger cables and higher capacity PDUs.

1.4 Redundant Power

When you look at the power requirements for a server, do not be swayed or confused by whether redundant power is being used or not. The design must take into account that 1 power supply may have to provide the full power load by the server and that 1 circuit may have to supply the power for all servers connected to a redundant set of PDUs.

Consider that we have redundant water supplies to our home sprinkler system. We now have 2 50psi sources that can each provide 16 gpm each. Under normal circumstances, both valves operate simultaneously and end up delivering roughly 8 gpm each. However, that doesn't mean we can add 5 more heads to our zone. If something were to happen to our first water supply, then we have to be able to water the lawn using just 1 source so we are still limited to 5 heads.

Under normal circumstances, we are running at roughly 50% capacity (worst-case). The point we are trying to make is that while we are running at 50% in normal operation, we need to plan for 100% in a failure situation. See Section 3 for a further discussion of redundant power.

1.5 AC Power

There are two different types of power:

- **AC** – the type of power that gets distributed to the data center or your house
- **DC** – the type of power the server components run on

Single-phase AC power is what you have in your house. If you use an oscilloscope and look at the power found at a normal wall-plate outlet in your house, what you will find is that the power at the wall plate looks like a sine wave oscillates between -170 volts and 170 volts.

The effective (RMS) voltage is 120 volts. RMS stands for "Root Mean Square," and it is the standard way of measuring and reporting effective alternating current and voltage. It is not the peak; it is the average. The RMS value is found by multiplying the peak amplitude by the square root of 2 (approximately 0.707). This yields the actual, useable voltage.




Figure 3 Useful Voltage

As we work with power configurations, we can encounter both single phase and three phase power sources. Since AC power is a wave, the phases of these waves can be offset to deliver more power.




Figure 4 Three Phase Power

Power plants produce three different phases of AC power simultaneously, and the three phases are offset 120 degrees from each other. There are four wires coming out of every power plant: the three phases plus a neutral or ground common to all three. Three phase power is three single phases synchronized and offset by 120 degrees. The advantage is that each phase can be peeled off to power machines in the rack.

Why three phases? Why not one or two or four? In 1-phase and 2-phase power, there are 120 moments per second when a sine wave is crossing zero volts. In 3-phase power, at any given moment one of the three phases is nearing a peak. Four phase power delivery from the power plant would not significantly improve things but would add a fourth wire, so 3-phase is the natural settling point.

What is the real impact of three-phase versus single phase at the rack? A 60A single phase circuit can provide just that, 60A, which, as previously mentioned, gets derated down to 48A. The three phase circuit can provide 27A *per phase* for a total of 81A per circuit. This has to do with RMS values and the fact that the circuit spends less time at zero.

For more information on how power is delivered from the power plant to the local neighborhood, see 1.Appendix C.

1.6 Power Hungry Servers

A power supply converts AC to DC. It takes more AC input to make the DC output. The efficiency of the power supply dictates how much more. As an example, the power supplies in the IBM BladeCenter have a relatively high efficiency and convert 90%+ of AC power into DC power.

So where are server power requirements headed? You may hear 2 different messages:

- 1) CPUs are going to require more power short term, then they will begin to require less.
- 2) Servers are going to require more and more power in the future.

Well, both are essentially true.

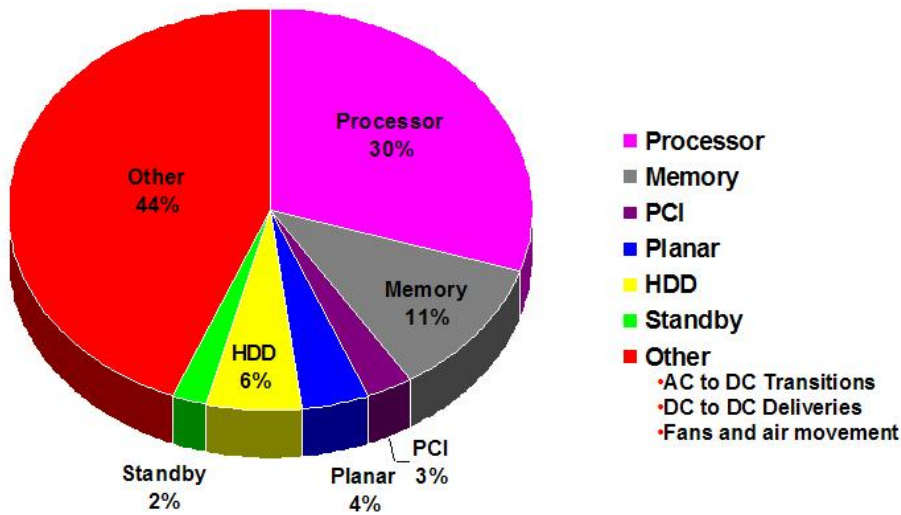


Figure 5 Power Usage by Component

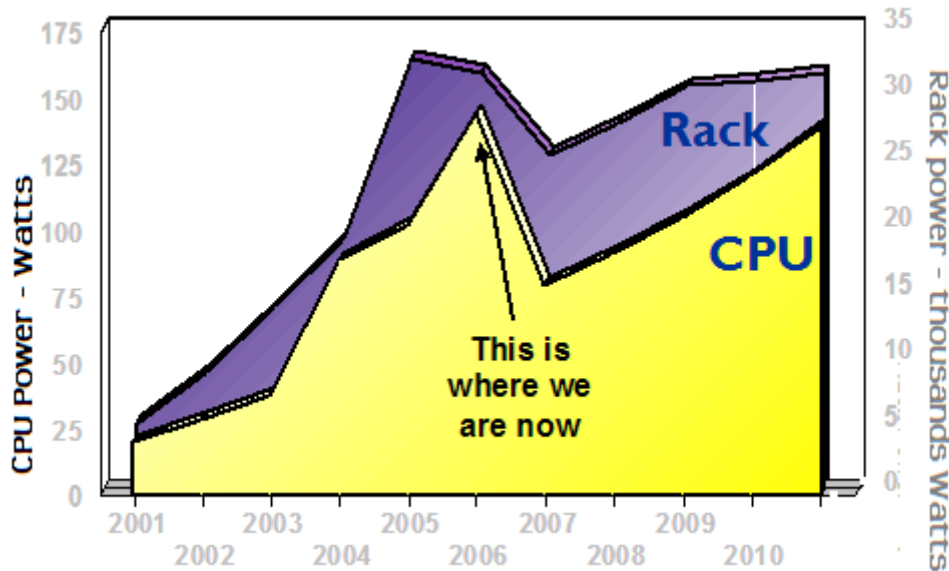


Figure 6 CPU and Rack Power Requirements over time

CPU Power requirements will decrease with new technology in the near future, but then the power curve tilts up again. As shown in the first chart, system level power has many components and some of these have decreasing power requirements, while others have increasing requirements. The rack level power consideration includes increasing density of components.

2 Delivering Power in the Data Center


2.1 Connectors

There are many different types of connectors in use. The ones shown below are the most common you will see. The choice of a connector is primarily driven by the maximum amount of power that it will be called on to deliver.

NOTE: NEMA vs. IEC. The larger size of the NEMA outlet means that fewer outlets can be incorporated into the same "real estate" on a piece of electronic equipment. As a result, products with NEMA outlets typically are less dense and not as flexible as products designed with IEC outlets.

1. NEMA 5-15 – Amperage limitation is 15A. This is the plug generally seen on home outlets.





Seen one of these? It is a NEMA 5-20p outlet. The "20" indicates that it can support up to 20A, but you'll need an appliance equipped with a 5-20p plug to take advantage of that. A 5-20p connector has 1 prong that is oriented horizontally.

2. C13 connector – 10A rated (female) rack, high density and enterprise PDU outlets. (see right, top row). The C13 outlet has been adopted throughout the world as the choice for internal rack power. It has universal support for both low and high voltage.
3. C14 connector – 10A rated (male). Inlet used on most rack and tower servers (power supplies less than 1500W) (see right, bottom row)




Figure 7 C13/C14 Connectors

Note: The lower number is female. Power flows from female connector to the male connector. Another way to visualize it is that the female connector will be hot and the male connector will not be hot. This makes sense since a C14 connector could be shorted with something like a screwdriver, but that couldn't be done with a C13 connector.

