Raritan
A brand of legrand

The iPDU Handbook

A Guide to Intelligent Rack Power Distribution

The iPDU Handbook

A Guide to Intelligent
Rack Power Distribution

Raritan began developing KVM switches for IT professionals to manage servers remotely in 1985. Today, as a brand of Legrand, we are a leading provider of intelligent rack PDUs. Our solutions increase the reliability and intelligence of data centers in 9 of the top 10 Fortune 500 technology companies. Learn more at Raritan.com

V1195R2

Table of Contents

1.0 **Introduction**

The rack power distribution unit (rack PDU) has emerged from obscurity. As the last link of the elaborate data center power chain, the traditional role of the rack PDU has been to deliver stable, reliable and adequate power to all the devices in the rack or cabinet—servers, storage, network equipment—which are plugged into it. Although it provides the electrical heartbeat to all the systems that run the critical applications that support the operation of the business (or that, in some cases, are the business); it was often considered a simple commodity, just a power strip. Typically, IT merely told facilities how much power was needed, based on device nameplate specs—and often with redundancy, so there was plenty of headroom and minimal risk of downtime. Little thought was given to efficiency or what other value a rack PDU could provide. That was yesterday.

Over the past few years, system availability has become a "given" and data center management attention is now being focused on operational costs, efficiency improvements and resource optimization. With the annual expenditure for powering the average data center surpassing the cost to purchase the IT equipment itself, the use (and waste) of energy is now targeted as a priority. Beyond the actual cost to power the data center there are the related issues that impact both current operations and future expansion—e.g. physical space and utility power availability, CO2 footprint and potential government regulation. Since almost all of the power delivered from the utility to the data center is consumed either directly by the devices plugged into rack PDUs or indirectly by the infrastructure to bring power to the rack and cool the devices, the once obscure rack PDUs have become visible on the data center management radar.

Not surprisingly, many of the major strategies to address the above issues and improve overall data center efficiency depend on new capabilities not available in the commodity outlet strips of a few years ago.

To consider a few of these capabilities:

- In order to maximize the use of data center space and other resources there has been a trend to deploy racks densely packed with 1U servers or power-hungry blade servers. Today's rack PDUs typically handle loads of 5-10 kW with 20 outlets and there are PDUs now designed to support 20+ kW and 36 or more outlets.
- To increase IT staff productivity and conserve energy by employing lights-out and/or remote data center operation, some rack PDUs provide real-time monitoring, reporting and alerts as well as secure, reliable outlet switching.
- To identify ghost (no definable function), underutilized or grossly inefficient servers for possible elimination, replacement, consolidation or virtualization, some rack PDUs provide individual outlet monitoring.
- To create individual awareness, accountability and/or chargeback for power usage and CO2 footprint, some rack PDUs are equipped with highly accurate, real-time power measurement capabilities at the PDU and outlet level.
- To optimize IT workload and make informed decisions for infrastructure capacity planning, IT and facilities managers need rack PDU management software that continually collects data on power consumption, analyzes trends and correlates with IT workload data.

Why Has the Selection of Rack PDUs Become So Important?

As you can imagine, a wide variety of rack PDU configurations is available based on parameters such as: number of phases, voltage, total amps, branch circuits, number of outlets, socket type, plug type, rack units consumed, and physical dimensions. Beyond the functions of the basic rack PDU, additional capabilities are available in rack PDU categories, we call metered, switched, and intelligent. Furthermore, if you cannot find an off-the-shelf rack PDU that matches your specific requirement, some vendors will assemble or even design a custom rack PDU (also called BTO/ETO: built-to-order/engineered-to-order).

In the following sections we will discuss the basic concepts, considerations and approaches in designing, selecting and deploying the appropriate rack PDU for typical data center applications. We will describe not only the basic ingredients for delivering adequate, reliable power to the rack, but also the factors and best practices that will contribute to a reliable, operationally efficient and environmentally sound data center for today and for the future.

2.0 **Fundamentals and Principles**

2.0 **Fundamentals and Principles**

IT equipment is normally mounted in racks or cabinets with provisions for all necessary cables, ventilation, cooling and convenient access. There are large data center PDUs that are used earlier in the power chain and take the form of panel boards mounted on walls or free standing pedestals. In this handbook we're discussing only the rack PDU, at the end of the chain, which supplies power to the IT equipment in the rack. Unless otherwise stated any reference to "PDU" for the remainder of this handbook means "rack PDU."

Rack PDUs come in many configurations with respect to number and type of receptacles, voltage, load capacity and physical mounting (horizontal or vertical). A unit may perform no function other than providing power to the devices plugged into it; or it may also provide additional functions—for example, turning power off and on remotely, monitoring power consumption and sensing the temperature in an IT equipment rack.

2.1 Overview and Class of Devices

A rack PDU is mounted in an IT equipment rack and provides electrical power to various IT devices such as servers, networking, and storage equipment. Today, rack PDUs are available in a number of configurations. We describe below the basic characteristics of four types of rack PDUs.

2.1.1 Types of Rack PDUs

Rack PDUs can be divided into two categories: Non-Intelligent PDUs and Intelligent PDUs

Non-Intelligent PDUs

- **Basic PDUs:** Are power strips that are used in critical environments such as data centers. They distribute correct voltage and current to multiple outlets to power IT equipment in racks.
- **Monitored PDUs:** Allow a user to view a local display that typically provides electric current information only. However, this information cannot be accessed remotely as the units have no network connectivity capabilities.

Intelligent PDUs

- **Metered Input PDUs:** Meter power at the PDU-level and can display the data both locally and over a network. Metering helps users to determine power usage and available capacity at the rack and facilitates provisioning. By metering at the input level, users can avoid overloading circuits and more easily calculate efficiency metrics like Power Usage Effectiveness (PUE).
- **Metered Outlet PDUs:** Meter power at the outlet-level and can display the data both locally and over a network. Like metered input PDUs, outlet-metered models help users to determine power usage and available capacity at the rack, and also facilitate provisioning. Most importantly, outlet-level metering allows users to understand power consumption at the device or server-level which makes it possible to allocate costs to specific business units or customers.

- **Switched PDUs:** Offer the features of Metered Input PDUs and also provide controlled on/off switching of individual outlets. They enable authorized users to securely power cycle devices remotely in a specific order, offer power sequencing delay to minimize inrush currents, prevent unauthorized device provisioning, and can power off devices that are not in use to conserve energy.
- **Switched PDUs with Outlet Metering:** Combine all of the capabilities of Switched PDUs with those of Outlet -Metered PDUs.

	Non-Intelligent PDUs		Intelligent PDUs			
Core Features	**Basic**	**Monitored**	**Metered Input**	**Metered Output**	**Switched**	**Switched and Outlet Metered**
Power Distribution	●	●	●	●	●	●
Input Metering	○	●	●	●	●	●
Outlet Metering	○	○	○	●	○	●
Network Connectivity	○	○	●	●	●	●
Switching	○	○	○	○	●	●
Secondary Features						
Environment Sensor Support	○	○	◐	◐	◐	◐
Strong Passwords	○	○	◐	◐	◐	◐
Encryption	○	○	◐	◐	◐	◐
User Permissions	○	○	◐	◐	◐	◐
Additional Ports, e.g. USB	○	○	◐	◐	◐	◐

● Always ◐ Sometimes ○ Never

Figure 1: Types of Rack PDUs

(Source: Raritan, Inc.)

To learn more about intelligent PDUs, visit: www.raritan.com/ipdus

2.2 Electrical Terminology

It will be helpful to review several electrical term definitions before discussing rack PDUs in more detail.

Voltage (Volt): Electromotive force, or difference in electrical potential, measured in volts and equal to the current times the resistance.

Current (Amp): The flow or rate of flow of electrons, ions, or holes in a conductor or medium between two points having a difference in potential, measured in amperes and equal to the ratio of the voltage to the resistance.

Active power (Watt): The real power drawn by an IT device. This determines the actual power purchased from the utility company and the heat loading generated by the equipment since, for IT equipment, 1W of electricity equals 1W of heat.

Apparent power (VA): The product of the voltage applied to the IT device times the current drawn by the equipment. The VA rating is used for sizing wiring and circuit breakers. The apparent power is always equal to or larger than the active power.

Power factor: The ratio of active power to apparent power.

Energy measurement (kW): Electrical energy is measured at one instant in time in watts or kilowatts. For example, a 100W light bulb consumes 100 watts at any instant in time, but energy is consumed, and billed for by utilities, over time. A kilowatt-hour (kWh) is a unit of electrical energy or work, equal to the power supplied by one kilowatt for one hour, e.g. a 1000W light bulb left on for one hour, or a 100W light bulb left on for ten hours.

Line: An electrical conductor which is a source of voltage, e.g. 120V. In a single-phase system there are one or two lines. In a three-phase system there are three lines.

Neutral: An electrical conductor that provides a return path for the voltage supplied by a line. The neutral itself is not a source of voltage.

Ground: A conducting body, such as the earth or an object connected with the earth, whose potential is taken as zero and to which an electric circuit can be connected. The purpose of a ground wire is to safely direct stray currents to ground rather than allowing them to pass through someone contacting the stray currents.

4-wire and 5-wire systems: A 4-wire rack PDU consists of one ground wire and three lines (see Three phase Delta below), each line carrying equal voltage but each voltage sine wave is 120 degrees out of phase with the others. The voltage of two lines is available (line to line, e.g., L1-L2). A 5-wire system is the same as the 4-wire system but with the addition of a neutral wire (see Three phase Wye below) so that the voltage of one line can be supplied (line to neutral) as well as the voltage of two lines (line to line).

Three-phase Delta "Δ": This configuration gets the name Delta because a schematic drawing of it has three transformers forming a triangle or the Greek letter Delta. The three lines connect to the three "corners" of the triangle.

Three-phase Wye "Y": This configuration gets the name Wye because a schematic drawing of it has three transformers meeting in the center forming the letter "Y". The three lines connect to the three "branches" of the "Y" and the neutral connects to the center.

2.3 Electrical Power Distribution to the Rack

2.3.1 Branch Circuits

Power is distributed to the rack over one or more electrical branch circuits. Branch circuits are power feeds that originate from a panel, switch or distribution board and terminate into an electrical receptacle mounted in a junction box near the IT equipment rack. Depending on the data center's layout, branch circuit wiring can be overhead, underneath a raised floor or both. The rack PDU itself could have multiple branch circuits. See Section 2.6 for details regarding branch circuit protection requirements.

2.3.2 Branch Circuit Load Capacity

The power that can be delivered by a branch circuit depends on the electrical characteristics of the circuit. A key factor in delivering power to a rack is whether the power is single phase or three phase. The amount of electricity delivered to a rack is often referred to as the load capacity and is the product of the rated voltage and the rated current and is presented as Volt-Amps (VA) or kVA (VA x 1000). Given the rated voltage and current, the load capacity that can be delivered by a branch circuit is determined using these formulas:

- **Single phase:** Load Capacity = Rated Voltage x Rated Current
- **Three phase:** Load Capacity = $\sqrt{3}$ x Rated Voltage x Rated Current

2.3.3 Branch Circuits: Rated Voltage

The rated voltage of a branch circuit specifies both its magnitude (volts) and number of phase conductors. Single-phase wiring is straight forward and consists of two wires (plus safety ground) where the AC voltage is a single sinusoidal wave as measured across the two wires.

Three-phase wiring is more complicated and consists of either three (three-phase conductors) or four (three-phase and one neutral) wires, plus safety ground. Three-phase branch circuits deliver more power, but require a rack PDU specially designed for three-phase branch circuits. Internally, a three-phase rack PDU divides the 3 or 4 branch circuit wires into pairs of single phase circuits – and these single-phase circuits are wired to the rack PDU's single-phase outlet receptacles.

The three-phase conductors have the same voltage magnitude but the sinusoidal AC waveforms are out of phase with each other by 120 degrees. Regardless of the number of wires, the rated voltage of three-phase wiring is always the measured voltage difference between any two-phase conductor wires – not the difference between a phase wire and neutral. Just as with single-phase power described above, connecting across one 120V hot line and the neutral provides 120V AC. Connecting across any two 120V hot lines, say L1 and L2, provides 208V AC, not 240V AC. Why? Because the phase of L1 is offset 120 degrees from L2 the voltage is not 240V (120V x 2), as it is for single-phase, but is 120V x $\sqrt{3}$ or 120V x 1.732 = 208V. A three-phase PDU can deliver three circuits of 208V each. Some rack PDUs take advantage of a neutral wire to provide three circuits of both 120V and 208V. But as mentioned in the preceding paragraph, regardless of the number of wires, or whether or not both a higher and lower voltage are supplied as outputs, a three-phase rack PDU is rated at the voltage between two phases, e.g., L1 and L2 which in the example here is 208V.

A rack PDU can also provide 400V AC. As with the 208V three-phase rack PDU, if one of those lines is connected to a neutral instead of another line, it provides a single-phase output circuit-- which for a 400V-rated PDU is 230V AC (400V / 1.732 = 230V). This is a common deployment in Europe and is becoming more common for high-power racks in North America.

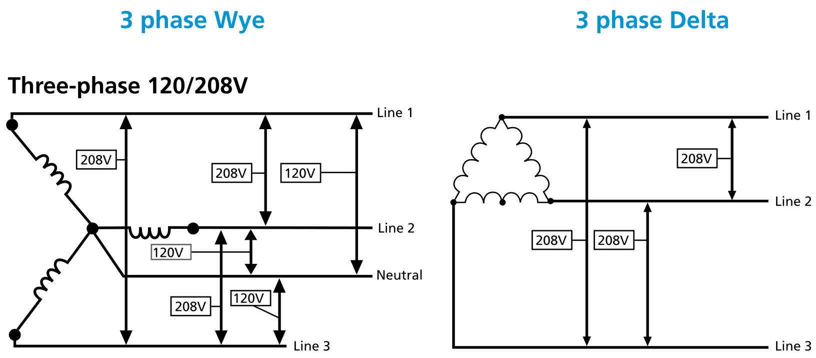

Figure 2: Three-Phase Wiring Diagram
(Source: Raritan, Inc.)

When looking at three-phase rack PDU specifications you will often see the terms Wye and Delta or the Greek letters Y and Δ. These terms or letters were chosen because the electrical configuration diagram of a Delta transformer looks like a Δ and the electrical configuration diagram of a Wye transformer looks like a Y. A rack PDU which does not convert a higher input voltage, e.g., 208V or 400V, to a lower output voltage, e.g., 120V or 230V, but instead retains the higher voltage throughout the PDU uses a Delta transformer. A Delta transformer has three connection points, one at each corner of the triangle. Each of these points is a connection for one of the three lines. Connecting any point to any other point provides a line-to-line connection, e.g., L1 to L2 and provides 208V or 400V as described in the examples above.

A rack PDU which does convert a higher input voltage to a lower output voltage uses a Wye transformer. A Wye transformer has three connection points for the lines, one at the end of each "arm" of the Y and one at the "foot" of the Y. The center intersection point of the Y is a fourth connection point and is where the neutral wire is attached. Connecting any two of the three line connections together, e.g., L1 and L2, provides 208V or 400V. Connecting any one of the three line connections to the neutral, e.g., L1 and neutral, provides 120V or 230V as described in the examples.

Rated Voltage	Location	Number of Wires	Outlet Voltage(s)
120V	North America	3 (line + neutral + ground)	120V
208V	North America	3 (line + line + ground)	208V
230V	International	3 (line + neutral + ground)	230V
208V 3Ø Δ	North America	4 (3 lines + ground)	208V
208V 3Ø Y	North America	5 (3 lines + neutral + ground)	Mixed 120V & 208V
400V 3Ø Y	International & North America	5 (3 lines + neutral + ground)	230V

Figure 3: Branch Circuit Rated Voltage and Wire Requirements
(Source: Raritan, Inc.)