4. C19 – 16A rated (female). Rack, high density and enterprise PDU outlets. FE PDU to rack PDU connector. BladeCenter line cord. (see right, top row)
5. C20 – 16A rated (male) BladeCenter power supplies and Rack PDU inlets. (see right, bottom row)




Figure 8 C19/C20 Connectors

6. L5-20P and L6-20P: Amperage limitation is 20A. Used on line cords for PDUs.
7. L6-30P: Amperage limitation is 30A. Used on line cords for PDUs. (see Figure 31 L6-30P)
8. IEC309 2P+G - Amperage limitation is 60A single phase. Used on line cords for PDUs. (see Figure 32 IEC309 2P+Gnd)
9. IEC309 3P+G - Amperage limitation is 60A three phase. Used on line cords for PDUs. (see Figure 33 IEC309 3P+Gnd)
10. Burndy - Amperage limitation: line cord dependent (very flexible) Used on BladeCenter H and IBM PDUs. (see Figure 12 Burndy Connector)

2.2 PDUs

2.2.1 Introduction

It is easy to get a little overwhelmed by the numerous PDU choices available for configuring the power delivery to your rack and servers.

Note: Underwriters Laboratories derates all circuits in the US by 20%. Thus:

- a 20amp circuit will only officially support 16A.
- a 30amp circuit will only officially support 24A.
- a 60amp circuit will only officially support 48A.

Finally, a 60A 3Phase PDU will support 27x3 or 81A

Figure 9 shows the simplest PDU configuration. In this configuration there is a rack PDU plugged into a wall outlet (or perhaps a front-end PDU). This simple topology provides non-redundant power for up to seven devices using a Universal Rack PDU such as 32P1736 (see Section 2.2.2).

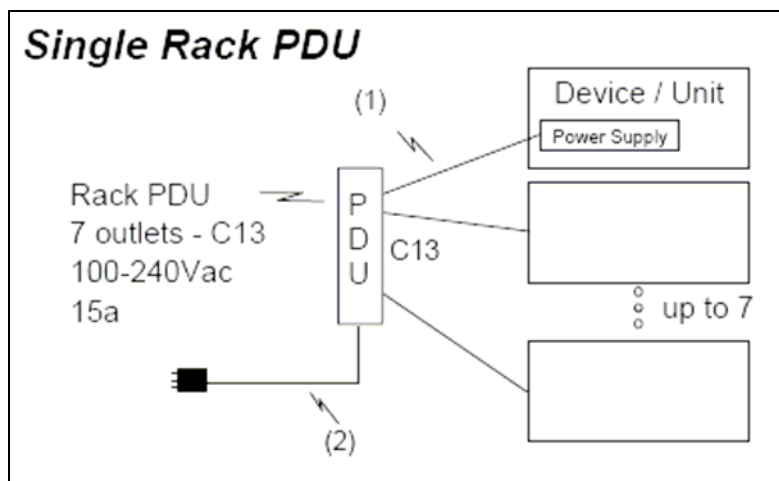


Figure 9 Single Rack PDU Option

The Universal Rack PDU is limited to 15A, so the maximum watts supported are either 1760W or 3328W depending on whether we have a low or high voltage circuit. ($W = I \times V$)

The addition of a second circuit and second PDU to the configuration provides power redundancy for our servers. With this topology, even if one power source is lost, the second power source and PDU are designed to power the load without bringing the solution down.

Again, we have the same power limitation as Example 1.

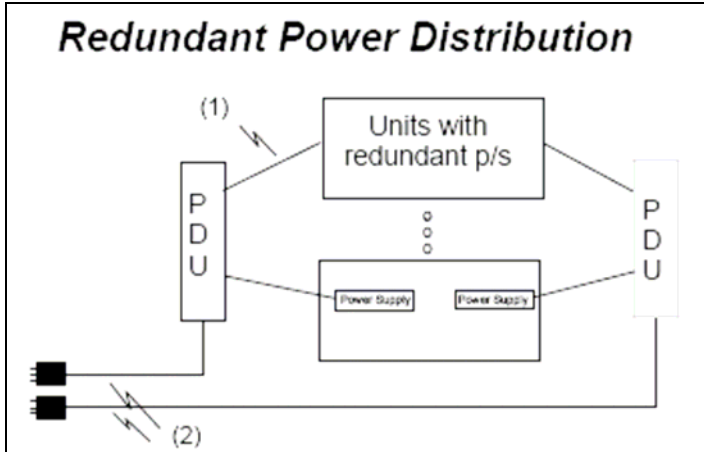


Figure 10 Redundant Power Distribution

The addition of front-end or enterprise PDUs to the power configuration allows the connection of multiple rack PDUs to higher-current source circuits.

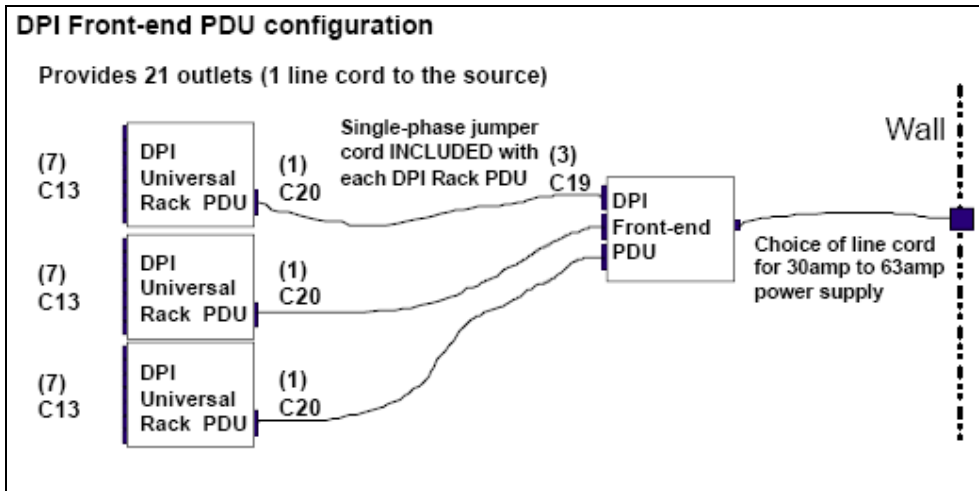


Figure 11 DPI Front-end PDU Configuration

2.2.2 PDUs and Racks

IBM 42U Enterprise rack allows 16 PDUs to be placed in side walls. IBM Standard Rack allows up to 12 PDUs in side walls.

NetBAY42 Enterprise Rack Cabinet (93084RX): This 42U high rack conforms to the industry-standard, EIA-310-D for 19-inch, type A rack cabinets. Four side-wall compartments support 1U PDUs and switches without taking up valuable rack space.

NetBAY42 Standard Rack Cabinet (93074RX): This 42U high rack conforms to the industry-standard, EIA-310-D for 19-inch, type A rack cabinets. Six side-wall compartments support 1U PDUs and switches without taking up valuable rack space.

NetBAY25 Standard Rack Cabinet (93072RX): Identical in design and features as the NetBAY42 SR, except it Provides 25U of rack space and contains two single side-wall compartments.


2.2.3 IBM PDU Information

IBM offers Distributed Power Interconnect (DPI) products, a family of modular "zero-U" power distribution units (PDUs) designed to be capable of supporting solutions from a handful of servers to an array of large enterprise servers—all without consuming valuable rack space. These units can be configured to provide redundant power distribution, thus helping to eliminate a point of failure and assure higher levels of data availability.

As the initial version of this document is being published, the entire industry is going through a RoHS part number transition. Both old and new part numbers are included in the descriptions and tables. The red part numbers are the new, RoHS part numbers.

2.2.3.1 DPI Universal Rack PDU

- Universal voltage (100-240Vac, 50-60Hz)
- Contains 7 C13 outlets.
- Power input is a C20 and is on the back of the unit
- The unit is 1U in height and 1/2 rack width. Installs in side panel
- Unit contains a 15A breaker for protection
- Always ships with a 2M IEC C19-C20 line cord for attachment to a Front End PDU. The NA version also ships with L5-20p and L6-20p line cords.
- Part numbers 32P1736 and **39Y8951**



2.2.3.2 DPI Front End PDU

- Voltage: Available in both low and high voltage versions.
- Contains 3 IEC C19 outlets.
- 1U height; 1/2 rack width, ships with side wall mounting hardware only. FE PDUs can be mounted in the EIA space using the rack PDU hardware.
- Can be mounted with Rack PDUs on same mounting plate for ease of installation- Side panel or EIA space
- The DPI Front End PDU has a special Burndy connector on the back of unit (see Figure 12 Burndy Connector). This connector allows the customer to attach line cords at their site. It also allows them to choose between LV 30A, HV 30A, or HV 60A via a line cord change.
- NOTE: Front End PDUs can be used to directly power high amperage servers such as IBM BladeCenter when required.
- Part numbers 32P1751, 32P1766, 32P1767 and **39Y8938, 39Y8939, 39Y8940**





Figure 12 Burndy Connector on Front End and Enterprise PDUs

2.2.3.3 DPI Enterprise PDU

The Enterprise PDU has the outlets, breakers and cord connections ALL on the same side for easy access.

- These are high voltage PDUs only
- Two categories:
 - One contains six C19 outlets. Unit contains a 16A breaker per outlet for protection
 - Other model contains twelve C13 outlets. Unit contains a 10A breaker for each pair of outlets
- The unit is 1U in height and full rack width. Hardware included for both EIA mounting and side pocket mounting.
- Power input on each unit is a changeable cord which varies depending on model. (Burndy connector)
- Units are available in 60A single (IEC309 2P+G connector) and three-phase (IEC309 3P+G connector) versions to maximize power capacity

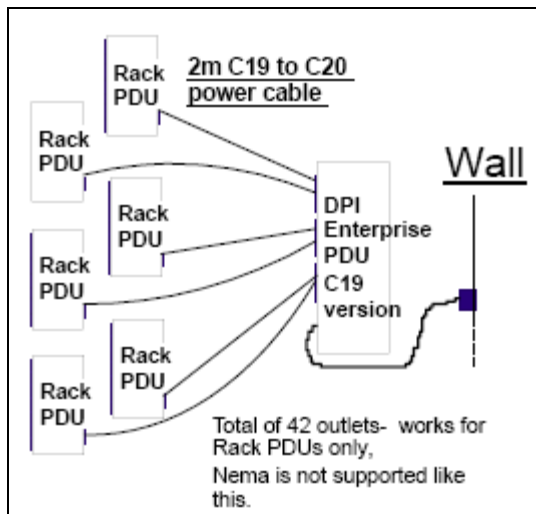


Figure 13 Enterprise PDU Layout Example

- Part numbers: The new RoHS compliant Enterprise DPI PDUs do not (in all but one case) ship with Line Cords. The line cords must be ordered separately.

PDU Description	Previous Part Number	New PDU Part Number	New Line Cord Part Number
30A Single Phase, C13	26K4246	39Y8941	40K9614
60A Single Phase, C13	26K4242	39Y8941	40K9615
60A Single Phase, C19	26K4256	39Y8948	40K9615
60A Three Phase, C13	26K4244	No replacement - 60A 3P C13 PDUs not widely used	
60A Three Phase, C19	26K4258	39Y8923	Comes with fixed line cord.

Table 1 Enterprise PDU Part Number Transitions

2.2.4 PDU Summary

Model	Part Numbers	Outlets	Line Cord	Capacity
DPI Universal Rack PDU	32P1736, 39Y8951	7 IEC C13 or 6 NEMA 5-15	IEC C19 to L5-20 and IEC C19 to L6-20	1800W (120V), 3120W (208V)
DPI Front End PDU (30A, low voltage)	32P1751, 39Y8938	3 C19 outlets	Burndy to L5-30R (120V)	2880W (120V)
DPI Front End PDU (30A, high voltage)	32P1766, 39Y8939	3 C19 outlets	Burndy to 30A L6-30P (208V)	4992W (208v)
DPI Front End PDU (60A)	32P1767, 39Y8940	3 C19 outlets	Burndy to IEC 309 60A 2P + Gnd	9360W (208V)
DPI 30A 1 Phase C13 Enterprise PDU	26K4246 pdu 39Y8941 line cord 40K9614	12 C13 outlets	Burndy to L6-30P (208V)	4992W (30a 208v)
DPI 60A 1 Phase C13 Enterprise PDU	26K4242 pdu 39Y8941 line cord 40K9615	12 C13 outlets	Burndy to IEC309 2P+G (208V)	9984W (208v)
DPI 60A 1 Phase C19 Enterprise PDU	26K4256 pdu 39Y8948 line cord 40K9615	6 C19 outlets	Burndy to IEC309 2P+G (208V)	9984W (208v)
DPI 60A 3 Phase C13 Enterprise PDU	26K4244 n/a	12 C13 outlets	Burndy to IEC309 3P+G (208V)	14976W (208v)
DPI 60A 3 Phase C19 Enterprise PDU	26K4258 39Y8923	6 C19 outlets	Burndy to IEC309 3P+G (208V)	14976W (208v)

Table 2 PDU Summary

94G6667	39Y7931	Power Cable - C13 / NEMA 5-15P 14ft
94G7448	39Y7932	12ft Power Cable C13-C14
90P4849	39Y7916	IEC 309 C19 to C20 intra rack cable
23K4809	39Y7938	IEC309 C20 to C13 rack jumper cable

Table 3 Additional line cord part numbers

3 Designing the Power Solution

In order to correctly and completely design our power solution, there are several configuration "parameters" that we need to know.

- 1) Power draw of the systems that we are configuring the power solution for. Note, in most situations, the number of systems we include per rack is going to be driven by power requirements, not "U" requirements.
- 2) High voltage (208V-220V) or low voltage (100V-120V) power source.

Note: all System x servers (with the exception of BladeCenter) are capable of running with either high voltage or low voltage sources. Some non-xSeries components may be restricted to low voltage power. BladeCenter must be connected to high voltage power. However, the ability to run a server on a single power supply at a low voltage may be limited. See Section 3.1.2.
- 3) Amperage rating of the power source. This is usually going to be one of the following:
 - a. 30A single phase (low or high voltage)
 - b. 60A single phase (high voltage)
 - c. 60A three phase (high voltage)
- 4) Connectors – need to know what we are plugging it into, and what is plugging into it. This is largely going to be determined by the volts/amps rating of our power source, and the power draw of our PDUs and servers.
- 5) Number and type of PDUs supported by the racks being used.

3.1 System x Power Requirements

3.1.1 System x Power Configurator

The Power Configurator can be used to determine the power requirements of each component in a configuration. A sample of the output is shown in the next section. This information is also available in most product announcements.

<u>Quantity</u>	<u>Description</u>	<u>Idle Power</u>	<u>Maximum Measured Power</u>	<u>Max Measured Input Current</u>	<u>Rated System Power</u>
1	xSeries 366 – x366 Sample (4) 3.00 GHz 2 x 2 MB L2 Cache Dual Core Xeon Processor MP 2 x 1300 W Hot-Swap (16) 4096 MB Dimm(s) (6) 73 GB 10K-rpm 2.5" SAS HDD (6) PCI Adapters	778 W	1088 W	5.2 A	1575 W
Power Estimates for Total Configuration					
Date & Time:		4/13/2006 10:46:01 PM			
Country:		United States			
Voltage:		208 V			
Based on system(s) running at:		<u>IDLE</u>	<u>MAX MEASURED</u>	<u>SYSTEM MAX</u>	
Power:		778 W	1088 W	1575 W	
Input Current:		3.7 A	5.2 A	7.9 A	
BTU/HR:		2653 BTU/Hr	3710 BTU/Hr	5371 BTU/Hr	

3.1.2 System x Power Considerations

To illustrate the importance of high voltage power, consider that a single x366 power supply is rated at 1300W when used with a 220V source. This single power supply is capable of powering a fully loaded server without redundancy. With 110V, this power rating for the single power supply falls to 650W.

A single power supply running on 110V can support roughly a maximum configuration of 1 microprocessor, 2 PCI-X adapters, 3 hard disk drives, and 4 DIMMs. If this configuration is exceeded, a second power supply must be installed or 220V input must be used.

If two power supplies are installed on 110V, the server will not support the above configuration with redundancy since if a one of the power supplies fails, the server may require 1088W but the single remaining power supply can only provide 650W. At 220V, there will be full redundancy since a single power supply at 220V can provide 1350W.

We are going to use the "maximum values" for our design. Why? The maximum values are what the power supplies are capable of pulling from the input circuits, and using those will ensure that any subsequent upgrades to an existing server won't overload our power solution. Additionally, *"electrical codes and regulations specify that all electrical equipment shall be used in accordance with listing and labeling requirements. Electrical service must be wired according to the system rating label, not the measured current."*

The max measured can be used for HVAC planning.

There is a new IBM Director plug-in called PowerExecutive. PowerExecutive enables you to monitor and collect power-consumption data from chassis and servers, manage trend data, and export data. Later this year, PowerExecutive 2.0 will also give clients the ability to cap the amount of power used by a single server or groups of servers to efficiently plan and utilize available power for maximum application performance.