2.3.4 Branch Circuits: Rated Current

The amount of current that can flow in a circuit is determined by the size (thickness) of its wire and terminating receptacle. All branch circuits are required to be over current protected using a circuit breaker (or fuse). The rating of the circuit breaker is sized to the current carrying capacity of the branch circuit's wiring and receptacle. For example, 10 AWG (American Wire Gauge) wire and a NEMA L21-30R receptacle are both specified at 30A – so a branch circuit using these components must be protected by a 30A circuit breaker.

In North America, the national electric code for data centers (NEC Article 645) requires branch circuit wiring to be rated 125% greater than the total connected load. To ensure this requirement is met without having to run heavier gauge wires, all electrical devices (rack PDUs, computers, etc.) used in North American data centers must be certified to UL 60950-1. UL 60950-1 limits a device to draw no more than 80% of the rating of its input plug. For example, a rack PDU (or any other device) containing a 30A NEMA L21-30P plug must not draw more than 24A. This 80% limitation is commonly known as "derated" current.

The following table summarizes power available for various branch circuits.

Location	Rated Voltage	Nominal Current	Rated Current	Available Power/ Branch Circuit
North America	**120V**	**20A**	**16A**	**1.9kW**
	208V			**3.3kW**
	208V 3Ø			**5.8kW**
International	**230V**	**16A**	**16A**	**3.7kW**
	400V 3Ø			**11.1kW**
North America	**120V**	**30A**	**24A**	**2.9kW**
	208V			**5.0kW**
	208V 3Ø			**8.6kW**
International	**230V**	**32A**	**32A**	**7.4kW**
	400V 3Ø			**22.2kW**

Figure 4: Branch Circuit Available Power
(Source: Raritan, Inc.)

2.4 Plugs, Outlets, and Cords

Rack PDUs are available with several types of plugs and receptacles (or outlets), designed so that only the appropriate rack PDU plug will fit into the appropriate circuit outlet and only the appropriate device plug will fit into the appropriate rack PDU receptacle. This is done to protect equipment, e.g. so that a device that is designed for 120 volts only isn't accidentally plugged into a 208-volts circuit, and for safety reasons; e.g. so that a server that draws 30 amps doesn't overload a circuit designed to handle only a maximum of 15 amps.

The two major classifications of plugs and receptacles used in data centers are defined by NEMA (National Electrical Manufacturers Association) and IEC (International Electrotechnical Commission). NEMA plugs and receptacles are most common in North America and IEC plugs and receptacles are most common in Europe. However, many data centers in North America use IEC plugs and receptacles and there are many families of plugs and receptacles in use in data centers around the world.

A significant concern in data center power distribution is unintentional disruption of power by accidental disconnecting of cords. A variety of solutions exist that lock the plug into the receptacle and prevent the cord separating from the receptacle.

There are three methods of securing the plug in the receptacle.

1. Plug with tabs snaps into the receptacle locking them together. One example of this system is Raritan's SecureLock™.
2. Plug inserted into a receptacle with a locking mechanism that grips the plug ground blade.
3. Wire retention clips mounted to the PDU chassis hold the plug in the receptacle.

The higher the current carrying capability of a plug, receptacle, or cord, the greater the amount of wire conducting material, typically copper, is required to prevent overheating the wire that could lead to an electrical fire. Note that the smaller the wire gauge number, the greater the diameter of the conductor.

The conductors are surrounded by insulating material and jacket, which may have special properties. For example, the jacketing may be designed to resist damage from exposure to oil. Typical insulating and jacket materials are PVC, rubber, and neoprene.

The number of wires in a cable can vary. Below are some typical data center configurations:

- Two wires: One hot and one neutral wire without a ground wire.
- Three wires: One hot, one neutral, and one ground wire.
- Four wires: Three hot wires (L1, L2, L3) and one ground wire.
- Five wires: Three hot wires (L1, L2, L3), one neutral wire, and one ground wire.

RECEPTACLE	PLUG	RATING
IEC 60320, C-13	IEC 60320, C-14	15 Ampere 250 Volt UL/CSA 10 Ampere 250 Volt International
IEC 60320, C-19	IEC 60320, C-20	20 Ampere 250 Volt UL/CSA 16 Ampere 250 Volt International
IEC 60320, C-15 IEC 60320, C-15	IEC 60320, C-14 IEC 60320, C-16	15 Ampere 250 Volt UL/CSA 10 Ampere 250 Volt International
IEC 60309 4H-R	IEC 60309 4H-P	20 Ampere 125 Volt UL/CSA
IEC 60309 6H-R	IEC 60309 6H-P	20 Ampere 250 Volt UL/CSA 16 Ampere 230 Volt European "CE" Mark, VDE

RECEPTACLE	PLUG	RATING
IEC 60320, C-5	IEC 60320, C-6	2.5 Ampere 250 Volt UL/CSA 2.5 Ampere 250 Volt International
IEC 60320, C-7	IEC 60320, C-8	2.5 Ampere 250 Volt UL/CSA 2.5 Ampere 250 Volt International
IEC 60320, C-13 IEC 60320, C-17	IEC 60320, C-18 IEC 60320, C-18	15 Ampere 250 Volt UL/CSA 10 Ampere 250 Volt International
IEC 60309 4H-R	IEC 60309 4H-P	30 Ampere 125 Volt UL/CSA
IEC 60309 6H-R	IEC 60309 6H-P	30 Ampere 250 Volt UL/CSA 32 Ampere 230 Volt European "CE" Mark, VDE

Figure 5: IEC Plugs and Receptacles

(Source: Raritan, Inc.)

RECEPTACLE	PLUG	RATING
NEMA 5-15R	NEMA 5-15P	15 Ampere 125 Volt USA & Canada Receptacle and plug. Polarized (UL 498)
NEMA 6-15R	NEMA 6-15P	15 Ampere 250 Volt USA & Canada Receptacle and plug. Polarized. (UL 498)
NEMA 5-20R	NEMA 5-20P	20 Ampere 125 Volt USA Receptacle and plug. Canada - Plug only. Polarized. (UL 498)
NEMA 6-20R	NEMA 6-20P	20 Ampere 250 Volt USA Receptacle and plug. Canada - Plug only. Polarized. (UL 498)
NEMA L5-15R	NEMA L5-15P	15 Ampere 125 Volt USA & Canada Locking receptacle and plug. Polarized. (UL 498)

RECEPTACLE	PLUG	RATING
NEMA L6-15R	NEMA L6-15P	15 Ampere 250 Volt USA & Canada Locking receptacle and plug. Polarized (UL 498)
NEMA L5-20R	NEMA L5-20P	20 Ampere 125 Volt USA & Canada Locking receptacle and plug. Polarized. (UL 498)
NEMA L6-20R	NEMA L6-20P	20 Ampere 250 Volt USA & Canada Locking receptacle and plug. Polarized. (UL 498)
NEMA L5-30R	NEMA L5-30P	30 Ampere 125 Volt USA & Canada Locking receptacle and plug. Polarized. (UL 498)
NEMA L6-30R	NEMA L6-30P	30 Ampere 250 Volt USA & Canada Locking receptacle and plug. Polarized. (UL 498)

Figure 6: NEMA Plugs and Receptacles

(Source: Raritan, Inc.)

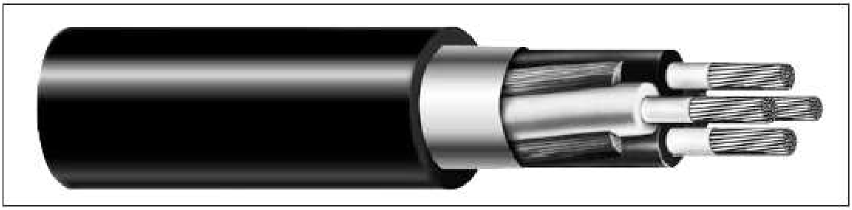

TYPE SOOW - 600 VOLT - UL/CSA										
CATALOG NUMBER	**NO. OF COND.**	**AWG SIZE**	**COND. STRAND**	**NOM. INS. THICKNESS**		**NOMINAL O.D.**		**CURRENT AMPS**	**APPROX. NET WEIGHT LBS/M**	**STD. CTN.**
				INCHES	**mm**	**INCHES**	**mm**			
02763	2	18	16/30	0.030	0.76	0.345	8.76	10	70	250'
02769	3	18	16/30	0.030	0.76	0.365	9.27	10	80	250'
02770	4	18	16/30	0.030	0.76	0.390	9.91	7	95	250'
02722	2	16	26/30	0.030	0.76	0.370	9.40	13	80	250'
02765	3	16	26/30	0.030	0.76	0.390	9.91	13	95	250'
02766	4	16	26/30	0.030	0.76	0.420	10.67	10	115	250'
02723	2	14	41/30	0.045	1.14	0.510	12.96	18	135	250'
02762	3	14	41/30	0.045	1.14	0.535	13.59	18	170	250'
02768	4	14	41/30	0.045	1.14	0.575	14.61	15	205	250'
02724	2	12	65/30	0.045	1.14	0.570	14.48	25	195	250'
02725	3	12	65/30	0.045	1.14	0.595	15.11	25	225	250'
02726	4	12	65/30	0.045	1.14	0.650	16.51	20	270	250'
02767	2	10	104/30	0.045	1.14	0.620	15.75	30	250	250'
02728	3	10	104/30	0.045	1.14	0.660	16.76	30	290	250'
02727	4	10	104/30	0.045	1.14	0.715	18.16	25	355	250'
16063	3	8	133/29	0.060	1.52	0.855	21.72	40	485	250'
16064	4	8	133/29	0.060	1.52	0.980	24.89	35	670	250'
16065	5	8	133/29	0.060	1.52	1.075	27.30	28	790	250'
16073	3	6	133/27	0.060	1.52	0.980	24.89	55	700	250'
16074	4	6	133/27	0.060	1.52	1.080	27.43	45	875	250'
16075	5	6	133/27	0.060	1.52	1.200	30.48	36	1015	250'
16083	3	4	133/25	0.060	1.52	1.140	28.96	70	920	250'
16084	4	4	133/25	0.060	1.52	1.260	32.00	60	1150	250'
16085	5	4	133/25	0.060	1.52	1.365	34.67	48	1400	250'
16093	3	2	133/23	0.060	1.52	1.330	33.78	95	1355	250'
16094	4	2	133/23	0.060	1.52	1.460	37.08	80	1690	250'
16095	5	2	133/23	0.060	1.52	1.580	40.13	64	1960	250'

Figure 7: Cord Specifications

(Source: General Cable)

2.5 Ratings and Safety

Rack PDUs, like all other electrical equipment, are subject to many general and specific safety standards. There are general industry terms and conventions that should be understood in order to ensure a reliable and safe data center. These are discussed in detail below.

2.5.1 Nameplate Data

Nameplate data is the electrical power consumption information specified by the equipment manufacturer. It is typically a conservative estimate of the maximum amount of power the device could draw. This information is found on a label near the electrical power input to the

IBM Server Model 520 - Rack Mounted Drawer

Configuration	Condition								
	Typical Heat Release (Voltage 110V)	Airflow				Weight		Overall System Dimensions (W x D x H)	
		Nominal		Maximum at 35°C					
	watts	cfm	(m³/hr)	cfm	(m³/hr)	lbs	kg	in	mm
Minimum	420	25	44	40	68	117	53	25 x 37 x 23	630 x 939 x 584
Full	600	30	51	45	76	117	53	25 x 37 x 23	630 x 939 x 584
Typical	450	25	44	40	68	117	53	25 x 37 x 23	630 x 939 x 584

ASHRAE Class	Airflow Diagram Rackmount Cooling Scheme F-R	Configuration	
3	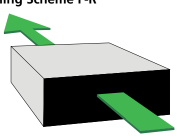 Front to Rear F-R	Description	Model
		Minimum	1-way, 1.5 GHz processor 15 GB memory
		Full	2-way, 1.65 GHz processor maximum memory
		Typical	1-way, 1.65 GHz processor 15 GB memory

device. There will be more discussion of the use of nameplate data below.

Figure 8: Nameplate Data

(Source: IBM)

2.5.2 Power Rating vs. Load Capacity

There can be confusion about power capacities and load capacity. The confusion stems from a misunderstanding of approval agency regulations and from some manufacturers who may use misleading terminology. In North America typical circuits have a maximum current carrying capability, and use circuit breakers or fuses rated at 15A, 20A, 30A, etc. In other words, a 20A fuse will blow or a 20A circuit breaker will trip if a 20A circuit experiences more than 20A for some period of time. The period depends on the magnitude of the current and the type of fuse or circuit breaker protecting the circuit.

In North America, circuits are to be loaded to 80% of their maximum capacity. For example, a 15A circuit should not carry more than 12A, a 20A circuit not more than 16A, a 30A circuit not more than 24A, etc. The 80% value, e.g., 16A for a 20A circuit, is often referred to as the derated value or the load capacity. In North America a rack PDU vendor's specifications sheet may have a few current carrying specifications.

The specifications provided and the terminology used may vary by vendor but the following is a typical example:

- **Maximum line current per phase:** 30A
- **Rated current:** 24A (30A derated to 80%)
- **Maximum current draw:** 6 x 16A (six circuits, each capable of carrying up to 16A)

In Europe and other parts of the world circuits are simply described at their rated capacity, e.g., 16A and 32A.

As mentioned above, apparent power is specified in volt-amps (VA) which is volts x amps. Load capacity is specified in VA where the amps are the rated current, i.e., the derated value. For example, for a single-phase rack PDU with a nominal voltage of 208V and a rated current (not the maximum current) of 24A, the load capacity is 5.0 kVA (208V x 24A).

2.5.3 Approval Agencies

In order to meet applicable local and national electric codes, rack PDUs must be safe and not emit electromagnetic radiation. Recognized approval agencies are contracted by manufacturers to test products according to existing standards. A product that passes agency testing receives an approval listing number, and the manufacturer can then affix the agency approval listing logo on each product. The listing logo is your assurance that the product meets applicable safety and electric codes. The manufacturer is required, upon request, to provide you the listing number and a copy of the testing report. You can also submit the listing number to the approval agency to verify compliance.

2.5.4 Proper Grounding

The National Electric Code (NEC Article 645.15) requires all exposed non-current-carrying metal parts of an information technology system to be grounded. This means all equipment within a rack and the metal rack itself must be grounded.

The inlet plug of a PDU contains a ground pin. When this plug is connected to a properly wired receptacle, the PDU becomes the grounding point for the equipment plugged into the rack PDU. The PDU can also be used to ground the metal rack and most PDUs contain a special threaded hole for this purpose. Typically, a grounding wire is connected to the rack and the PDU using screws. Care should be taken to make sure any paint on the rack is scraped off where the grounding wire is attached to ensure proper electrical conduction. There are also special grounding screws with teeth under the head of the screw to ensure a good ground connection.

Approval	Description	Standard/Revision/Year	Comment
UL	Safety	UL 60950-1	Required in USA
cUL/CSA	Safety	CAN/CSA-C22.2 No. 60950-1-03	Required for Canada
CB	Safety	IEC 60950-1	Common replacement for UL, CSA & CE in countries that accept CB
CE	Electromagnetic Interference (EMC)	EN 5502:2006	Europe
CE	Safety	EN 60950-1	Europe
FCC-A or B	Electromagnetic Interference (EMC)	FCC 47 CFR Part 15	USA
ICES-003	Electromagnetic Interference (EMC)	ICES-0003 issue-004	Canada

Figure 9: Safety and Electromagnetic Approval Agencies

(Source: Raritan, Inc.)

2.6 Overload Protection

The Underwriters Laboratories standard UL 60950-1 applies to the safety of information technology equipment (ITE) and requires the use of branch circuit over-current protection for ITE PDUs greater than 20 amps. Typically, ITE PDUs greater than 20 amps and certified after April 2003 must have built-in UL 489 circuit breakers or fuses (e.g., UL 248-5 fuses) suitable for branch circuit protection.

UL 60950-1 permits products at a maximum current of 15 and 20 amps without circuit breakers or fuses, since the 15 or 20 amp circuit breakers in the building are considered sufficient to protect the PDU; however, supplementary protection in the PDU provides additional protection. UL also “grandfathers” PDUs at more than 20 amps that were certified prior to April 2003. Although such PDUs are still being sold, their use should be avoided if they are to be incorporated in larger ITE systems designed to the latest UL 60950-1 standard.

Newly certified ITE PDUs at more than 20 amps are required to use over-current protection that meets branch circuit protection requirements in accordance with the National Electrical Code, ANSI/NFPA 70. In effect, this means these products are required to have branch circuit breakers listed under UL 489, “Standard for Molded-Case Circuit Breakers, Molded-Case Switches and Circuit Breaker Enclosures” or fuses listed for branch circuit protection, such as those listed to UL 248-5, “Low-Voltage Fuses – Part 5: Class G Fuses.”

In addition to standard UL 489, Underwriters Laboratories also publishes the standard UL 1077, “Standard for Supplementary Protectors for Use in Electrical Equipment.” Devices certified to this standard are called “Supplementary Protectors” and are called “Recognized” components, not “Listed” devices, as are UL 489 breakers. UL Listed Circuit Breakers meet more stringent requirements for branch circuit protection than Supplementary Protectors with UL Recognition.

Circuit breakers are used in a variety of ways. They are mounted in panel boards (also referred to as building PDUs) and rack PDUs to protect branch circuit wiring. They are also built into equipment to protect components and systems. Interrupting a short circuit – current flow limited only by the resistance of wiring – is a severe test of a circuit breaker. If the interrupting capacity of the breaker is not adequate, the device can literally explode.

UL 489 requires the breaker to be functional after being subjected to a short-circuit test. UL 1077 and the IEC standard EN 60934 allow for breakers to clear a short-circuit condition but become safely destroyed in the process. UL 489 breakers can interrupt short circuits of 5,000 amps or more. Typically, UL 1077 breakers can interrupt fault currents of 1,000 amps.

Overloads can be short term or long term. The protective device must not trip with a momentary or short-term over-current event that is normal for the piece of equipment being protected. Servers, for example, may create inrush currents as their internal power supply and filter circuits start. These inrush currents typically last only a fraction of a second and seldom cause a problem. If an overload lasts longer than a few minutes, the breaker should open to prevent overheating and damage. What gives a breaker the ability to discriminate between normal and damaging over-currents is its delay curve.

If you would like more information on the topic of data center overload protection see, “Data Center Power Overload Protection: Circuit Breakers and Branch Circuit Protection for Data Centers” at www.raritan.com/high-power-white-paper.

Chapter 2 Summary

Rack PDUs come in many configurations with respect to number and type of receptacles, voltage, load capacity and physical mounting.

- Rack PDU can be divided into two categories:

  Non-Intelligent PDUs

  Basic PDUs - Are power strips that are used in critical environments such as data centers. They distribute correct voltage and current to multiple outlets to power IT equipment in racks.

  Monitored PDUs - Allow a user to view a local display that typically provides electric current only. However, this information cannot be accessed remotely as the units have no network connectivity capabilities.

  Intelligent PDUs

  Metered Input PDUs - Meter power at the PDU-level, and can display the data both locally and over a network. Metering helps users to determine power usage and available capacity at the rack, and facilitates provisioning. By metering at the input-level, users can avoid overloading circuits and more easily calculate efficiency metrics like Power Usage Effectiveness (PUE).

  Metered Outlet PDUs - Meter power at the outlet-level, and can display the data both locally and over a network. Like metered input PDUs, outlet-metered models help users to determine power usage and available capacity at the rack, and facilitates provisioning. Most importantly, outlet-level metering allows users to understand power consumption at the device or server-level which make it possible to allocate costs to specific business units or customers.

  Switched PDUs - Offer the features of Metered Input PDUs and also provide controlled on/off switching of individual outlets. They enable authorized users to securely power cycle devices remotely in a specific order, offer power sequencing delay to minimize inrush currents, prevent unauthorized device provisioning, and can power off devices that are not in use to conserve energy.

  Switched PDUs with Outlet Metering - Combine all of the capabilities of Switched PDUs with those of Outlet-Metered PDUs.

- Power is distributed to the rack over one or more electrical branch circuits.
- Single-phase wiring is straight forward and consists of two wires where the AC voltage is a single sinusoidal wave as measured across the two wires.
- Three-phase wiring is more complicated and consists of either three or four wires, plus safety ground. Many PDUs support both single-phase and three-phase power.
- When looking at three-phase rack PDU specifications, a Delta transformer looks like a Δ and the electrical configuration diagram of a Wye transformer looks like a Y.

- The amount of current that can flow in a circuit is determined by the size (thickness) of its wire and terminating receptacle.
- Rack PDUs are designed so that only the appropriate rack PDU plug will fit into the appropriate feeder circuit outlet and only the appropriate device plug will fit into the appropriate rack PDU receptacle.
- In North America circuits are to be loaded to 80% of their maximum capacity. For example, a 15A circuit should not carry more than 12A. This limited value is known as the "derated" current.
- The National Electric Code (NEC) requires all exposed non-current-carrying metal parts of an information technology system to be grounded.
- Underwriters Laboratories requires the use of branch circuit over-current protection for PDUs greater than 20 amps.
- Overloads can be short term or long term. The protective device must not trip with a momentary or short-term over-current event that is normal for the piece of equipment being protected.

3.0 **Elements of the System**

3.0 **Elements of the System**

Rack PDUs are the final endpoint of power supplied to IT equipment from incoming building feeds through a chain of equipment including UPS, transformers, and larger PDUs and circuit panels. IT and facilities management are increasingly viewing their rack PDUs not merely as a collection of power outlets for IT equipment but as a network of critical devices that significantly impact the overall efficiency and effectiveness of the data center. As such they need to be properly managed like the IT equipment they power. This is driving the trend for use of more intelligent PDUs in data centers with environmental sensors and even integration with higher level data center management systems. This section describes not only the components of the physical rack PDU and basic environmental sensors, but also the rack PDU management system that leverages the intelligence in PDUs for operational improvements and energy use reduction. Further, this system can interface with and become part of a larger ecosystem of enterprise IT and facilities management systems.

3.1 Rack PDU

Over the past few years, average power consumption per server has rapidly increased with the adoption of high-power computing equipment like blade servers and data center containers. In addition, ongoing deployment of densely packed storage, virtualization, and cloud computing results in data centers with greater watts per sq. ft. requirements from more densely packed racks such as a rack filled with 1U servers. To support new, power-hungry IT equipment, data center managers have to deliver more power to the IT equipment rack. Over the last decade, the typical power required at a rack has increased from 2 kilowatts to 12 kilowatts and continues upward.

3.1.1 Single-Phase or Three-Phase Input Power for Rack PDUs

To accommodate the increased power demands at IT equipment racks, data center managers are deploying rack PDUs capable of supplying multiple circuits, higher voltages and higher currents. One way to increase the power at the rack is to increase the number of circuits and the voltage coming to the rack.

The amount of power available for use is referred to as apparent power and is calculated as volts x amps and is described as volt-amps or VA. A 120V, 20A circuit has an apparent power of 2400VA or 2.4kVA. A 208V, 20A circuit has an apparent power of 4160VA or 4.2kVA. Thus one 208V circuit provides almost twice as much power as one 120V circuit assuming the current (amperage) remains the same. With three 208V circuits it is clear that a substantial amount of power can be deployed in one three-phase PDU.

The cable to provide power to a three-phase PDU is thick and heavy but not as thick and heavy as the multiple, individual cables required to provide the same amount of apparent power using either single-phase 120V or single-phase 208V. Running a single three-phase power cable to each three-phase rack PDU reduces both the number of cables, making installations easier, and the physical bulk of the cables, so less space is filled with cables that can block necessary cooling airflow under raised floors and within IT equipment racks.

In cases where power needs to be provided at 120V for devices such as routers, hubs and switches, as well as at 208V for demanding servers, three-phase PDUs can provide outlets with both 120V (one of the three lines and a neutral) and 208V (two of the three lines). Three-phase power at the IT equipment rack is a convenient way for data center managers to efficiently deploy both greater power capacity and flexibility.

3.1.2 Form Factor

Rack PDUs are available in heights of one rack unit (1U; 1.75 inches) or two rack units (2U; 3.5 inches) for horizontal mounting in a 19-inch equipment rack.

Zero U rack PDUs mount vertically, typically to the vertical rails at the back of the rack. This can offer advantages. Zero U PDUs don't consume any rack unit spaces, and since the receptacles on the Zero U PDU line up better with the power cords for each IT device in the rack, they allow for the use of shorter power cords. This results in neater cable arrangements contributing to better airflow within the rack, which can improve cooling efficiency.

Depending on the rack cabinets, Zero U rack PDUs can be mounted with screws or hung into the cabinet via buttons, that are spaced 12.25 inches apart.

Higher-power rack PDUs will commonly be equipped with circuit breakers for branch circuit protection. These circuit breakers may cause the rack PDUs to extend deeper into the racks. Considerations should be made for how these PDUs are mounted in the rack and whether outlets are facing center or back to allow for cable management, airflow, and easy accessibility and serviceability of the IT equipment.

3.1.3 Outlet Density and Types

Rack PDUs vary in the number of outlets supported based on the physical size (length, width and depth) and the total space available for mounting outlets and internal components and the power handling capacity of the PDU. For example, a 1U rack-mount PDU may have enough space for eight 120V/15A NEMA 5-15R outlets; whereas a 2U rack-mount PDU may have enough space for twenty 120V/15A NEMA 5-15R outlets. On the other hand, a Zero U PDU may have 24 IEC C-13 230V/10A outlets or just four 250V/30A NEMA L15-30R outlets to support blade servers.

In the case of a large number of devices, each demanding a moderate amount of power, a large number of moderate power outlets are required. A typical dense "pizza box" deployment would include two rack PDUs for redundant power where each PDU is loaded to 40% so that if one power feed fails, the other feed will not exceed the NEC requirement of 80% (for North America). Typical outlets for "pizza box" servers are IEC C-13 (up to 250V, 16A) and NEMA 5-20R (up to 125V, 20A, 16A rated).

In the case of high power consumption at a rack for a few devices, each of which consumes a lot of power (blade servers, storage, or network devices), the total amount of power required might be comparable to the high outlet density example above, but the number and type of outlets may be different. Density for devices such as blade servers depends on their number of power supplies (often between two and six for redundancy), how the power supplies are configured (power supplies are most efficient when they operate close to their maximum level), and how many devices will be deployed in a rack.

In the case of a few devices demanding a lot of power, a large number of outlets may not be needed but outlets capable of delivering substantial power may be required. Typical outlets for high-demand devices, such as blade servers at 208V or 230V, are IEC C-13 (up to 250V, 16A) or C-19 (up to 250V, 32A) or, less commonly, NEMA L6-20R (up to 250V, 20A, 16A rated) or L6-30R (up to 250V, 30A, 24A rated) locking outlets.

3.1.4 Connectors: Ethernet, Serial, Sensor, USB and Other

Today, only the very basic rack PDUs have no external connectors. Most now include a variety of connectors based on application requirements. Below we describe four rack PDU connector configurations and general applications.

- **No connectors:** Cannot be managed externally and may not feature a local display.
- **Local navigation:** Provides users with a navigable display via local buttons. User can see basic data at the PDU or outlet-level.
- **Serial RS232 Connector:** Used for local metering; may be an LCD or LED monitor. Can be plugged into a terminal or console server for Telnet or SSH remote access. Access via a menu or command line interface using terminal emulation. Local buttons allow navigation to see basic unit data. No SNMP support available for alarms, unless via a specially developed serial console server. Typically non-switched.
- **Ethernet (RJ-45) and RS232 Serial (DB-9M) Connectors:** For remote metering for PDUs, circuit breakers, and outlets. USB-A (host) and USB-B (device) connectors to support PDU-to-PDU cascading, webcams and wireless networking. SNMP support available for alarms, Telnet or SSH access possible for command line access. Support for environmental sensors like temperature, humidity, airflow, air pressure and others may be available on the PDU or with an add-on external device. Remote metered models typically have an LCD or LED display with buttons for navigation to see basic unit and outlet data.

3.1.5 Branch Circuit Protection

Since April 2003, Underwriters Laboratories (UL) has required branch circuit protection, either circuit breakers or fuses, for PDUs where the inlet current is greater than the outlet current, e.g., 30A (24A rated) plug, 20A (16A rated) outlets. 15A and 20A (12A and 16A rated) rack PDUs can be supplied without branch circuit breakers because circuit breakers in upstream panel boards are deemed to provide the necessary protection. Rack PDUs with breakers or fuses are like mini-subpanels. For example, a 208V 30A (24A rated) 3-phase PDU has 3 circuits and each circuit / set of outlets has a 20A circuit breaker.

There are four types of circuit breakers--thermal, magnetic, thermal-magnetic, and hydraulic-magnetic.

- **Thermal circuit breakers** incorporate a heat-responsive bimetal strip or disk. This technology has a slower characteristic curve that discriminates between safe, temporary surges and prolonged overloads.
- **Magnetic circuit breakers** operate via a solenoid and trip nearly instantly as soon as the threshold current has been reached. This type of delay curve is not ideal for servers since they typically have inrush currents anywhere from 30 percent to 200 percent above their normal current draw.
- **Thermal-magnetic circuit breakers** combine the benefits of thermal and magnetic circuit breakers. These devices have a delay to avoid nuisance tripping caused by normal inrush current, and a solenoid actuator for fast response at higher currents. Both thermal and thermal-magnetic circuit breakers are sensitive to ambient temperature. A magnetic circuit breaker can be combined with a hydraulic delay to make it tolerant of current surges.

- **Hydraulic-magnetic breakers** have a two-step response curve. They provide a delay on normal over currents, but trip quickly on short circuits and are not affected by ambient temperature.

Circuit breakers used in rack PDUs are typically thermal-magnetic or hydraulic-magnetic with delay curves that allow for reasonable inrush currents while protecting devices from excessive fault currents.

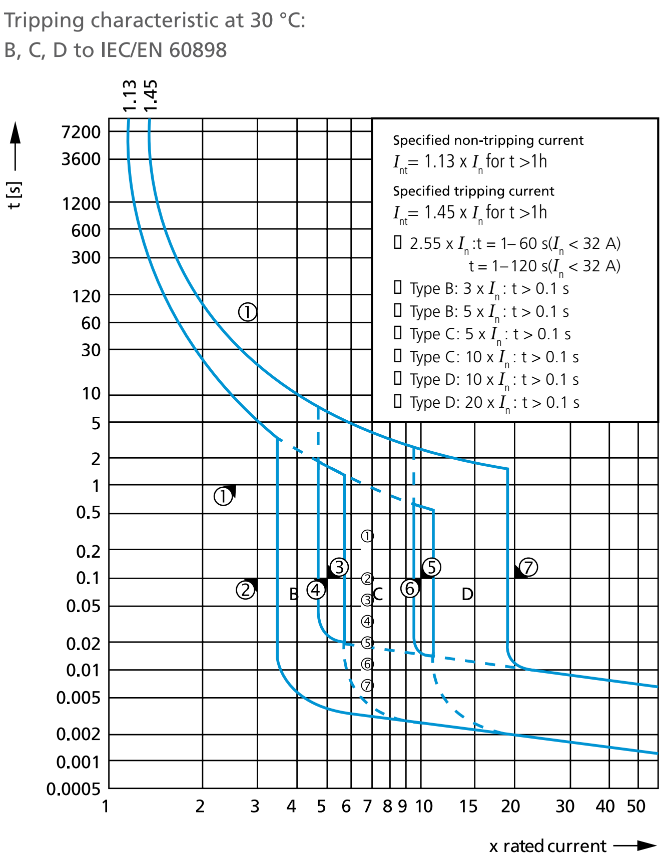

Figure 10: Trip Delay Curve

(Source: Moeller Electric)

Fuses are also acceptable for PDU circuit protection. However, replacing a fuse can be time consuming and may require an electrician leading to longer mean time to repair (MTTR). Spare fuses must be stocked in inventory and the correct fuse must be used to ensure reliability and protection.