3.2 Sample Solution Configuration

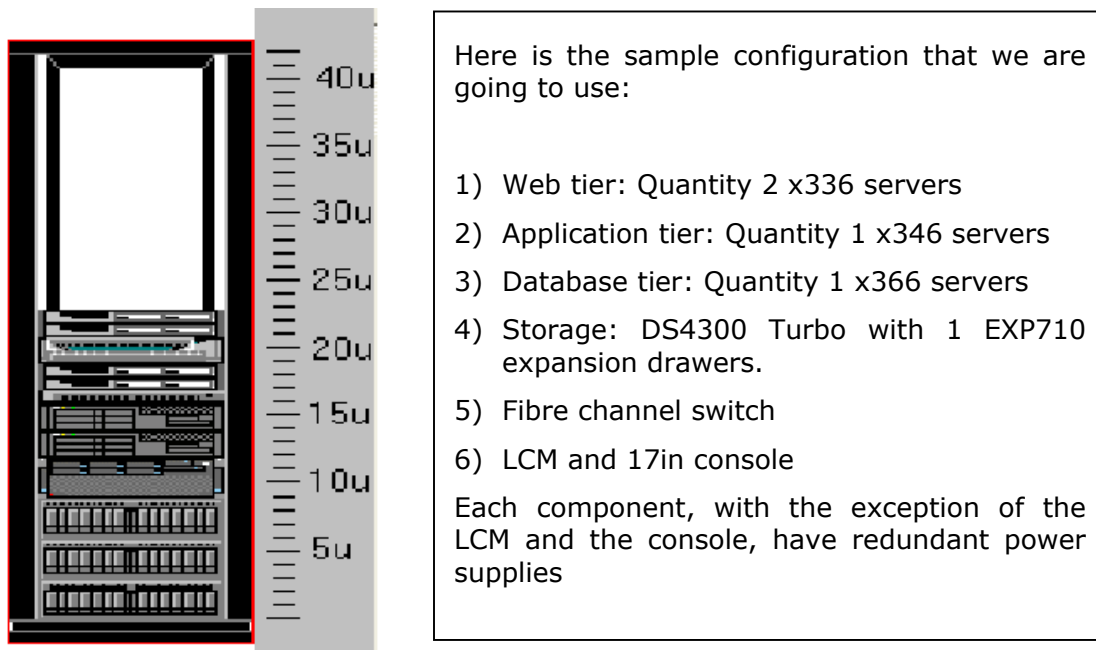


Figure 14 Sample Configuration

The values from the power configurator are shown in the tables below.

	Idle W	Measured W	Amps – Meas.	Max W	Amps – Max	BTU	# of cords
x336 Qty 2	554 (2x277)	844 (2x422)	4.0 (2x2)	1100 (2x550)	5.3 (2x2.6)	2872 (2x1436)	4
x346 Qty 1	278	434	2.1	830	2.0	1480	2
x366 Qty 1	572	799	3.8	1575	7.5	2725	2
DS4300	272	340	1.6	390	1.9	1159	2
EXP710 Qty 1	320	400	1.9	450	2.1	1364	2
Fibre Channel Switch	200	200	1	200	1	682	2
LCM/console	100	100	0.5	100	0.5	341	2
Total	3448	4795	14.9	7025	20.3	16339	16

Table 4 Power Requirements for Sample Configuration 1

Shown below are the configurations for 3 different power source types: 30A, Single Phase; 60A, Single Phase; and 60A, Three Phase. Not all possible arrangements are presented - these are to be used as guidance only.

First we will consider the 30A single phase solution. A pair 30A circuits are sufficient. Remember, the 30A circuit is derated by 20% and will provide us with 24A.

- 1) One choice would be to have 2x front end PDUs (32P1766/[39Y8939](#)) and connect 2 rack PDUs (32P1736/[39Y8951](#)) to each of these. This would look similar to Figure 11 DPI Front-end PDU Configuration on each side. When doing this, keep in mind that the x366 will require half of the amps available in the rack PDU it is connected to, so x336 and x346 should go into the second PDU.
- 2) This can also be done with 2x 30A Single Phase C13 Enterprise PDUs (26K4246/[39Y8941+40K9614](#)). These have twelve outlets each and 8 will be used 8 per PDU. This isn't a bad choice since most of the outlets are used and also most of the available amps have been allocated.

Next, let's look at 60A single phase:

- 1) Again this can be done with 2x C13 Enterprise PDUs, this time the 60A Single Phase version. (26K4242/[39Y8941+40K9615](#)). As before, most of the outlets on these PDUs are used, but in this case there are still have quite a few unused amps at the rack. We'll see in the next example that this is not the optimal approach.
- 2) A choice that gives us greater expandability would be to have 2x C19 Enterprise PDUs (26K4256/[39Y8948+40K9615](#)) to act as FE pdus for 4 rack pdus. (32P1736/[39Y8951](#)) This way we'll still have the ability to add additional rack PDUs in the future. This would look similar to Figure 13 on each side. We can also use the 60A FE PDUs (32P1767/ [39Y8940](#)) although future expansion is more limited.

With 60A, three phase power available:

- 1) We *could* do this with 2x 60A Three Phase C13 Enterprise PDUs (26K4244), but get your order in quickly as this PDU will not be available for much longer.
- 2) The better choice is to have 2x C19 Enterprise PDUs (26K4258/[39Y8923](#)) to act as FE pdus for 4 rack PDUs (32P1736/[39Y8951](#)). This way we'll still have the ability to add additional rack pdus in the future.

For the next example, we will expand upon our configuration for example number one. In this example, we increase the number of x336, x346 and EXP710 components:

- 1) Web tier: Quantity 4 x336 servers
- 2) Application tier: Quantity 2 x346 servers
- 3) Storage: DS4300 Turbo with 2 EXP710 expansion drawers.

	Idle W	Measured W	Amps – Meas.	Max W	Amps – Max	BTU	# of cords
x336 Qty 4	1108 (4x277)	1688 (4x422)	8 (4x2)	2200 (4x550)	10.6 (4x2.6)	5744 (4x1436)	8
x346 Qty 2	556 (2x278)	868 (2x434)	4.2 (2x2.1)	1660 (2x830)	8. (2x4.0)	2960 (2x1480)	4
x366 Qty 1	572	799	3.8	1575	7.5	2725	2
DS4300	272	340	1.6	390	1.9	1159	2
EXP710 Qty 2	640 (2x320)	800 (2x400)	3.8 (2x1.9)	900 (2x450)	4.3 (2x2.1)	2728 2x1364	4
Fibre Channel Switch	200	200	1	200	1	682	2
LCM/console	100	100	0.5	100	0.5	341	2
Total	3448	4795	22.9	7025	33.7	16339	24

Table 5 Power Requirements for Sample Configuration 2

Again, we will use the “maximum values” for our design.

Consider the 30A single phase solution. This cannot be done with only two 30A circuits since we need more than 24A to power all of the components, so 4 circuits are needed for full redundancy.

- 1) This can be done with 4x 30A Single Phase, C13 Enterprise PDUs (26K4246/[39Y8941+40K9614](#)).
- 2) Another choice might be to have 4 front end PDUs (32P1766/[39Y8939](#)) and connect 3 rack PDUs (32P1736/[39Y8951](#)) to each these. We need to go with a total of 6 rack PDUs since each rack PDU is limited to 15A.

Next, let’s look at 60A single phase:

- 1) This can be done with 2x C13 Enterprise PDUs, this time the 60A Single Phase version. (26K4242/[39Y8941+40K9615](#)). In this case, all of the available outlets are used on these 2 PDUs, yet there are still have quite a few unused amps at the rack, so this is not an optimal approach.
- 2) A better choice might be to have 2x C19 Enterprise PDUs (26K4256/[39Y8948+40K9615](#)) to act as FE PDUs for the 6 rack PDUs. (32P1736/[39Y8951](#)) This way there is still a way to add additional rack PDUs in the future. This would look similar to Figure 13 on each side. Using the 60A FE PDUs (32P1767/[39Y8940](#)) would also work, but similar to option 1, all of the outlets on the FE PDUs are allocated but there are still amps available.

Finally, with 60A, three phase power:

- 1) This *could* be done with 2x 60A Three Phase C13 Enterprise PDUs (26K4244), but get the order in quickly as this PDU will not be available for much longer. Plus, just like the previous example, all of the outlets are used and a lot of amps are unallocated.
- 2) The better choice is to have 2x C19 Enterprise PDUs (26K4258/[39Y8923](#)) to act as FE PDUs for 4 rack PDUs (32P1736/[39Y8951](#)). This way there is still the ability to add additional rack PDUs in the future.

Other considerations:

In most cases, the options for each category have similar list prices. The other thing to watch is whether we are going to fill up the available side pockets in the rack. All of the above configurations will fit in an IBM Enterprise Rack.

4 BladeCenter Power Considerations

The power supply modules in the BladeCenter unit must be connected to high voltage power only (between 200V and 240V AC). The BladeCenter unit comes with two power cords (C19/C20) to connect the BladeCenter unit to the power distribution units (PDU). Non-redundant power is not supported in BladeCenter products. Power modules must always be present in power bays 1 and 2. The additional power supply option which consists of a pair of power supplies also includes a pair of cables.

4.1 BladeCenter Power Domains

The BladeCenter chassis is divided into two power domains. The first two (standard) redundant power supplies (domain 1) provide, in the current maximum load scenario, 2250W of DC power to the blowers, management modules, switches, and blades 1 thru 6.

The second pair of redundant power supplies (domain 2) provides, in the current maximum load scenario, 2350W of DC power to Blades 7 thru 14. The power supplies are load balancing and are capable of providing up to 4000W DC power. Currently that much power never needed since if we look at the Domain 2 example above the max draw across the two power supplies doesn't exceed 2350 Watts which works out to less than 1200W per power supply.

Please make sure the power supplies are cross cabled as shown in Figure 15 Power Domains.




Figure 15 Power Domains

4.2 Designing Power Solutions for BladeCenter

If you take nothing else from this section: 2 x 30A circuits for 1 fully populated BladeCenter.

Here an example of the steady state power usage of a fully populated BladeCenter:

		Circuit 1	Circuit 2
Power Domain 1	Power Supply 1	~6A	
	Power Supply 2		~6A
Power Domain 2	Power Supply 3	~6A	
	Power Supply 4		~6A
	Total	12A	12A

Table 6 BladeCenter Power Requirements 1

As we did in Section 3.1, what power is available to be provided to the rack?

- 1) 30A, Single Phase
- 2) 60A, Single Phase
- 3) 60A, Three Phase

First, let's consider option 1 where we have 30A circuits available at the rack. Taking a look at the above table, it would appear we could use two 30A circuits to supply not just one but two BladeCenters. However, if there is an issue with Circuit 2, we now have this:

		Circuit 1	Circuit 2
Power Domain 1	Power Supply 1	12A	
	Power Supply 2		0A
Power Domain 2	Power Supply 3	12A	
	Power Supply 4		0A
	Total	24A	0A

Table 7 BladeCenter Power Requirements 2

And we've immediately maxed out Circuit 1.

Figure 16 shows the power distribution for a single BladeCenter chassis with 30A single phase power sources. Power sources A and B represent separate AC power feeds to the rack. The 24A PDU is the front end PDU 32P1766.




Figure 16 BladeCenter with 30A Circuits

Next, let's look at 60A single phase:

1) We can do this all with 2 Enterprise PDUs (26K4256).




Figure 17 Two BladeCenters with 60A Single Phase Power

If we are fortunate to have 60A, three phase power available, the PDU configuration will be relatively straightforward. Figure 18 is the preferred power distribution for three BladeCenter chassis of any power. Power sources A and B represent separate AC power feeds to the rack. The 60A 3-P PDU is part number the high density PDU 73P5847.

This can also be done 2 Enterprise PDUs (26K4258). If we have any other components in the rack that require a C13 outlet, we add rack PDUs to the configuration and connect them to the ePDU.




Figure 18 Three BladeCenters with 60A – Three Phase power

4.3 Power Management in BladeCenter

If there is a power related failure and a single 2000W power supply is called upon to do all the work, then the BladeCenter Management Module will manage the power situation. If the domain that is operating with a failed power supply is drawing less than 2000W, then nothing needs to be done. However, if the impacted domain is drawing more than 2000W, the three possible scenarios are available:

- **Redundant with performance impact:** Blades will be powered even if there is not enough power to fully power them with the loss of a power supply. In the case of a failure, Certain blades will be directed to “throttle down” or temporarily kick themselves down to a lower clock speed so that all blades continue to function until the redundant power supply is back online. **This is the standard configuration.**
- **Redundant without performance impact:** A new blade will not be powered unless it can continue to run at full performance even with a power supply failure. For customers that are insistent on no possible performance loss this is the option to choose.
- **Non-redundant mode:** Blades are powered without consideration of power supply redundancy. All blades will power on – in the event of a power loss some blades may go down until chassis is at a supportable power level




Figure 19 Power Management Policy Settings

Here are 2 power examples that show how this power management helps preserve a datacenter's power infrastructure:

Without power management:

- The max AC draw could be as much as 2650W / power supply
- 2650W / 200V is 13.3 Amps
- A 30A PDU is UL derated by .8 and provides up to 24A. Since 13.3A/power supply x 2 is greater than 24Amps, we would not be able to use a 30A PDU in this scenario.

With power management:

- Maximum AC power draw is managed to be no more than 2220W / power supply.
- 2220W / 220V is 11.10 Amps.
- A 30A PDU is UL derated by .8 and provides up to 24A. So 11.10A/power supply x 2 is less than 24Amps so the BladeCenter can be run on 30A circuits. It also allows us to place 4 instead of 3 power supplies on a 60A 1phase (48A UL derated) PDU.

Figure 20 shows the warning message that will be seen in the Management Module Event Log in an oversubscribed situation.

Index	Sev	Source	Date/Time	Text
1	I	BLADE_02	03/20/05, 12:41:55	(SN#ZJ1V5T4461LF) Blade Server Powered Down
2	I	SERVPROC	03/20/05, 12:41:53	Recovery Demand exceeds a single power module. Throttling can occur in power domain 1.
3	W	SERVPROC	03/20/05, 12:40:21	Demand exceeds a single power module. Throttling can occur in power domain 1.
4	I	BLADE_02	03/20/05, 12:40:21	(SN#ZJ1V5T4461LF) Blade Server Powered Up

Figure 20 Throttling warning message

Note: There is no way prioritize Blade Throttling. The MM goes through an algorithm that gets the blades under the power budget by throttling blades as soon as a power supply failure detected.

This approach saves money in the data center. In an N+N data center, one can never exceed what a single power source can give. So for a customer with a 200KW (2x100KW) UPS, they can only run 100KW of server gear. With this power management, customers can exceed the 100KW limit if they understand that, at times of a rare power crisis, they will have to accept some performance throttling. Most customers have welcomed and accepted this approach once they understand the benefits. This IBM breakthrough has unlocked hidden power assets in the data center.

4.4 BladeCenter H

4.4.1 BCH Power Introduction

The BladeCenter H holds up to 4 2900W power supplies.




Figure 21 BladeCenter H

Blue Power Supplies	Red Power Supplies
<ul style="list-style-type: none"> • Attached to first line cord via the right inlet (as viewed from the back of the chassis) • Maximum draw from each 2900W power supply is 38A • Two left supplies operate as N+N Redundant Pair • Provides redundant power for Blades 1-7 + blower • Also provides shared power for common infrastructure in chassis (MM, switching) 	<ul style="list-style-type: none"> • Attached to second line cord via the left inlet (as viewed from the back of the chassis) • Maximum draw from each 2900W power supply is 38A. • Two right supplies operate as N+N Redundant Pair • Provides redundant power for Blades 8-14 + blower • Also provides shared power for common infrastructure in chassis (MM, switching)

4.4.2 BladeCenter 2900W Power Supply

The IBM BladeCenter H 2900W power supply is unique to the BC-H and is not compatible with the other BladeCenter chassis. BladeCenter H also has an improved design for powering the chassis. Instead of a power cord for each individual power supply, each pair of supplies is joined internally and there is one cord connector for each domain. The rear of the chassis contains two Burndy connectors. (see item 10 in Section 2.1).



Figure 22 BladeCenter H Power Supply

Because there are two ways to power the chassis (direct or via PDU) BladeCenter H does not ship with any power line cords, they must be ordered separately.

- BladeCenter H can be plugged into an IBM Enterprise PDU with C19 outlets using the Triple C20 power cord option. For proper redundancy, input currents requirements must be met.
- BladeCenter H can also be connected directly to a properly sized and protected outlet (without the need for PDUs). Only single phase power is supported by the power supplies. For three phase power sources, properly designed PDUs must be used.

The power supplies are front loading for easier service. No cable removal is required for replacement or service of a supply

Each option kit includes a single line cord. Cords can be mixed for applications that require PDU and UPS, or direct wall power and UPS.




Figure 23 BladeCenter H Power Connections

The connector on the back of the BladeCenter assures that the cable can not be installed incorrectly. These cables work in the same fashion as the connectors on many PDUs in the IBM PDU family.