The following are some points to consider when selecting a rack PDU:

- Compliance with the latest fuse and circuit breaker standards
- The acceptable MTTR for fuse replacement vs. circuit breaker resetting
- Impact on uptime service level agreements if a fuse blows vs. a circuit breaker trip.

3.1.6 Circuit Breakers: Single Pole vs. Double and Triple Pole

An important consideration is the reliability and flexibility of the branch circuit breaker configuration. Typically, circuit breakers are available as single, double or triple-pole devices. Single-pole breakers are appropriate for circuits comprised of a hot wire and neutral, e.g., 120V at 20A or 230V at 16A. Single-pole breakers provide a disconnect for the single hot wire used in circuits with a hot wire and neutral. Double-pole breakers provide a disconnect for circuits comprised of two hot wires, e.g., 208V at 20A. Some PDU designs use double-pole (or triple-pole) breakers to provide protection for 2 different circuits, e.g. two different hot wires. Since one double-pole breaker is less expensive than two single-pole breakers, this will lower the cost of the design. Double-pole breakers will trip if either of the two circuits they protect are overloaded. It is less expensive than 2 (or 3) single-pole breakers but, unless the poles can be operated independently, in a maintenance shutdown or trip, all 2 or 3 circuits are de-energized.

For example, assume a rack PDU with six branch circuits is protected by circuit breakers. Some rack PDUs in this configuration may protect the six circuits with three double-pole circuit breakers – one double-pole circuit breaker for the circuits with Line 1, one for the circuits with Line 2 and one for the circuits with Line 3. It is less expensive to use double-pole circuit breakers but there are some drawbacks. As noted, double-pole breakers will trip if either of the two circuits they protect are overloaded. This means double-pole breakers are less reliable. Double-pole breakers are also limiting because if for instance you choose to shut off a circuit for maintenance, you have no choice but to shut off both circuits. Alternatively, some rack PDUs protect the six circuits with six single-pole circuit breakers – one breaker per circuit. This is more expensive but single-pole breakers are more reliable and less limiting. Thus, consider rack PDUs that allow only one circuit to be de-energized for improved reliability and flexibility.

3.1.7 Circuit Breaker Metering

Circuit breaker metering is a useful feature on any rack PDU. It is particularly important when dealing with high power because the consequences of tripping a breaker can be disastrous if it means losing several blade servers. With circuit breaker metering the end user sets a threshold. When that threshold is crossed, an alert is delivered so the end user knows power demand needs to be reduced or there is the risk of tripping a circuit breaker. Monitoring branch circuit breakers is important since high power draw means a greater chance of tripping a breaker.

Line metering, intended for three-phase rack PDUs, is very useful for balancing the power drawn over each line. Overdrawing power from one line relative to another line wastes available power and unbalanced lines can place excessive demands on the neutral in Wye-configured PDUs.

3.1.8 Cord Length and Feed

Rack PDU power cords vary in length depending on the whips (power cables from a building PDU) and the location of the racks. The rack PDU power cord must be long enough to reach its power source, which is typically a whip located under the raised floor or an outlet just above the rack. A common power cord length is ten feet (3m), but other lengths can often be specified to a UL maximum of 15 feet (4.5m).

Rack PDU power cords may exit the PDU itself from the rear, the front, the top, or the bottom. With the power cable exiting the bottom of a Zero-U PDU, the data center manager will need to ensure sufficient space for the bend radius of the cable. In general, a bend radius of 5.25 inches (3U) will be sufficient, but this should be confirmed as bend radii will depend on the gauge (AWG) of the cord. A smaller bend radius may be acceptable for thin cables and a larger bend radius may be required for heavy-duty cables. The orientation of the PDU power cord may seem trivial but it can be a potential problem depending on the physical rack and the location of the power source for the rack. Consider the orientation of the power cord and how it will be routed to connect to the whip. For example, does the power come up from the raised floor or down from cable trays above the racks and is there room inside the rack to route the cable so as not to block airflow?

3.1.9 Cord Retention

Proper PDU cord retention practices, just like rack cable management, can make a big difference in operational efficiency and reliability. Taking steps to support, organize and secure the many power cords using some method of cord retention will dramatically improve your ability to access and manage the equipment connected to PDUs inside the rack. This will also minimize the chance of inadvertently unplugging power cords from rack PDUs. You should neatly arrange and secure the power cords between equipment and the rack PDU to allow for maximum airflow.

There are a few different methods to secure IT equipment power cords. Two of the most common are to design the rack PDU with locking IEC C-13 and C-19 outlets that work on standard IEC power cords. Another option, which is less expensive, is to use specially designed outlets and power cords with tabs that securely hold the cord in place until the tabs are released. (See Figure 11) These cords can be purchased in different colors to help identify which power supply is plugged into the PDU.

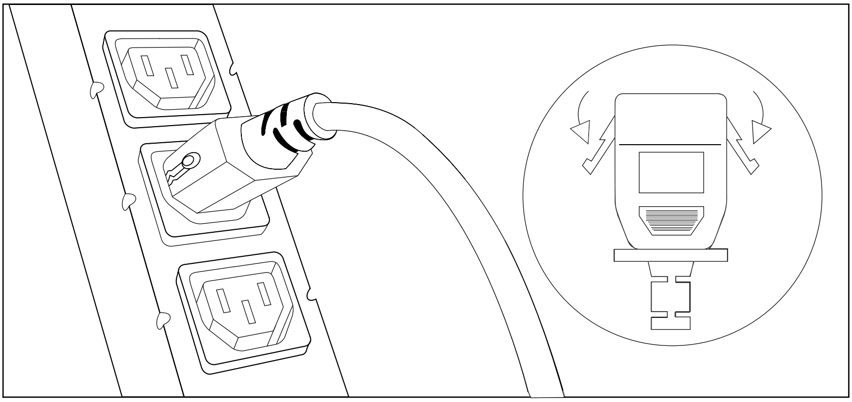

Figure 11: SecureLock™

(Source: Raritan, Inc.)

3.1.10 Local Display and User Interface

Virtually all rack PDUs designed for data center use have built in displays, typically LEDs, to show current draw for the entire PDU unit. Local displays have limited functionality compared to the information and control available from a remote interface, but they can be convenient and useful when working at the rack itself. The local display might allow an IT admin to toggle between current draw and voltage; or, for those rack PDUs that monitor individual outlets, to sequence through the outlets to determine the current being drawn by each IT device. Some switched and intelligent models will have LED indicators next to each outlet to display status, whether it is on/off, booting, firmware upgrade, or fault.

In addition to a local display on the rack PDU itself, some PDUs offer a serial interface for local terminal connectivity via a laptop for configuration, diagnostics or connectivity to a serial console server that concentrates multiple connections.

3.1.11 Remote User Interface

For a remotely accessible rack PDU (all but the basic PDU or PDU with metering and only local display) there are typically two choices for a remote user interface to the rack PDU over an IP network. The most common is a web-based graphical user interface to an Ethernet-enabled PDU. Some PDUs support SSL encrypted access (using https), while others support only unencrypted access (using http). Check your organization's security requirements when selecting a PDU.

The PDU can also be accessed via Ethernet over IP using SSH (encrypted) or Telnet (unencrypted) with a Command Line Interface (CLI). Security considerations should be kept in mind before enabling/disabling Telnet access. Some PDU manufacturers provide a serial console server that connects to the PDU locally via serial (RS232) and allows access to the unit remotely using SNMP or CLI.

One factor to consider is integration with central directory services for user authentication and access control. This becomes especially important when the rack PDU offers the ability to remotely turn on/off or recycle individual outlets or groups of outlets. Remote access to the PDU does not eliminate the need for some local access to the PDU with an LED/LCD and associated buttons.

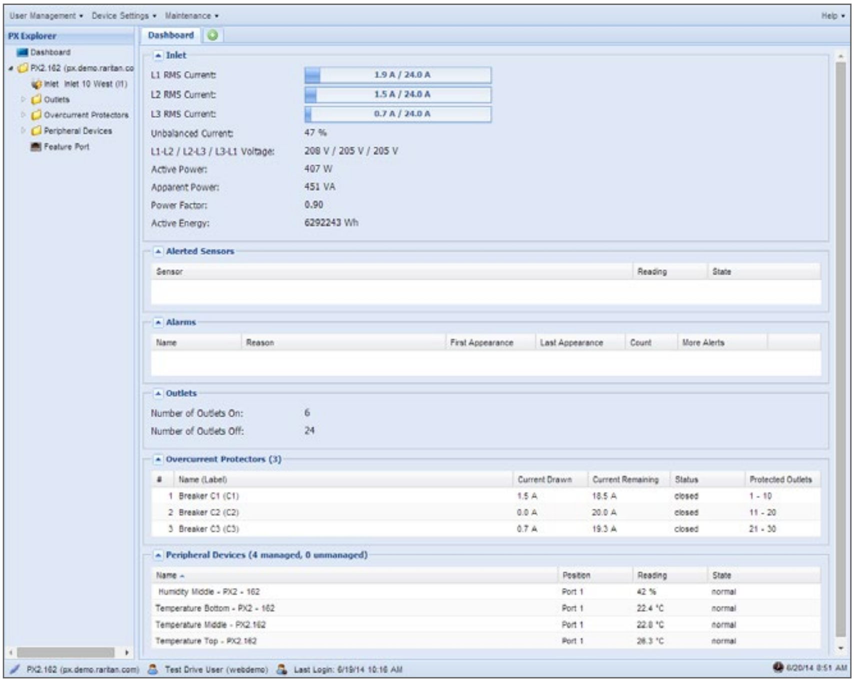

Figure 12: Remote User Interface

(Source: Raritan, Inc.)

3.1.12 Inline Meters

For data centers with existing PDUs that lack any metering capability, adding in-line meters can be useful to monitor power consumption per line or per circuit. In-line meters are usually one line or circuit in and one line or circuit out. Some vendors offer models that support up to four in's and four out's. There are basic inline meters that provide simple current metering, with or without IP connectivity. Others are more sophisticated and provide richer data like active power, apparent power, and kWh metering; and some might even include integrated environmental monitoring. By upgrading older, basic PDUs with in-line meters with IP connectivity, they can be managed along with metered and intelligent rack PDUs by a rack PDU management system so that you have a comprehensive view of the health and usage of power for day-to-day operations and planning.

3.2 Environmental Management

With the IT industry's increasing focus on improving data center efficiency, more rack PDU manufacturers are offering environmental sensors. These include sensors to measure rack air temperature at the server inlets, humidity, airflow, vibration, smoke, water, and air pressure. Some PDUs will have pre-installed sensors; others provide for optional, plug-in external sensors. Another common approach is to deploy a completely independent rack management system, choosing from a wide range of environmental sensors; however, this has the disadvantage of consuming additional rack space for the rack management system as well as the cost of a separate infrastructure—e.g. IP addresses, Ethernet ports, and cabling. Connectivity for sensors is typically either via RS485 or 1-Wire®.

3.2.1 Temperature Sensors

Temperature sensors monitor the air inlet temperatures at IT devices such as servers (See the sensor placement diagram on the next page). Since IT equipment generates considerable heat, manufacturers specify a range of acceptable temperatures for proper operation. A sensor-capable PDU should allow thresholds to be set for sending automatic alerts when the inlet temperature approaches the vendor-specified maximum to prevent servers from shutting down or failing due to overheating. In addition, it is also a good practice to set a minimum temperature threshold to provide alerts when the inlet temperature is colder than necessary. From a data center plant perspective, the cost of cooling and moving air is the largest infrastructure expense, so maintaining IT inlet air temperatures colder than necessary merely wastes energy and money. Temperature sensors at the rack also provide early warning about temperature extremes, hotspots or cold spots, and can help identify when an HVAC system is becoming unbalanced. To ensure that IT equipment is getting enough cool air, the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) has recommended that temperature probes be placed at specific locations at the inlets of equipment racks.

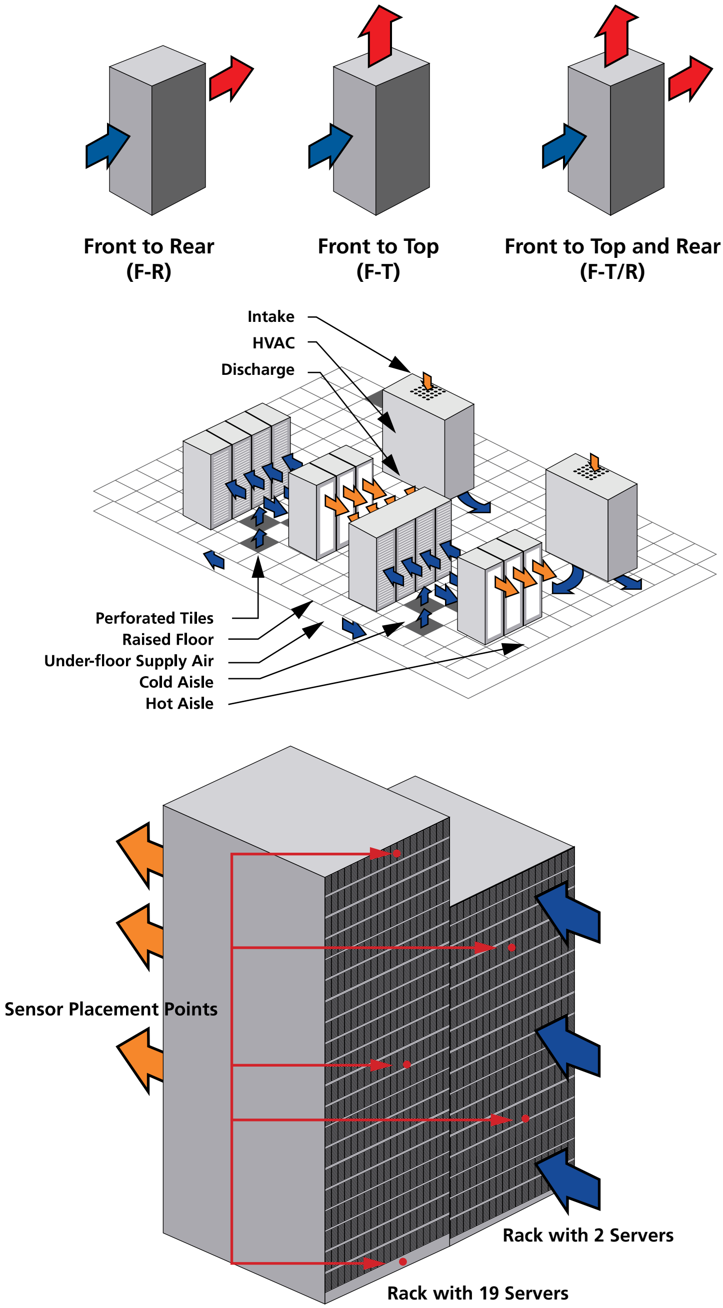

Figure 13: Recommended Temperature Sensor Placement

(Source: ASHRAE)

3.2.2 Humidity Sensors

Understanding the basics of what humidity is and how it affects your server room can impact how long your computer equipment lasts and how much your electricity bill costs. Humidity is a measurement of moisture in the air. High humidity can cause condensation buildup on computer components, increasing the risk of an electric short. Likewise, if the humidity is too low, data centers can experience electrostatic discharge (ESD). Humidity can be monitored per area or zone to ensure that it is in the safe range. ASHRAE has recommended ranges for the data center that should be consulted. Appropriate thresholds and alarms should be set to indicate a potential problem.