Part Number	Description	Picture	Where Used
25R5785	Triple 16A IEC 320-C20 (200-240V) 2.8M		This cable is required for use with IBM PDUs and UPS. It will allow connection to any C19 based PDU/UPS. This cable will be used for larger configurations where PDUs make sense or battery back power conditioning is needed via a UPS
25R5783	Double 30A L6-30P (208V) 4.3M		This cable contains two L6-30 plugs and will allow direct connection to a 208V 30A NEMA wall outlet. This cable will be used for small to medium installations or for customer that can not provide larger power feeds to the rack.

Table 8 BladeCenter H Line Cords

4.4.3 Wiring Examples for BCH

Consolidated with PDU - Redundant

Single chassis configuration: 200-208V Single Phase 60A PDU connection redundant solution, outlets are IEC 309 60A single phase, PDU is IBM part number 32P1767, 26K4256, or 73P5855




Figure 24 BCH with 60A Single Phase

Multiple chassis: 200-208V Three Phase 60A direct connection redundant solution, outlets are IEC 309 60A three phase, PDU is IBM part number 26K4258 or 73P5847

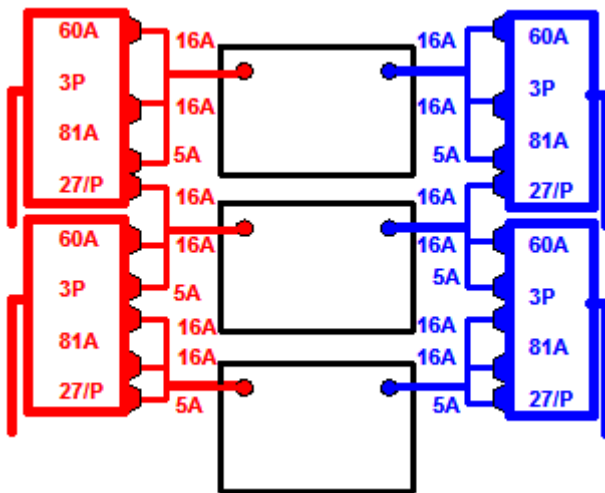


Figure 25 BCH with 60A Three Phase

Looking closely at the diagrams above, you'll see that two of the C20 connectors are pulling up to 16A and the third one, which directly powers the blowers, only pulls 5A. A 60A-3P PDU can provide 3 x 27A of power. In this example, we need to make sure that the connections are split between the phases correctly: PDU1: 16A, 16A, 21A; PDU2: 16A, 21A, 21A.

Direct, Redundant Power

Top: Single chassis with 208V Single Phase 30A direct connections, redundant solution. Outlets are L6-30.

Bottom: Multiple chassis with 200-208V Single Phase 30A direct connections, redundant solution. Outlets are L6-30

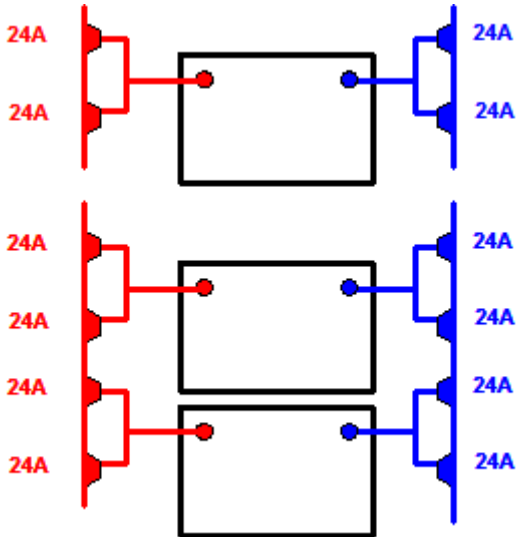


Figure 26 BCH with Direct Connections

4.4.4 FAQ

- If I'm using two 60 Amp FE Enterprise PDUs and I lose one do I lose a whole domain?
 - No, with all four power supplies installed, if you lose an entire circuit (DPI cable fails - Triple C19 or L6-30P twin-tail, or FE PDU), the chassis will stay up with both domains running in non-redundant mode.
- If I'm using the L6-30P twin-tail, and I lose one of the 30 Amp circuits, do I lose the entire Power Domain, or does one supply stay up? I guess the question is: how are the 2 power supplies and blower wired to the twin-tail L6-30P plugs?
 - One of the plugs powers the blower and one of the power supplies. The other L6-30 is attached to the second power supply. So if you lose one of the L6-30s, you will only lose one power supply in the domain. The other power supply in the domain will remain up.
- If I'm using lower voltage blades, do the power supplies step down in power required so I don't need so much power? If I know I'm using lower power blades, well under the 2900 watts, do the power supplies "step down" and use less amperage, and thus I can get buy with less power (like can I have two 30 A circuits, two 220V FE PDUs with 3 C19s, and just plug in the triple C19 to the FE PDU for each power domain? Instead of having to have 4 30A or 2 60Amp 3 phase PDUs?
 - The power draw is dictated by the options and blades not by that 2900W power supply. You can certainly design a fully redundant power solution that uses less power than we show in the PDU guide released.

5 UPS

Now that we have an understanding of power, how it is distributed and what the requirements are of our target systems, we can take a look at how we can provide uninterrupted power to these devices.

The table below shows the low- and high-voltage rack-ready UPS models from IBM. Each model also supports the addition of external battery packs to extend the amount of run time provided by the UPS. For more detail on these UPS models and estimating run time for the configuration, please see the Configuration and Options Guide (link in Appendix D).

Model	Part Number	Size	Line Cord	Power Outlets	Rated Output	Battery Pack
3000XLV	2130R30 21301RX	2U	Fixed with L5-30 connector	6 IEC C13 & 1 L5-20R	2850W (24A)	2U; 32P1692, 40K9620
3000XHV	2130R31 21302RX	2U	HV C19/C-20, L6-20p	7 IEC C13 & 1 IEC C19	2850W (16A)	
7500XHV	21086RX 21303RX	6U	208-250V Hard wired	Four IEC C19	6000W (38A)	3U; 25R5582, 39Y8857
10000XHV	21088RX 21304RX	6U	208-250V Hard wired	Four IEC C19	8000W (50A)	

Table 9 UPS Specifications

As indicated above, the 7500XHV and 10000XHV must be hardwired directly into the room power source. There is no plug included in the ship group. See 1.Appendix B for more information on hardwiring these units. This work must be done by an electrician.

5.1.1 BladeCenter and 7500 XHV or 10K XHV UPS

The recommended approach is to use the 10K UPS when design a UPS solution for the BladeCenter. In a comparison of the 10K UPS vs. the 2x3000XHV approaches, the results are very similar. All other things being relatively equal, the 10K provides more capacity and better expandability while using 2U less space, so that is the recommended choice.

This setup should provide about 12 minutes of runtime for the current configuration (1 BC, x366, DS4300). This extends to over 30 minutes with the addition of a 3U battery pack (total of 9U). The downside to this configuration versus the 3000XHV is that the 10K UPS has hardwired connection and an electrician required to complete this, but 2x 208V circuits should suffice.