Classes (a)	Equipment Environmental Specifications							
	Product Operations (b)(c)					Product Power Off (c) (d)		
	Dry-Bulb Temperature (°C) (e) (g)	Humidity Range, non-Condensing (h) (i)	Maximum Dew Point (°C)	Maximum Elevation (m)	Maximum Rate of Change (°C/hr) (f)	Dry-Bulb Temperature (°C)	Relative Humidity (%)	Maximum Dew Point (°C)
Recommended (Applies to all A classes; individual data centers can choose to expand this range based upon their analysis described in this document)								
A1 to A4	18 to 27	5.5°C DP to 60% RH and 15°C DP						
Allowable								
A1	15 to 32	20% to 80% RH	17	3050	5/20	5 to 45	8 to 80	27
A2	10 to 35	20% to 80% RH	21	3050	5/20	5 to 45	8 to 80	27
A3	5 to 40	-12°C DP & 8% RH to 85% RH	24	3050	5/20	5 to 45	8 to 85	27
A4	5 to 45	-12°C DP & 8% RH to 90% RH	24	3050	5/20	5 to 45	8 to 90	27
B	5 to 35	8% RH to 80% RH	28	3050	NA	5 to 45	8 to 80	29
C	5 to 40	8% RH to 80% RH	28	3050	NA	5 to 45	8 to 80	29

Figure 14: Recommended and Allowable Humidity and Temperature Ranges

(Source: ASHRAE)

3.2.3 Airflow Sensors

Airflow sensors will detect a reduction of air movement that might create the potential for overheating, which can destroy IT equipment. There are two primary areas for monitoring airflow in the data center–above the floor (monitored at a number of points), and below the floor (monitored at select points). Differential airflow sensors are used to ensure that the pressure differential between the subfloor and the floor is sufficient to control air flowing from the subfloor to the floor above. Blockages in under-floor supply plenums can cause high pressure drops and uneven flow, resulting in cold spots in areas where cooling air is short circuiting to the return path. Airflow sensors should have thresholds set, and alarms enabled, like other environmental sensors, to ensure that data center managers are alerted when conditions are less than optimal for efficient cooling.

3.2.4 Air Pressure Sensors

It is important to have the appropriate air pressure in under-floor supply plenums, but sometimes this is treated as an afterthought. Air pressure that is too high will result in both higher fan costs and greater leakage that can short circuit cooling air, while pressure that is too low can result in hot spots at the areas most distant from the cool-air supply point. This can lead to poor efficiency 'fixes' to correct the problem, such as a lowering the supply air temperature or overcooling the full space just to address a few hot spots. Differential air pressure sensors can be used to ensure that the pressure differential between the subfloor and the floor is sufficient. Maintaining appropriate room pressure prevents airborne particulates from entering the data center.

3.2.5 Contact Closure Sensors

Contact closure sensors can be used for a variety of applications. For example, a contact closure could send an alert when a cabinet door is opened and trigger a webcam to take a picture. Contact closure sensors can be connected to any device that can open or close a contact--making them very versatile.

3.2.6 Other Sensors

There are a variety of other sensors that can be used in the data center. Examples include in-cabinet smoke, water, and vibration sensors. Like the other sensors mentioned above, these are used to send alarms when measured conditions are outside the range for proper data center operation.

3.3 System Connectivity

3.3.1 Physical Topology

Like many functions of data center management, best practices for the physical topology and, by extension, remote management of rack PDUs, is evolving. The current best practice is to connect all remotely accessible rack PDUs to the "management network" (separate from the "production network") directly in order to collect periodic meter readings, get immediate notifications of any faults or potential problems and enable remote power cycling of IT equipment (depending on the intelligence of the rack PDU). When planning for a new facility it is good practice to provide for a minimum of two Ethernet drops for each cabinet or rack, since each will typically require two PDUs.

However, some data centers try to reduce the number of Ethernet drops and IP addresses to minimize cabling and costs. In these cases the data center deploys PDUs that can be daisy chained or cascaded so that one Ethernet connection can communicate with all the PDUs in the chain.

3.3.2 Communication Protocols

The communications protocols used are typically TCP/IP when PDUs are Ethernet connected, and proprietary protocols for PDUs serially connected to a console server, which, in turn, connects to the TCP/IP network via Ethernet. Most often, SNMP protocol is used for management, while LDAP and Active Directory are used for authentication, authorization, and access control. SSH and Telnet may be used for command line management and HTTP/HTTPS for web-based access, in some cases with SSL encryption protocols.

There are rack PDUs now with USB-A (host) and USB-B (device) ports that can be used to support devices such as webcams and Wi-Fi modems. Some rack PDUs support MODBUS, a common, older building management communication protocol, and some rack PDUs support the GSM modem protocol so that cell phones can receive one-way text alerts.

3.3.3 Managing the Rack PDU

As mentioned earlier, the management system for data center power is often run on a "management network" separate from the production network. This reduces the likelihood of a Denial of Service (DOS) or other attack that would affect this critical function. In mission critical facilities, there are often two connections to each rack PDU equipped with remote communications: one for syslogging, SNMP traps, access via Web browser, and kWh logging, and another for critical functions like remote power cycling, status of circuit breakers, and load monitoring (normally at the unit or circuit level). In some cases, administrative functions, like rack PDU configuration, are performed via command line scripting through a secondary interface such as a serial port, while Ethernet remains the primary interface for all other functions.

Some important management functions are listed below:

1. **Audit logging:** To track activity like switching of outlets, configuration changes, etc. Two or more syslog servers are often used for this function.
2. **Fault management:** Done via SNMP with tools like HP Openview, IBM Tivoli, and others. SNMP V2 is still the most commonly used, but SNMP V3, with its built-in security, is recommended for applications requiring outlet control.
3. **Configuration:** Setup via Web browser, SNMP, command line, or a central software tool.
4. **Firmware upgrades:** Not an issue for older PDUs with minimal functionality, but something that may be required for Ethernet-enabled PDUs. A central tool is essential to manage large numbers of PDUs and to simplify management and reduce cost of ownership.
5. **Alerts:** Normally via SMTP messages.

A combination of some or all of the above capabilities is required to effectively manage a data center. Check your application requirements and choose the PDU type appropriate for your application. If you have multiple rack PDUs (40+) you will want to consider a comprehensive rack management system, as discussed in the following section.

3.4 Rack PDU Management System

A rack PDU management system is a software application (sometimes delivered as a software appliance) that consolidates all communications with your rack PDUs and in-line meters equipped for remote communication. Its main functions are data collection, reporting, power control, element management and fault management. The system collects and converts raw power data into useful information and provides a central point for secure access and control across multiple rack PDUs with operation validation and an audit log. Consequently, it simplifies the management of rack PDUs and alerts you to potential incidents. We include it in this section because for larger data centers with more than 40 racks it is a "must have" in order to realize several of the benefits offered by metered, switched or intelligent PDUs--improved energy efficiency, increased uptime and lower operational costs.

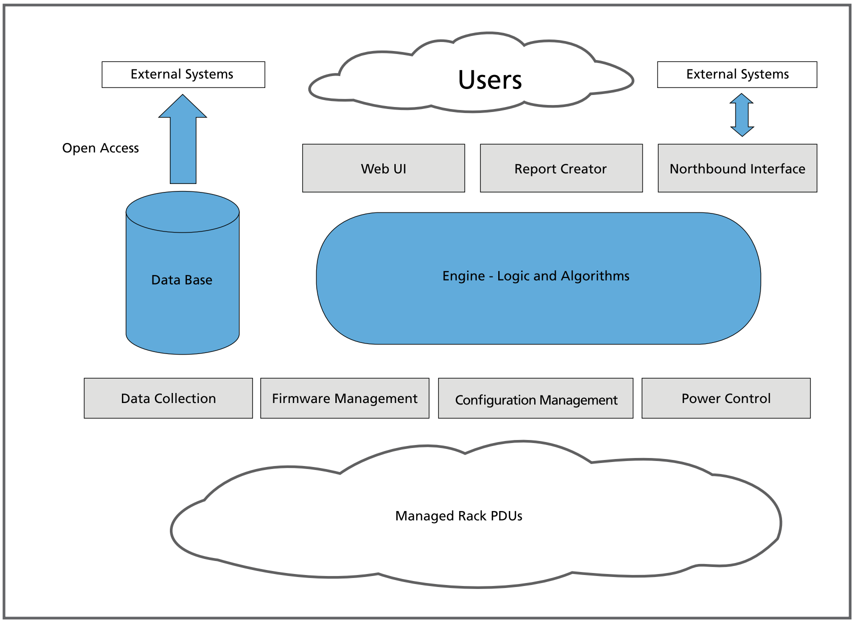

Figure 15: Rack PDU Management Topology

(Source: Raritan, Inc.)

3.4.1 Data Collection

Data collection is the fundamental component that enables all reporting and most other management functions. Of course, the management system can collect only the data elements provided by the managed PDUs. As discussed earlier, basic PDUs provide no data, metered PDUs provide total unit data and intelligent PDUs may provide individual outlet data and more, so it is important to understand what data you will want to analyze when selecting rack PDUs.

Typical data elements you will probably want to collect, as available from the PDUs, include:

- Total unit active and apparent power.
- Line current and capacity.
- Outlet-level current and active power.
- Environmental sensor data.
- Real-time kWh metering data.

Next you will want to determine the granularity of the data you need. Your management system should offer a user-configurable data polling interval. For most applications, a normal polling interval is 5 minutes, which means the system will collect data points every five minutes, but if you require greater granularity, your rack PDU may need to be able to store data readings so that the network is not overloaded with polling traffic. You will want to use a roll-up algorithm to collect data for long periods without causing the database to balloon and effect performance. For most energy management applications, actual data readings are rolled up to a maximum, average and minimum, on an hourly, daily and monthly basis.

Advanced polling options enable a customer to minimize network traffic while still enabling granular data collection. Advanced polling requires a rack PDU that has the memory capacity to record and store readings called samples. The rack PDU management system should offer the ability to configure optional sample rates for each rack PDU and also set optional polling intervals for the management system itself to collect the stored samples at each rack PDU not previously collected. For example a rack PDU can be configured to record and store samples every minute. The rack PDU management system can be configured to poll the rack PDU once an hour. In each poll, it will pull the 60 one-minute samples since the last poll, with the intelligence to know the last reading it recorded on the previous poll.

3.4.2 Reporting and Analytics for Power Monitoring and Measurement

Reporting and graphing should include active power, current, temperature, humidity, and information derived from the basic collected data--such as energy usage, cost and carbon emitted due to the energy consumed--for standard and selected time periods.

Reports on maximums and minimums for current, temperature, humidity and active power simplify key tasks and ensure that you are not in danger of exceeding circuit breaker ratings, overcooling or undercooling. For instance, environmental information can give data center operators the confidence to raise temperature set points without the risk of undercooling IT equipment. Analyzing trend line graphs, reports, and "what-if" modeling can help you perform capacity planning based on real world data.

Outlet-level data and reporting granularity can help you become more energy efficient. It enables you to determine the potential savings of upgrading to more energy efficient servers or the benefits of server virtualization. Consolidating several low-utilization physical servers as virtual servers on one high-utilization physical server can reduce overall expenses, but you will need to understand the resulting power demand of the host servers. You can also establish objectives, report on usage and implement changes for both physical components of the data center--floor, room, row, rack and IT device; and, also logical groupings like customer, department, application, organization, and device type. This level of detail creates visibility and accountability for energy usage. Some IT organizations issue energy bill-back reports to users/owners of the IT equipment.

3.4.3 Graphical User Interface (GUI)

The graphical user interface is your window into all of the rack PDU management system functions. This should be clean, intuitive, and web-based, functioning with all major web browsers. A web-based system provides you more remote access options and is easier to support and upgrade. The GUI will most likely include a user-configurable dashboard. The dashboard can be displayed in the data center network operations center for an easy, at-a-glance view of the status of the data center power and environmental conditions. This will give your customers, internal and external, a good indication of your data center management capabilities.

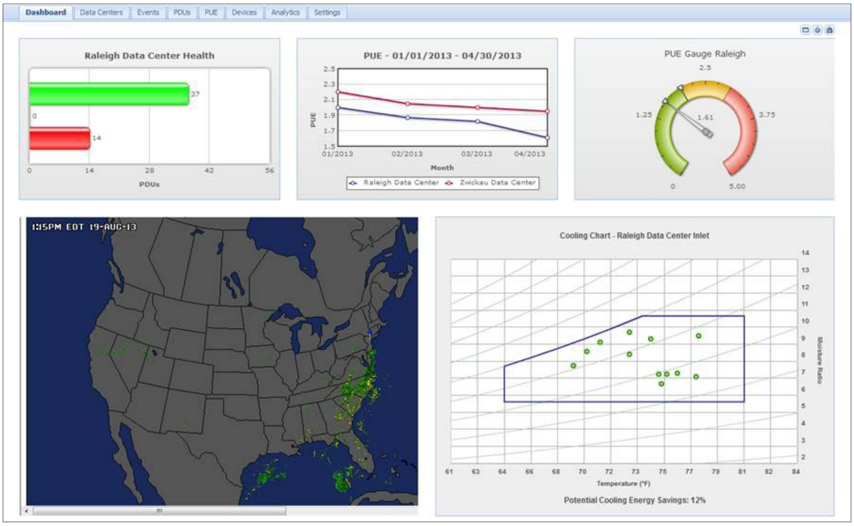

Figure 16: Rack PDU Management System Graphical Interface
(Source: Raritan, Inc.)

3.4.4 Element Management

The main components of element management include centralized rack PDU access and control, firmware management, and bulk configuration. You can view all your managed rack PDUs from one web browser window and get a summary view of the name, location, status, manufacturer model, and firmware level. You will want to be able to drill down to manage at the PDU unit level, line level and, in many cases, outlet level. One click, sign-on access to each managed PDU can give you control through the PDU's own GUI.

Since intelligent PDUs run firmware with many configuration options, the rack PDU management system should allow you to centrally store rack PDU firmware/configuration versions and facilitate distribution to multiple PDUs. Configuration template storage and distribution will simplify initial PDU installation as well as future unit additions and replacements.

3.4.5 Fault Management

Rack PDU management systems often provide both a map view and a floor layout view, and use a color scheme to provide an at-a-glance view of the health of all managed PDUs. Health problems are discovered in several ways. The system can receive an SNMP trap or a syslog event so that you are alerted of the problem as it happens. A management system can also poll the rack PDU at set intervals to collect the heath status of the communication path to the PDU or critical failures, and forward events to a higher-level enterprise management system.

3.4.6 Local and Remote Control/Switching

Switched rack PDUs allow for outlet control including on/off power cycling. However, most IT devices have more than one power supply for power redundancy purposes and these supplies are connected to outlets on separate Rack PDUs. Through the management system, you can power cycle at an IT device level, which will programmatically switch outlets from multiple PDUs. The system should also allow for grouping of IT devices into racks so you can control a full rack. Finally, any switched PDU must allow for flexible sequencing and delay so that an in-rush current spike does not trip a circuit breaker and so that application intra-system dependencies are taken into account during startup and shutdown.

3.4.7 Security of Data and User Access

Remote monitoring, metering and management require secure remote access via Ethernet and/or serial connections. To ensure security an intelligent rack PDU should have strong encryption and passwords and advanced authorization options including permissions, LDAP/S and Active Directory. A web session timeout will protect against a user accidentally leaving an authenticated session live while not in use.

3.4.8 Administration and Maintenance

Most of the administration takes place during the initial set-up. All systems will allow for GUI entry of this data, but that can be time consuming. Systems should also allow for the import of configuration information; e.g. via CSV files. During the set-up you will add your rack PDUs and the hierarchical and logical relationships. Hierarchical relationships include data center, floors, rooms, rows, racks, rack PDUs and IT devices. Logical associations include owners/customers of the IT device and IT device type. The administrator will also set the data pruning intervals to ensure unnecessary data is pruned from the system.

3.4.9 Open Point of Integration

Most data centers have some other management systems already in use, so it is important that the rack PDU management system can be integrated into these systems to minimize the amount of duplicate data entry and collection. Asset management and enterprise reporting systems are two typical systems that should logically interface with the rack PDU management system. The asset management system will automatically add rack PDUs, IT devices and their associated connections to the rack PDU management system's inventory of managed devices. Integration with an enterprise reporting system enables the creation of custom reports with the additional ability to correlate data that exits in other systems.

In recent years, a more comprehensive class of products called data center infrastructure management (DCIM) have been introduced that normally include the above functions, overall capacity planning tools and more.

Learn more about DCIM at www.RaritanDCIM.com

Chapter 3 Summary

- Rack PDUs are the final endpoint of power supplied to IT equipment from incoming building feeds. As a result, IT data centers increasingly view rack PDUs as an integral part of managing their network of critical devices.
- Data centers are deploying rack PDUs capable of supplying multiple circuits, higher voltages and higher currents.
- Rack PDUs are available in a variety of rack form factors, that lead to neater cable arrangements and better airflow along with more efficient device powering.
- Basic PDUs are falling out of favor as intelligent PDUs can collect power and environment information and warn of potential problems.
- Circuit breakers used in rack PDUs are typically thermal-magnetic or hydraulic-magnetic with delay curves that allow for reasonable inrush currents while protecting devices from excessive fault currents.
- Typically, circuit breakers are available as single, double or triple pole devices. Single-pole breakers are optimal because they allow a single circuit to be de-energized.
- Circuit breaker metering is important, particularly when dealing with high power because tripping a breaker can be disastrous if it means losing several blade servers.
- Rack PDU power cords may exit the PDU itself from the rear, the front, the top, or the bottom. This will affect how the PDU fits into your rack and the bend radius of the cable.
- Virtually all rack PDUs designed for data center use have built-in displays to show current draw for the entire PDU unit.
- To remotely access a rack PDU there are typically two choices: a graphical user interface (GUI), or a Command Line Interface (CLI).
- Integration with central directory services is important when the PDU offers the ability to remotely turn on/off/recycle individual outlets or groups of outlets.
- PDU manufacturers are offering environmental sensors to measure: rack air temperature at the server inlets, humidity, airflow, vibration, smoke, water, and air pressure.
- The management system for data center power is often run on a "management network" separate from the production network. This reduces the likelihood of a Denial of Service (DOS) or other attack that would affect this critical function.
- A rack PDU management system is a software application that consolidates all communications with your rack PDUs. Its main functions are data collection, reporting, power control, element management and fault management.

4.0 Considerations for Planning and Selecting Rack PDUs

4.0 Considerations for Planning and Selecting Rack PDUs

There are several approaches to deploying power to IT equipment racks, which affect rack PDU selection and configuration. Some approaches provide degrees of redundancy and hence higher reliability/availability than others, but may not be appropriate for certain types of equipment. Redundancy and higher availability require resources, so managers of data centers that have limited power resources need to decide what IT equipment justifies redundant power, e.g. production servers, and what equipment does not, e.g. non-production equipment being tested or evaluated.

4.1 Power Available and Distributed to Racks

4.1.1 Single Feed to Single Rack PDU

The simplest power deployment to an IT equipment rack is a single appropriately sized power feed to a single rack PDU. IT equipment with one or more power supplies would plug into this single rack PDU. If that single feed or single rack PDU should fail, for whatever reason, power to the equipment in the rack will be lost. The failure could occur at the rack PDU itself or farther upstream (perhaps a main feed fails or a building PDU circuit breaker trips).

As noted earlier, the NEC requires that circuits be loaded to no more than 80% of their maximum capacity. For example, if a 30A feed and rack PDU were deployed in this configuration, the load allowed (the rated current) would be 24A (30A x 80%). NEC would expect the feed and PDU to handle a maximum of 30A but the circuit should be loaded to only 24A.

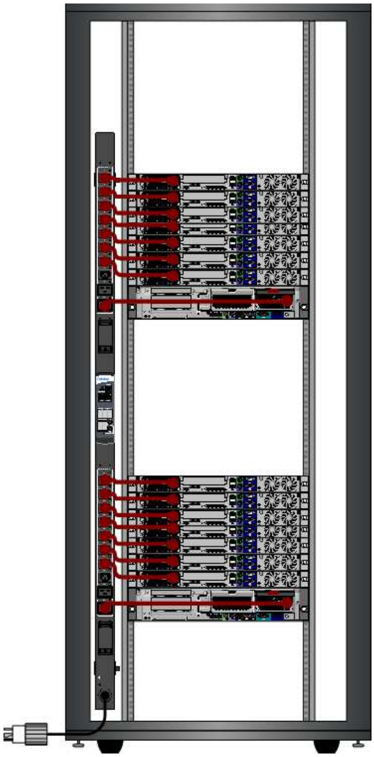

Figure 17: Single Feed to Single Rack PDU

(Source: Raritan, Inc.)

4.1.2 Dual Feed to Single Rack PDU with Transfer Switch

The next step up in availability is still a single feed to a single rack PDU, but with the addition of a transfer switch, which typically has two feeds from the same or from different building feeds. If one of the feeds to the transfer switch fails, the transfer switch automatically switches to the other power feed and the rack PDU continues to supply power to the IT equipment. However, if the single rack PDU fails, the power to the IT equipment is lost.

Learn more about transfer switches at www.raritan.com/TransferSwitch

There are two types of transfer switches: static transfer switch (STS) and automatic transfer switch (ATS). A STS is based on static electronic component technology (silicon controlled rectifier or SCR), which results in faster and better controlled transfer between sources. An ATS is less expensive and is based on electro-mechanical relay technology, which results in slower transfer times.

Again, with this arrangement, the rack PDU is still loaded to 80% of the maximum, but the electrical power capacity required has doubled – one feed is operational and the second feed is a backup. It has also doubled the amount of upstream equipment necessary to supply the additional feed.

Two power feeds to an ATS and then to a single rack PDU is generally used only where reliability is a concern and the IT equipment itself, e.g. a server, has only one power supply.

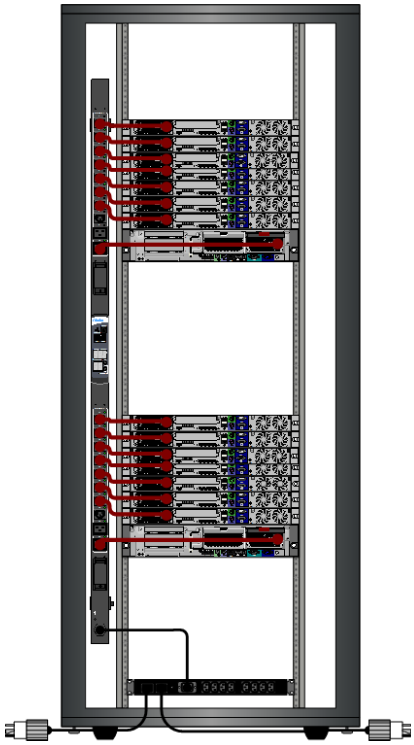

Figure 18: Dual Feed to Single Rack PDU with Transfer Switch

(Source: Raritan, Inc.)

4.1.3 Dual Feed to Dual Rack PDUs

Today, many servers, network devices, storage systems, even keyboard, video, mouse (KVM) switches and serial console servers, are available with dual power supplies. Some larger servers may have as many as four or even six power supplies. The most reliable deployment here is to use two power feeds to two rack PDUs. With this configuration, if one rack PDU or power feed fails there is a second one available to maintain power to the IT equipment in the rack. A common practice when using dual feeds is to use rack PDUs with colored chassis, such as red and blue. The colored chassis enables a visual control for installation of or changes to the PDU and connections. The rack will have a red chassis PDU fed by input circuit "A" and a blue chassis PDU fed by input circuit "B". The colored chassis helps to eliminate confusion about which PDU is fed by circuit "A" or "B".

It is important to remember the requirement that each circuit be loaded to no more than 40%. If the two circuits feeding the rack are both loaded to 80%, the NEC requirement will be met, but think about what would happen if one of the circuits failed. The power demand to the second circuit would jump from 80% to 160% and the circuit breaker for that feed would trip so the second circuit to the rack would also lose power. To prevent this both feeds should be loaded to no more than 40%, so that if one fails, the remaining circuit won't be loaded to more than 80%. Compared to the previous case with ATS (Section 4.1.2) where one feed is the backup, in this configuration, both feeds are powering IT equipment.

Note that if you intend to perform remote switching for IT equipment with dual power supplies, you will want to use a rack PDU that supports outlet grouping; i.e., two or more outlets are controlled as though they are a single outlet.

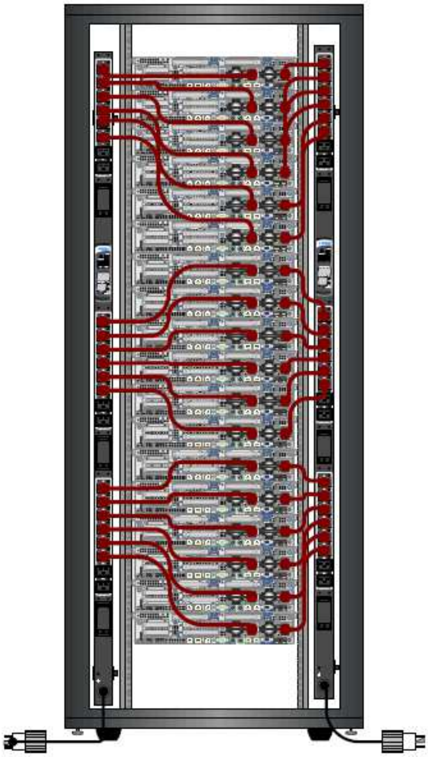

Figure 19: Dual Feed to Dual Rack PDUs

(Source: Raritan, Inc.)

4.1.4 Multiple Power Supplies

IT equipment with two or more power supplies can vary in the way power is delivered to the equipment. Some devices have a primary and backup power supply; some alternate between the power supplies; and some devices share power demand across all the power supplies. For example, a blade server with four power supplies in a 3+1 redundancy configuration would draw one third of its power from each of its three primary power supplies, leaving one for redundancy in the event any one of the three fail. Some more sophisticated devices have multiple power supplies, which are designed for both redundancy and efficiency. For example, some devices might drive utilization rates higher on specific power supplies to drive higher efficiency. You will need to check with each equipment manufacturer to understand how the power supplies work so that optimal balanced load configurations can be achieved on the rack PDU, especially those with branch circuits and 3 phase models.

4.1.5 Load Balancing

Load balancing is a procedure which attempts to evenly distribute the rack equipment's current draw among the PDU's branch circuits so that as you come closer to perfect balance, more total current can be supplied with the greatest headroom in each branch circuit. For example, consider a PDU with two 20A circuit breaker protected branch circuits – where each branch contains a number of outlets. The total current capacity of the PDU is 40A with the limitation that no branch circuit of outlets can exceed 20A. If the total load of all devices plugged into the PDU is 30A, perfect balance is achieved when the load is exactly divided between the two branches (15A each branch). The headroom in each branch is then 5A (20A circuit breaker less 15A load). Any other distribution of the load (16A:14A, 17A:13A) results in less headroom.

Load balancing has similar benefits for three phase PDUs. As the load comes closer to perfect balance, the current draw is more evenly distributed among the three phase lines (more upstream headroom) and total current flowing in the three lines is minimized. For example, consider a 24A three-phase delta wired PDU with three branch circuits. When an 18A load is balanced across the three branch circuits (6A load in each branch), the current flowing in each input phase line is 10.4A (6 x sqrt 3 or 6 x 1.732) and the total current in all three lines is 31.2A (10.4 x 3). If the entire load was carried by one branch circuit (totally unbalanced), the current in the three phase lines are 18A, 18A, and 0A, respectively and the total current is 36A. When the load is balanced across all three lines, the PDU has 7.6A (18.0A – 10.4A) more headroom.

Load balancing can be tricky because many IT devices draw power in varying amounts depending on the computational load. For devices with single power supplies an estimate of the power consumption should be made for each device and then the devices plugged into the circuits, so that the circuits are loaded evenly. This is true both within a rack and across multiple racks. For devices with dual power supplies they should be plugged into different circuits. A typical deployment would be the dual feed to dual rack PDUs mentioned above.

For IT devices with more than two power supplies, such as blade servers, load balancing can become even more complicated, especially if the rack PDUs are three-phase models. As an example, assume four blade chassis are to be installed in a rack, each chassis has six power supplies and two three-phase rack PDUs will be installed in the rack for redundant power. The first blade server will have power supplies (PS) # 1, # 2 and # 3 plugged into circuits (C) # 1, # 2 and # 3 respectively on PDU A and power supplies (PS) # 4, # 5 and # 6 plugged into circuits (C) # 1, # 2 and # 3 respectively on PDU B. Since we want to try to balance the load across all circuits and lines, and we can't be sure that each of the four blade servers will be performing tasks that equally load the circuits, we will stagger the second blade server power supplies. So the second server will have PS # 1 plugged into C # 2; PS # 2 plugged into C # 3; PS # 3 plugged into C # 1 on PDU A; PS # 4 plugged in C # 2; PS # 5 plugged into C # 3, and; PS # 6 plugged into C # 1 on PDU B. Circuit level metering, phase level metering, and outlet level metering will be very helpful for (re) balancing loads in the rack.

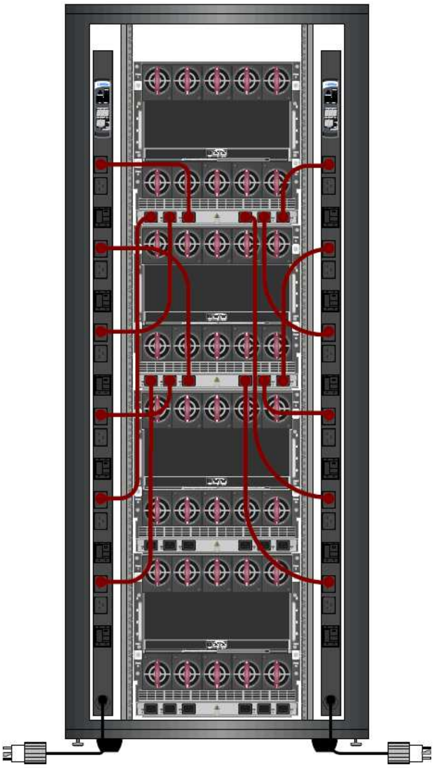

Figure 20: Multiple Power Supply Configuration

(Source: Raritan, Inc.)

4.1.6 Inrush Current

Servers draw substantially more current when they are first turned on. This is known as inrush current. As discussed in the section on overload protection, rack PDUs with circuit breakers are designed not to trip during very short periods of high currents. However, it's easier on upstream circuits if the sudden surge when equipment is turned on is minimized. For this reason, some rack PDUs provide outlet sequencing and allow users to configure both the sequence and also the delay time in which the outlets are turned on. Some rack PDUs even allow programming of outlet groups and allow sequencing of groups of outlets.

4.2 Power Requirements of Equipment at Rack

Section 4.1 deals with ways to deploy electrical power to a rack. This section deals with determining how much power to deploy to a rack. Typically, the starting point is an IT device's nameplate power requirement data (see section 2.5.1) which specifies a voltage and current (amps), which is typically higher than what is usually seen during actual deployment. As a result, there is a convention of using a percentage of this value, e.g. 70%, when computing the maximum PDU load capacity required: PDU load capacity = sum (device nameplate (V x A)) x 70%). For example, 208V x 2.4A x 70% x 14 servers = 4.9kVA.