Redundant, UPS for BladeCenter, etc. 60A + 10K UPS

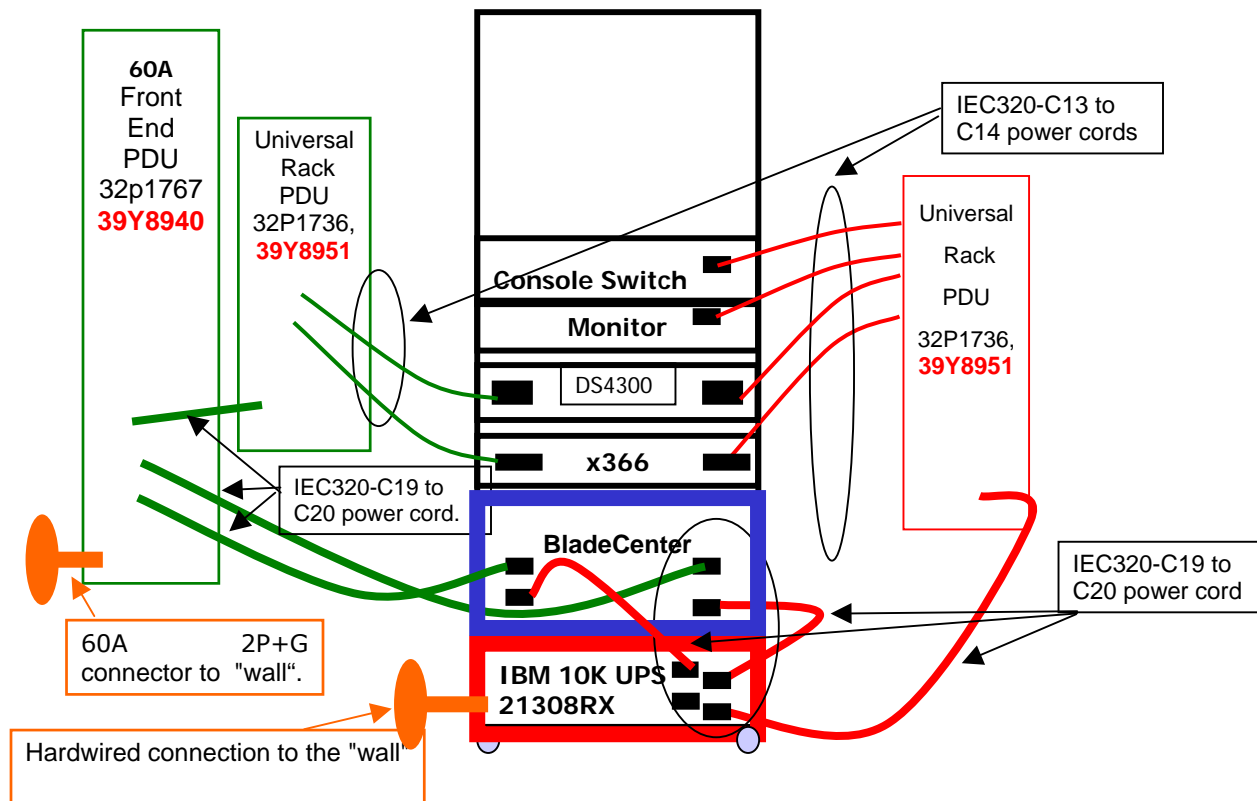


Figure 27 BladeCenter and 10K UPS – Sample Configuration

Here is the part list and bill of materials:

21308RX	IBM UPS 10000XHV	Qty 1
32P1736/39Y8951	IBM DPI Universal Rack PDU	Qty 2
32P1766/39Y8939	IBM DPI 30A/208V US Front End PDU	Qty 1

Cables:

Front End PDU (32P1766) will plug directly into the wall. The connector is an L6-30P and the cable is included with the PDU.

10K UPS will plug directly into the wall. 10K does not come with a line cord. There are terminal blocks at the back of the UPS that must be connected by an electrician. One way to do this is to have a pigtail built with a standard wall connector on the other end.

- 1 C19/C20 cable: FE PDU to Rack PDU - included with Rack PDU
- 1 C19/C20 cable: 10K UPS to Rack PDU - included with Rack PDU
- 4 C19/C20 cables: 2 connect to FE PDU, 2 connect directly to 10K UPS - all 4 included with BladeCenter
- 6 C13/C14 cables: 1 to each Rack PDU from DS4100 and x366, one from monitor kit to Rack PDU (UPS side), one from console switch to Rack PDU (UPS side)- all included with rack mounted components.

Note: Four IEC power cables are included with the 10K UPS to attach devices or PDUs to the UPS outlet, which are 2m long with IEC 320-C19/C20 connectors.

5.1.2 BladeCenter and UPS 3000 XHV

It is possible to configure a UPS solution for BladeCenter with with 2 x 3000XHV. Note, the COG has a notation that only the 7500 or 10K should be used with the BladeCenter. The 3000XHV has only 1 C19 outlet, but with 7 C13 outlets, we don't need the Rack PDU for the UPS side.

A benefit of this approach versus the 10KXHV is that no electrician required since the 3000XHV have C19/L6-30P cables to connect to wall. However, 4x 30A 208V circuits are needed for this setup. One battery pack per UPS (total of 8U) would provide roughly 30 minutes of runtime with current configuration. Without the battery packs, we'll only get about 5 minutes.

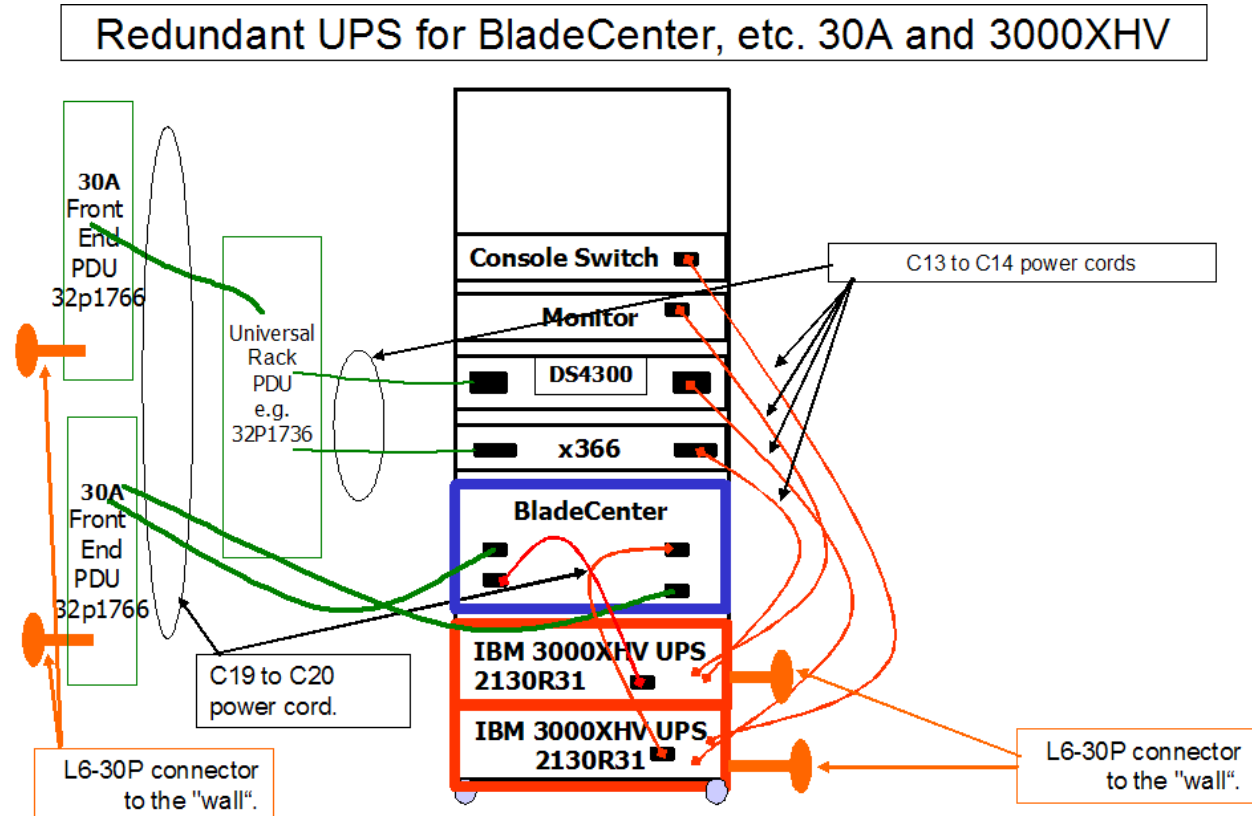


Figure 28 BladeCenter and 3000XHV – Sample Configuration

Here is the part list and bill of materials:

21308RX	IBM UPS 3000XHV	Qty 2
32P1736/39Y8951	IBM DPI Universal Rack PDU	Qty 1
32P1766/39Y8939	IBM DPI 30A/208V US Front End PDU	Qty 1

Cables:

Front End PDU (32P1766) will plug directly into the wall. The connector is an L6-30P and the cable is included with the PDU.

The 2 3K UPS will plug directly into the wall using the C19 – L6-30P provided with the UPS.

- 1 C19/C20 cable: FE PDU to Rack PDU - included with Rack PDU
- 4 C19/C20 cables for BladeCenter: 2 connect to FE PDU, 1 each to connect directly to each 3K UPS - all 4 included with BladeCenter
- 6 C13/C14 cables: 1 each to Rack PDU, and to UPS from DS4100 and x366, plus monitor kit to UPS and console switch to UPS.

6 Cooling

"So-called blade servers can be packed into cabinets like books on a shelf. But this configuration generates more heat, increasing the demand for energy to run cooling systems. In fact, 50% of a center's monthly energy bill can go for cooling."

-"Servers As High As An Elephants Eye", Business Week, June 12, 2006

The basic rule in the datacenter is that **POWER IN = HEAT OUT**. So as computers require more power, they also require more cooling. So the above quote is essentially stating the obvious, without accounting for any efficiencies in the power and cooling systems themselves.

The BTU is a measurement of heat output from a device. The AC wattage x 3.41 = BTU where 3.41 is a constant value.