For the example above, if you run 208V you need a 30A (5kVA) rack PDU since you will load it to 80% to meet North American requirements (4.9kVA / 208V = 23.5A, 23.5A is approximately 80% of 30A). If you want redundancy, add a second 5kVA rack PDU and load both PDUs up to 40%. You will need to specify the appropriate number of outlets. It is a good idea to have a few spare outlets for other devices, even if the rack PDU will be at its maximum capacity. More efficient or different equipment might be installed in the rack in the future or servers may not run near full capacity, leaving additional power capacity to power more equipment. The current best practice is to standardize on IEC C-13 and/or C-19 PDU outlets, and 208V. Most servers and data center devices can run at 208V (even up to 240V).

Remember the derating factor of 70% was just an estimate. Research has been done with sophisticated rack PDUs that accurately measure power at the outlet. The findings were surprising. Even at peak power consumption, 15% of the servers drew 20% or less of their nameplate rating. Equally surprising was that nearly 9% drew 81% or more of their nameplate rating. The point here is that the actual power consumed as a percentage of nameplate rating can vary widely. Ideally, data center managers should measure the actual power consumption rather than use a rule-of-thumb average such as 70%. If the actual overall average is closer to 40%, as it was in the study, deploying power at 70% of nameplate is wasteful and strands unused power.

If a cabinet populated with 30 1U servers has dual power feeds and the servers require an average of 150W each, then the total power requirements for a cabinet are 150W x 30 servers = 4.5kVA. Assuming 250VA for additional equipment, like an Ethernet switch and a KVM switch, this brings the total to 4.75kVA. So a 208V 30A PDU, which is rated at 5kVA, would be sufficient. Such a PDU can carry the full load of 4.75kVA in a failover situation when the power feed to one side of the cabinet fails, or is taken down for maintenance. Typically, each PDU would be carrying only 40% of the 4.75kVA.

It is also important to note that three-phase Wye 208V rack PDUs are able to support both 120V and 208V in the same PDU. This can be handy for situations where a variety of equipment types with different voltage requirements need to be racked together.

4.2.1 Rack PDU 208V Single-Phase vs. 208V Three-Phase

In a rack of 42 1U servers, if each server consumes an average of 200W then the total power consumption is 42 x 200W = 8.4kW. To allow for the NEC requirement of 80% the rack needs 10.5kVA (8.4kW / 0.8). To allow for redundant power feeds two rack PDUs able to provide 10.5kVA are required. 208V single phase at 60A (48A rated) can deliver 10.0kVA. This could suffice, particularly if the 200W per server estimate is on the high side. Another alternative is 208V three phase at 40A (32A rated) which can deliver 11.5kVA. The 208V three-phase alternative provides headroom to add higher power demand servers in the future and can handle the existing servers even if their average power consumption increases from 200W to 220W.

The use of three-phase power enables one whip or rack PDU to deliver three circuits instead of just one circuit. The whip or input power cord on the rack PDU will be somewhat larger for three-phase power than single phase power because instead of three wires (hot, neutral and ground) a three-phase cable will have four or five wires.

The two three-phase alternatives are Delta and Wye. A three-phase Delta system will have four wires – Line 1 (hot), Line 2 (hot), Line 3 (hot) and a safety ground. Individual circuits are formed by combining lines. Three circuits are available – L1+L2, L2+L3 and L1+L3. The power on each of the lines is a sine wave (this is also the case for single-phase power), but each of the three sine waves is 120 degrees out of phase with the other two.

For three-phase power, the sine waves are 120 degrees out of phase, so calculating VA is slightly more complex because we need to include the √3 which is 1.732. The apparent power formula for three phase is V x Derated A x 1.732 = VA. As an example, three phase 208V, 40A (32A derated) is 208V x 32A x 1.732 = 11.5kVA. A three-phase Delta deployment provides three separate circuits and over 70% more total power than a comparable single phase, single circuit.

A three-phase Wye system will have five wires – Line 1 (hot), Line 2 (hot), Line 3 (hot), a neutral and a ground. Individual circuits are formed by combining lines and by combining a line with the neutral. As an example, a three-phase 208V Wye rack PDU supports three 208V circuits (L1+L2, L2+L3, L1+L3) and three 120V circuits (L1+N, L2+N, L3+N). Three-phase Delta and three-phase Wye have the same apparent power, but the three-phase Wye provides two different voltages while the three phase Delta only provides one voltage.

In North America, there may be a requirement for 120V convenience outlets, such as NEMA 5-15R (120V, 15A, 12A rated) or 5-20R (120V, 20A, 16A rated). These can be supported by 208V three-phase Wye PDUs where wiring between lines (L1, L2, L3) and lines and the neutral can provide power to both 208V and 120V outlets. Whether the three-phase wiring is Delta or Wye, the voltage is always referenced to the line-to-line voltage, not the line-to-neutral voltage. This is even true in the 400V example on the next page (Section 4.2.2), where all the outlets are wired line-to-neutral.

Since the Wye system adds a neutral wire, many data centers are wired for Wye and use whips terminated with Wye receptacles, such as NEMA L21-30R. This means the data center can use Wye PDUs that support 120V/208V or use Delta PDUs that support only 208V without needing to change the data center wiring. A Delta PDU would use a NEMA L21-30P (the mating Wye plug) but would not use a neutral inside the PDU. This is a perfectly acceptable practice. For example, a data center could deploy Delta PDUs to racks where there is only a need for 208V and Wye PDUs to racks where there is a need for both 120V and 208V.

Three-phase cables may be slightly larger than single-phase cables but it is important to remember that one slightly thicker three-phase cable will be significantly smaller and weigh less than three single-phase cables for the same voltage and amperage.

4.2.2 Rack PDU 400V Three-Phase

As shown in the 208V/120V example (Section 4.2.1), three-phase Wye wiring is a convenient way to step down voltage. This is particularly true for 400V power. A generally accepted method of delivering substantial power to densely packed racks is via 400V three-phase Wye rack PDUs. 400V power distribution from panels to racks is now an accepted practice. A data center designer could specify 400V Wye whips to 400V Wye rack PDUs. Since most data center equipment can safely operate on voltages ranging from 100V to 240V, the 400V Wye PDU can provide three circuits – L1+N, L2+N, L3+N – each supplying 230V (400V / 1.732). 400V Wye rack PDUs do not lend themselves to supporting 120V outlets as 208V Wye rack PDUs do.

4.3 Rack PDU Selection

4.3.1 Rack PDU Selection and Special Application Requirements

There are many factors involved in selecting a rack PDU: data center location, application requirements, IT equipment requirements, available power, space in the cabinet, energy management, and efficiency objectives will combine to dictate what type of PDU should be used. Some of the considerations below will guide you to select the feature set and hence the type of PDU you will need to satisfy your requirements.

What type of equipment and how many devices are going into the cabinets, e.g. 42 x 1U servers with a single feed per device, versus three 10U high blade servers with six power supply feeds per server? The answer will help define the physical configuration, e.g. number and type of outlets, and capacity of your PDU(s), e.g. how much power (kW) the PDUs need to support. Average rack power requirements have risen from 6.0kW in 2006 to 7.4kW in 2009 and 12.0kW in 2011, and it is not unusual to see racks wired to provide as much as 30kVA.

Decision criteria for 24/7 manned sites will be different than remote management of lights-out facilities. If you need remote or lights-out management of a facility, then you will probably need a switched PDU, which will require more security and user access management. Remote applications may also call for SNMP management.

Integration with directory services, like LDAP or Microsoft's Active Directory, is increasingly a requirement for controlling access to resources, rather than requiring a separate access control system. This capability is applicable to all applications, requiring central authentication, local or remote. And for many data center applications, e.g. federal government and financial institutions, encryption and strong password support are necessary for remote access.

The rack PDU must supply uninterrupted power to each device plugged into it. You will want to prevent or mitigate any events that can potentially cause the circuit breaker on the rack PDU or any point upstream to trip. Outlet sequencing is a valuable feature to prevent inrush current from tripping a circuit breaker by establishing a sequence and appropriate delay for powering multiple devices. Outlet sequencing not only prevents the undesired tripping of a circuit breaker but also lets the user specify the order in which services (devices) come online or are shut down during power cycling. For example, you will want to power the database service before the web servers. This capability is most useful when used in conjunction with the outlet grouping capability.

For some applications and equipment, you many need a customizable alarm threshold for each outlet, with the capability to switch off an outlet should it exceed a certain power draw. This would prevent a temperature or other sensor (see Section 3.2) from causing a shutdown of servers. An advanced application is the control of the HVAC system using the temperature reported by PDU temperature sensors.

In many mission critical environments, managed devices often have multiple feeds which will be fed from different feeds or circuits for failover and redundancy. The managed device needs to be managed as a single device regardless of the number of power supplies/plugs, and all outlets must be handled simultaneously. This capability is applicable to all applications, local or remote.

Event-driven power cycling of an outlet/device is required for some applications, particularly for remote or unmanned sites. For example, if a device in a remote location fails to respond and the WAN is not operational, there are basically two options: first, an expensive, time-wasting truck roll to restart, or second, a rack PDU with the intelligence to trigger a restart of a malfunctioning device, e.g. if the device has not responded for 20 minutes, recycle power to the device.

If there is a need to maximize power efficiency, then rack PDUs can provide valuable data to support those efforts: current, voltage, and power factor measurements at the PDU, line, breaker, and outlet-level. Look for accurate kWh metering at the outlet level, especially if you intend to report or charge back individuals or groups for usage. Metering accuracy can vary significantly and, for some rack PDUs calculations may be based on assumptions and not actual real-time measurements.

4.3.2 Rack PDU Functionality

Rack PDUs can vary significantly, not only in the operational functions they offer, but also in their monitoring and data collection. Below is an overview of the strengths and weaknesses of the four types/classes of rack PDUs previously defined in Section 2.1.1. Clearly, our class definition is not rigid, since features offered by vendors will vary and you will want to select PDUs based on the total fit to your requirement, but this can be a useful guide in your selection.

Non-Intelligent PDUs

- **Basic PDUs -** are power strips that are used in critical environments such as data centers. They distribute correct voltage and current to multiple outlets to power IT equipment in racks.
- **Monitored PDUs -** allow a user to view a local display that typically provides electric current only. However, this information cannot be accessed remotely as the units have no network connectivity capabilities.

Intelligent PDUs

- **Metered Input PDUs -** Meter power at the PDU-level, and can display the data both locally and over a network. Metering helps users to determine power usage and available capacity at the rack, and facilitates provisioning. By metering at the input level, users can avoid overloading circuits and more easily calculate efficiency metrics like Power Usage Effectiveness (PUE).
- **Metered Outlet PDUs -** Meter power at the outlet-level, and can display the data both locally and over a network. Like metered input PDUs, outlet-metered models help users to determine power usage and available capacity at the rack, and facilitates provisioning. Most importantly, outlet-level metering allows users to understand power consumption at the device or server-level which make it possible to allocate costs to specific business units or customers.
- **Switched PDUs -** Offer the features of Metered Input PDUs and also provide controlled on/off switching of individual outlets. They enable authorized users to securely power cycle devices remotely in a specific order, offer power sequencing delay to minimize inrush currents, prevent unauthorized device provisioning, and can power off devices that are not in use to conserve energy.
- **Switched PDUs with Outlet Metering -** Combine all of the capabilities of Switched PDUs with those of Outlet-Metered PDUs.

4.3.3 Benefits of an Intelligent Rack PDU

Over the past decade, the IT industry has chosen to move to more sophisticated, manageable systems. This fact is nowhere more evident than the trend toward the use of intelligent PDUs.

An intelligent PDU will provide, at minimum, real-time outlet-level and/or PDU-level power monitoring, rack temperature and humidity monitoring, and may provide remote outlet switching. For top-tier data centers, deployment of intelligent PDUs can make a significant difference in the ability of IT administrators and facilities managers to improve uptime and staff productivity, efficiently utilize power resources, make informed capacity planning decisions, and save money. And, in so doing, they will operate greener data centers. Clearly, if your data center has dozens of racks, then the greatest benefits will be realized by using a rack PDU management system to consolidate data acquisition, reporting, PDU administration, and control. Let's examine the benefits of deploying intelligent PDUs in depth.

Improve Uptime and Staff Productivity

- Monitoring power at a PDU and individual outlet level, with user-defined thresholds and alerts via e-mail or SNMP, provides awareness of potential issues before they occur.
- Remote reboot of servers and IT equipment from anywhere in the world via a web browser reduces downtime and personnel costs.

Use Power Resources Safely

- User configurable outlet-level delays for power sequencing prevent circuits from tripping from IT equipment in-rush currents.

- Control of outlet provisioning prevents accidentally plugging IT equipment into circuits that are already heavily loaded and are at risk of tripping circuit breakers.

Make Informed Power Capacity Planning Decisions

- Outlet-level monitoring may identify some simple rearrangements of equipment to free up power resources by balancing power demands across racks.
- Monitoring power at the outlet-level can identify equipment that may need to be changed to stay within the margin of safety of defined thresholds.
- Monitoring rack temperature and other environmental conditions can prevent problems, especially when a data center is rearranged and airflow patterns change.

Save Power and Money

- Monitoring power at the outlet-level combined with trend analysis can identify ghost or underutilized servers that are candidates for virtualization or decommissioning.
- Remote power cycling enables IT managers to quickly reboot hung or crashed IT equipment without incurring the cost of site visits.
- Temperature and humidity sensors help data center managers optimize their air conditioning and humidity settings and avoid the common practice of overcooling and related waste of energy.

4.4 Power Efficiency

4.4.1 Introduction to Power Usage Effectiveness (PUE)

It is important to understand the role the rack PDU can play in providing accurate data that IT and facilities management will use to improve efficiency. We all know the maxim, "You can't manage what you don't measure," but be aware that unless the data is clearly understood, it can unwittingly lead to false conclusions and inappropriate actions.

The objective is to maximize the amount of useful computational work by IT equipment for total energy consumed by the equipment itself and the infrastructure that supports the equipment. The most commonly used metric is The Green Grid's Power Usage Effectiveness or PUE:

PUE = Total Facility Power ÷ IT Equipment Power

"Total Facility Power" in this equation is all the power required to operate the entire data center, including the IT equipment items: servers, storage, network equipment and other IT equipment; and the support infrastructure items: CRAC units, fans, condensers, UPS and lighting. "IT Equipment Power" is the power required to operate the servers and IT equipment alone. Theoretically PUE can range from 1.0 (where all the power is consumed by IT equipment only) to infinity. A PUE = 2.0 means that IT equipment is consuming 50% of data center power. Another commonly used metric is Data Center Infrastructure Efficiency (DCiE), which is the inverse of PUE. Since they are derived from the same data, there is no substantive difference in the measurement or its usage.

A recent Lawrence Berkeley National Laboratory study of data center power allocation showed an average data center PUE of 2.18 and a similar result was found by EYP Mission Critical Facilities, Inc. (now an HP company), with the following allocation:

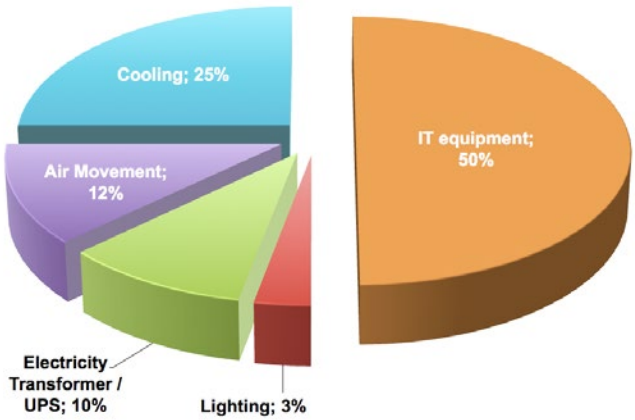

Figure 21: Data Center Power Consumption

(Source: EYP Mission Critical Facilities, Inc.)