- x346 AC draw is 523W max BTU is therefore 1783

Based on the heat output, or BTUs of a server, it is possible to calculate the tonnage of AC needed. The BTU divided by 12000 = AC tons to cool.

An example with BladeCenter 2000W power supplies

- 2600W per domain is theoretical max (5000/chassis)
- 5000W (entire chassis) x 3.41 = 17,050BTU
- 17,050/12,000 = 1.42 tons

As mentioned previously, the growth in power requirements for servers is not coming solely from the CPUs. Servers with more function typically mean more power drawing component including larger hard drivers, more memory, and more PCI and I/O functions. Add to that greater rack level density, especially with blade based solutions, and your physical planning team is a nervous wreck.

IBM servers are designed to work with front to back airflow. Inlet air temperatures must be maintained at or below 35C temperature and 8-80% relative humidity which is actually slightly higher than other server vendors due to our enhanced airflow dynamics within the servers. Add to that the other basic physical requirements such as plenty of space in front and rear as well as over head to handle thermal physics and enough CFM of air to cool the server.




Figure 29 Front to back cooling

Not every data center in the world can house full density servers today. What is limiting density in the data center?

- Real estate in the rack (easy)
- Power input (manageable)
- Thermal output (complex)
- Airflow requirements (can be much harder)

Physics dictates that to maintain no more than 30 degrees temperature rise across a server, each kilowatt of power requires 120 Cubic Feet per Minute (CFM) of cool air moving through the server. So if a cabinet is running at 20 Kwatt of power it requires $(20 \text{ Kwatt} \times 120 \text{ CFM/Kwatt}) = 2400 \text{ CFM}$ of air.

The higher the server is located in the cabinet, the more difficult the cooling problem, as cool air is typically supplied from the floor of the data center. **For example:** Imagine a server rack that has 600cfm of cool supply air available, and is populated with 10 servers, each of which draws 100cfm at its air intake. The lowest server draws 100cfm of the rising air, leaving 500cfm to cool the remaining nine servers. The next server up in the enclosure draws another 100cfm, leaving 400cfm to cool the remaining eight servers, and so on. Obviously, in this example, there will be no cool air left for the seventh through tenth servers at the top of the rack. Typically, the remaining servers simply draw air from whatever source possible, often times the hot exhaust from and around the ceiling.




Figure 30 Hot at the top of the rack

Appendix A. Connectors




Figure 31 L6-30P




Figure 32 IEC309 2P+Gnd




Figure 33 IEC309 3P+Gnd



Figure 34 IEC309 3P+Gnd Connector

Appendix B. Hardwired UPS Connections

Figure 35 and Figure 36 shows what a hardwired connection for the UPS7500XHV looks like. More detail on what is required can be found in the *IBM UPS 7500XHV IBM UPS 10000XHV Uninterruptible Power Supply Operation and Setup Guide*. (ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers_pdf/25r5636a.pdf)



Figure 35 UPS 7500 Hardwired



Figure 36 Hardwired Wall Connection (bottom box)

Appendix C. Power Transmission

Power travels from the power plant to the data center or your house through the power distribution grid. Commercial electrical generators generate 3-phase AC power.

AC has at least three advantages over DC in a power distribution grid:

1. Large electrical generators generate AC naturally, so conversion to DC would involve an extra step.
2. Transformers must have alternating current to operate, and we will see that the power distribution grid depends on transformers.
3. It is easy to convert AC to DC but expensive to convert DC to AC, so if you were going to pick one or the other AC would be the better choice.

The power plant produces three different phases of AC power simultaneously, and the three phases are offset 120 degrees from each other. There are four wires coming out of every power plant: the three phases plus a neutral or ground common to all three.

In order to reduce line, typical voltages for long distance transmission are in the range of 155,000 to 765,000 volts. As it reaches the local delivery points, it is transformed to lower voltages. The power loss across a wire is equal to $I^2 \cdot R$, so by running these transmission lines at high voltage, and low amps, the amount of power lost to heat is reduced




Figure 37 Transmission and Distribution Grid

For power to be useful in a home or business, it comes off the transmission grid and is stepped-down to the distribution grid. This may happen in several phases. The conversion from the transmission grid to the distribution grid occurs at a power substation. A power substation has transformers that step down transmission voltages to distribution voltages, typically less than 10,000 volts.

At this point, the typical line voltage at something like 7,200 volts running through the neighborhood on three wires with a fourth ground wire lower on the pole. (Figure 38)

A house needs only one of the three phases, so typically you will see three wires running down a main road, and taps for one or two of the phases running off on side streets.



Figure 38 Three Phase Distribution



Figure 39 Transformer

And finally we are down to the wire that brings power to your house! Past a typical house runs a set of poles with one phase of power (at 7,200 volts) and a ground wire (although sometimes there will be two or three phases on the pole, depending on where the house is located in the distribution grid). At each house, there is a transformer drum attached to the pole. The transformer's job is to reduce the 7,200 volts down to the 240 volts that makes up normal household electrical service.

There are two wires running out of the transformer and three wires running to the house. The two from the transformer are insulated, and the third one is bare. The bare wire is the ground wire. The two insulated wires each carry 120 volts, but they are 180 degrees out of phase so the difference between them is 240 volts. This arrangement allows a homeowner to use both 120-volt and 240-volt appliances.

Appendix D. Related Publications and Tools

Related Publications

Related publications include the following:

- *Configuration and Options Guide (COG)* - The COG contains general configuration information about IBM products and supported options, such as IBM System x, BladeCenter, xSeries, IntelliStation and other options. External PDF version: www.ibm.com/support/docview.wss?rs=1201&uid=psg1SCOD-3ZVQ5W. The COG is now available in HTML: <http://www.ibm.com.com/servers/eserver/xseries/cog/>
- *Decision Makers Guide to Power Distribution Units*: This guide provides an introduction to the IBM DPI family of PDUs and explains various factors to be considered when designing the layout of your power system. (<http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-46119>)
- *Decision Makers' Guide to selecting the right 110V-208V uninterruptible power supply* - (<http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-46116>)
- *IBM BladeCenter Type 8677 Planning and Installation Guide*: This manual provides physical planning information for the IBM BladeCenter products. It describes the BladeCenter components, explains deployment and installation considerations, and provides worksheets that you can use to determine the configuration, power, weight, and cabling requirements for your BladeCenter unit. (<http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-53670>)

Related Tools

- *Power Configurator* - IBM has released a Power Calculator for accessing accurate power and thermal information for BladeCenter and System x servers. This tool is available from ibm.com at the following link:

<http://www-03.ibm.com/systems/bladecenter/powerconfig/>

- *IBM PowerExecutive Download*: (<http://www-1.ibm.com/support/docview.wss?uid=psg1MIGR-62411>) *IBM Power Executive Installation and User's Guide* (<http://www-1.ibm.com/support/docview.wss?rs=0&uid=psg1MIGR-62399>)
- *Rack Configurator* - While this Power Guide provides a basic understanding of the way in which IBM DP Interconnect PDUs are used and installed, the complexity of today's power solutions and the variety of sizes, access options and cabling requirements of servers, accessories and storage components make configuring a complex networked system a challenge. The IBM Rack Configurator is an easy-to-use tool that helps you design the optimum layout. The configurator helps check, correct and report the following:
 - Components by product number and position
 - Infrastructure specifications, including weight, power, volt-amps (VA), heat (BTU/hr), bays, EIA, outlets and console ports
 - Width and depth; front, rear and side clearances; total weight and top clearance
 - Cables and connectors by component position

The configurator software can be downloaded from <http://www-1.ibm.com/servers/eserver/xseries/storage/rack.html>. The configurator is a good source of accurate information on designing a redundant, reliable power distribution topology. Individual User and Installation Guides also provide detailed information on the use and mounting of for DPI PDU products.

Index

BladeCenter	NEMA 5-15.....	11
Power Configuration	Cooling	32
Power Domains.....	Line Cords.....	16
Throttling.....	Ohm’s Law	5
UPS	PDU	
BladeCenter H	DPI	14
Power Configuration	Enterprise	15
Power Cords.....	Front End.....	14
Power Domains.....	RoHS Transitions.....	16
Power Supplies	Summary.....	16
BTU	Universal Rack.....	12, 14
CFM	PowerExecutive	18, 39
Connectors	Rack Configurator	39
Burndy	Racks	13
C13/C14.....	Related Documents	39
C19/C20.....	System x Power Configurator.....	17, 39
IEC309 2P+G	Three Phase Power.....	9
IEC309 3P+G	UPS	29, 35
L6-30P	and BladeCenter	29