4.4.2 PUE Levels

The Green Grid defines three levels of PUE: Basic or Level 1, Intermediate or Level 2, and Advanced or Level 3. Many industry analysts recommend measuring IT power consumption at the Intermediate, Level 2; i.e., at the PDU level. While it is true that PDU-level power consumption will provide the denominator needed to calculate PUE, this information alone is likely insufficient to drive the best efficiency improvement decisions. Regardless of the PUE Level you choose to employ, the best practice is to gather data over a time period of "typical' power usage to ensure that the peaks and valleys have been captured in order to establish a baseline and to track your improvements.

	Level 1 Basic	Level 2 Intermediate	Level 3 Advanced
IT equipment power measurement from	UPS	PDU	Server
Total facility power measurement from	Data center input power	Data center input less shared HVAC	Data center input less shared HVAC plus building, lighting, security
Minimum measurement interval	Monthly/weekly	Daily	Continuous

Figure 22: Definition of Green Grid PUE Levels

(Source: The Green Grid)

4.4.3 Why Advanced, Level 3 PUE?

An improved (lower) PUE may be misleading since it can result from inefficiencies in the power consumed by IT equipment which merely increases the denominator. A lower PUE is generally better than a higher one, but it is possible to implement measures that reduce data center energy consumption yet actually increase your PUE. For example, if you were to replace older, less efficient servers with more efficient ones, eliminate ghost servers, turn off servers that were idle during the night, or employ server virtualization, the net result would be power reduction, but your PUE would actually increase. The detailed IT load data from Level 3 provides the granularity of information to reduce energy consumption, not just improve the PUE metric. Clearly, PUE (and its inverse, DCiE) becomes a more useful beacon once you have built efficiency into the IT equipment performance; to do that you will want the granular power usage data for the Advanced, Level 3 PUE metric. Then you can attack the numerator and squeeze inefficiencies out of the infrastructure.

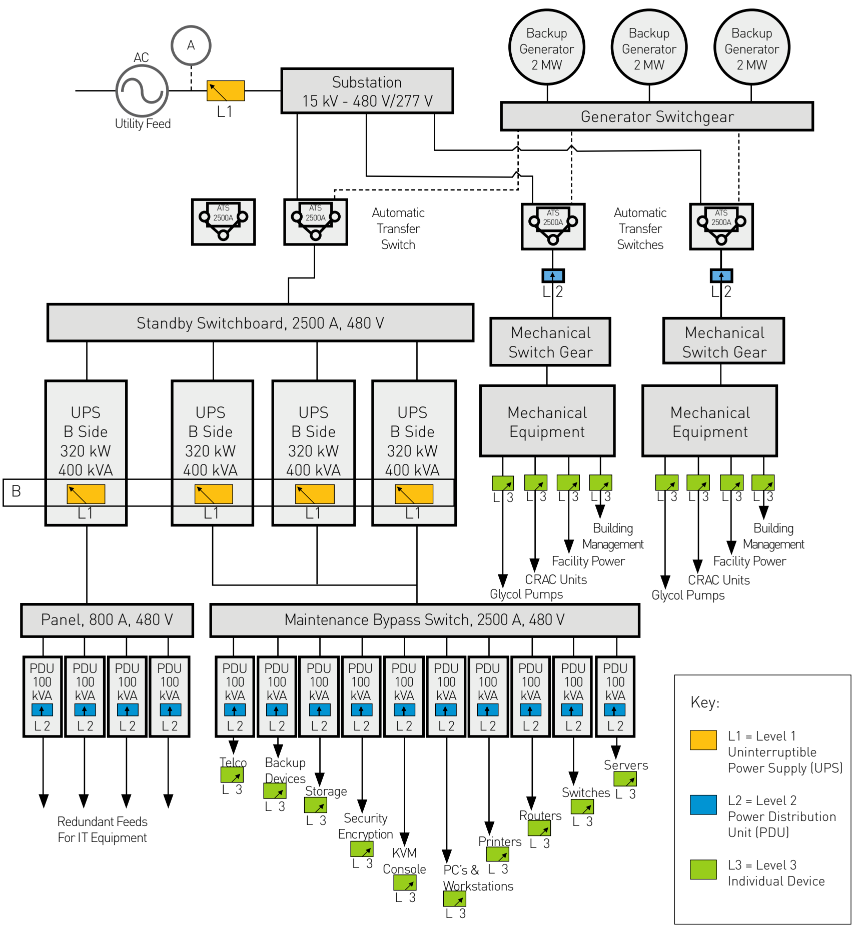

Figure 23: Measuring at the Three PUE Levels

(Source: The Green Grid)

4.4.4 The Advantages of High Power

A single-phase 120V at 100A (80A rated) circuit provides 9.6kVA. A single-phase 208V at 60A (48A rated) circuit provides 10.0kVA. A three-phase 208V at 40A (32A rated) circuit provides 11.5kVA. A single-phase 230V at 60A (48A rated) circuit provides 11.0kVA. A three-phase 400V at 20A (16A rated) circuit provides 11.1kVA.

Running higher voltages at lower currents means smaller cables which use less copper, weigh less, take up less space and cost less. Running three-phase power instead of single-phase power means fewer cables which simplifies deployment, and provides the benefits of smaller cables, less copper, less weight, and less cost.

Plugs and receptacles are also less expensive at higher voltages and lower current ratings. For example, a 30A 400V three phase Wye (16.6kVA) plug (Hubbell NEMA L22-30P) costs $32 and the receptacle costs $41. A 60A 208V three phase Delta (17.3kVA) plug (Mennekes IEC309 460P9W) costs $166 and the receptacle costs $216. The plug/receptacle combination is $73 vs. $382.

There are other benefits to higher voltages. A 400V power circuit will eliminate voltage transformations, and can reduce energy costs by approximately 2-3% relative to 208V distribution, and approximately 4-5% relative to 120V distribution.

Consolidating data centers will generally reduce total power consumption and may create opportunities for the use of high-density racks and high power rack PDUs. For example, a 42U rack filled with 1U servers consuming 250W each draws 10.5kW which would require two three phase 208V, 50A circuits providing 14.4kVA each. Taking advantage of blade servers might lead to deploying four blade chassis in one rack which would require two three phase 208V, 80A or two three phase Wye 400V, 50A rack PDUs. These examples allow sufficient headroom should one of the feeds fail. They also support the North American requirement for 80% derating.

High-density racks can be effectively deployed in small, medium, or large data centers. Even small data centers will benefit from racks delivering high power for multiple blade servers or densely packed 1U servers.

Chapter 4 Summary

- The simplest power deployment to an IT equipment rack is a single appropriately-sized power feed to a single rack PDU.
- The next step up in availability is still a single feed to a single rack PDU, but with the addition of a transfer switch, that typically has two feeds from the same or different building feeds.
- The most reliable deployment is to use two power feeds to two rack PDUs, so if one rack PDU or power feed fails there is a second one available to maintain power to the IT equipment in the rack.
- IT equipment with two or more power supplies can vary in the way power is delivered to the equipment, often requiring load balancing.
- Load balancing attempts to evenly distribute the rack equipment's current draw among the PDU's branch circuits so that as you come closer to perfect balance more total current can be supplied with the greatest headroom in each branch circuit.
- Determining how much power should be deployed to a rack requires you to know device voltage and ideally, actual current draw.
- The use of three-phase power enables one whip or rack PDU to deliver three circuits instead of just one circuit. One slightly thicker three-phase cable will be significantly smaller and weigh less than three single-phase cables for the same voltage and amperage.
- There are many factors involved in selecting a rack PDU: data center location, application requirements, IT equipment requirements, available power, space in the cabinet, energy management, and efficiency objectives.
- Deployment of intelligent PDUs can make a significant difference in the ability of IT administrators and facilities managers to improve uptime and staff productivity, efficiently utilize power resources, make informed capacity planning decisions, and save money.
- The most commonly used metric for data center power efficiency is The Green Grid's Power Usage Effectiveness or PUE. The best practice is to gather data over time of "typical' power usage to ensure that the peaks and valleys have been captured in calculating your PUE to establish a baseline and to track your improvements.
- Running higher voltages at lower currents means smaller cables which use less copper, weigh less, take up less space and cost less. Running three-phase power instead of single-phase power means fewer cables which simplifies deployment as well as the benefits of smaller cables, less copper, less weight and less cost.
- High-density racks can be effectively deployed in small, medium, or large data centers. Even small data centers will benefit from racks delivering high power for multiple blade servers or densely packed 1U servers.

5.0 **Future Trends for Rack PDUs**

5.0 **Future Trends for Rack PDUs**

5.1 Higher Density, Higher Power Rack PDUs with Sensors
5.2 Increased Intelligence at the Rack to Support Efficiency Initiatives
5.3 Integration with Higher Level Data Center Management Systems

Two primary forces are influencing rack PDU development and innovation trends. First is the demand for increasing power and density of IT equipment at the rack or compute density per U of rack space. Second is the industry-wide goal, even mission, to create energy efficient (often called "green") data centers, including carbon footprint reduction. Both trends challenge PDU vendors to improve both hardware and software design; and the second requires all IT and facilities organizations to better understand how the data center power is consumed and take active measures to reduce it.

The above trends are underscored in the IHS November 2013 report "The World Market for Rack Power Distribution Units – 2013 Edition." Compound Annual Growth Rate (CAGR), for rack PDUs between 2012-2017 is predicted to be:

- <5kW are expected to have unit growth of 4.4% (2.0% for the Americas).
- 5kW – 10kW are expected to have unit growth of 5.9% (5.4% for the Americas).
- >10kW are expected to have unit growth of 6.4% (7.5% for the Americas).

When comparing the growth rate of single-phase PDUs to three-phase PDUs, analysts predict:

- Single-phase PDU units are expected to grow 4.4%.
- Three-phase PDU units are expected to grow 7.9%.

When comparing the growth rate of non-intelligent PDUs to intelligent PDUs, analysts predict:

- Non-intelligent rack PDUs are expected to have unit growth of 4.1%.
- Intelligent rack PDUs are expected to have unit growth of 8.4%.

The factors that are driving PDU demand are increasing power consumption and higher densities along with increasing the need for PDU intelligence.

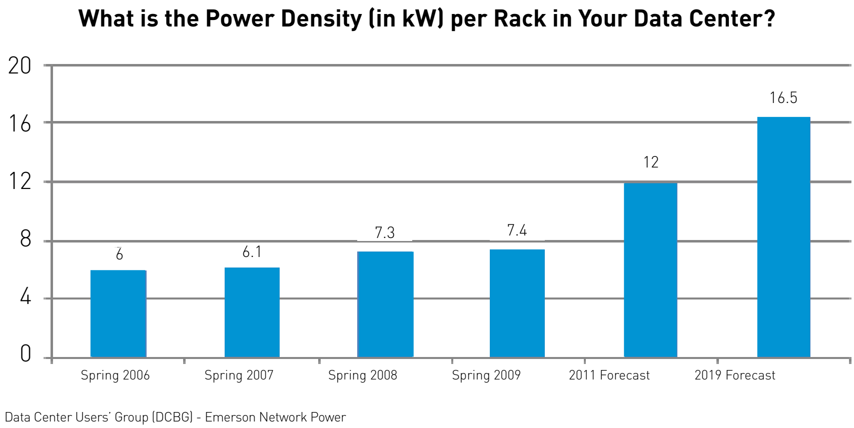

Figure 24: Average Power per Rack
(Source: Emerson Network Power)

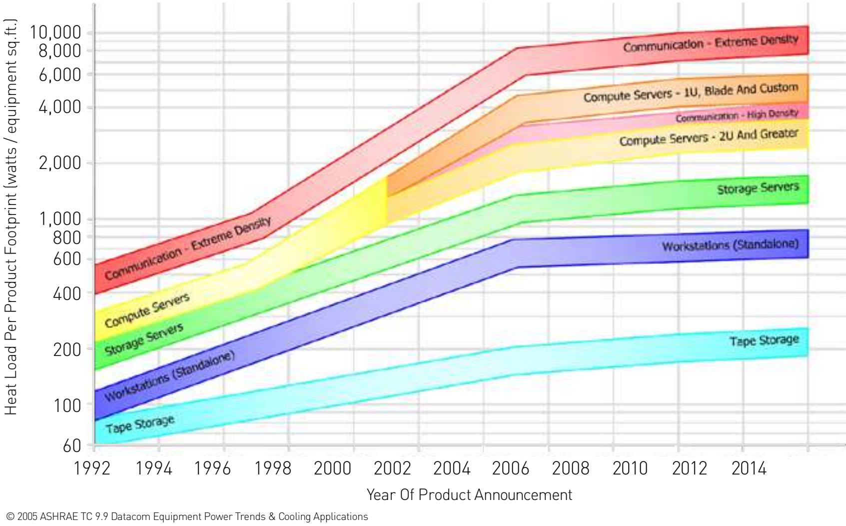

Figure 25: Equipment Power Density Trends

(Source: ASHRAE)

5.1 Higher-Density, Higher-Power Rack PDUs with Sensors

The growing popularity of 1U servers, blade servers, network-attached storage, storage area networks (SANs) and multi-gigabit, chassis-based network communications gear, place enormous demands on rack PDUs. For example, 4 blade server chassis in a single rack could draw in excess of 20kW of power creating both power and cooling challenges for data center managers. From a power perspective, racks will require three-phase power with 60A, 80A, even 100A of service. There are some data centers bringing 400 volt three-phase service to the rack, to accommodate this power demand while increasing efficiency from reduced voltage step downs. Similarly, end users are packing dozens of 1U servers into a single rack and pressing rack PDU vendors to design PDUs to support 40+ outlets and 20+ kW.

Server virtualization is another major trend in data centers and should lead to improved efficiency and cost reduction. However, running multiple virtual machines on one server will drive up its total power consumption; and a rack containing several such servers could potentially experience a lot more power consumption, thereby driving the need to gain additional visibility into power loads of servers in the rack to optimally manage power capacity.

More power consumption means more cooling to remove the additional heat. PDU vendors will be expected to supply the basic environmental sensors for heat, humidity, and air flow to help understand the overall environmental conditions and to identify zones that must be fine-tuned or supplemented with dedicated or specialized cooling.

5.1.1 Customizing IT Equipment for Power Efficiency

One trend to watch, especially for the largest data centers, is the design and deployment of custom servers, power supplies, rack PDUs, etc. to maximize power usage efficiency. For example, Facebook along with OpenCompute have begun to deploy 480V three-phase Wye power where each line is wired to the neutral so the outlets deliver 277V. This approach is very efficient but highly specialized since most IT equipment today is not built with power supplies that support 277V. Furthermore, common data center receptacles are IEC C-13 and C-19, which do not support 277V.

The savings and efficiencies (1% - 2% over 400V/230V three-phase systems) are enough that Facebook/ OpenCompute can justify building custom triplet racks, custom servers with custom power supplies, custom battery/UPS, and 480V/277V rack PDUs with custom Tyco 3-pin Mate-N-Lock outlets.

5.2 Increased Intelligence at the Rack to Support Efficiency Initiatives

Many data centers have grown larger and more complex in recent years as the consolidation trend continues. With increasing size and complexity there is a greater need to drive intelligence to the IT equipment at the rack to create what industry people are beginning to think of as the "Intelligent Rack."

Every data center, regardless of size, is designed to support the servers at the rack where the actual computing is taking place. It is also where the vast majority of the power is or should be consumed. Proper monitoring and metering of the IT equipment along with environmental sensors at the rack will collect the data necessary to produce the most significant overall efficiency, savings and operational improvement. Collection and analysis of actual energy data will enable you to maximize the use of current resource capacity and take advantage of capacity planning tools to "right-size" the data center for future requirements. This will allow you to eliminate or defer capital expenses of data center expansions; while improving day-to-day energy efficiency and overall IT productivity.

Capacity planning based on nameplate data is no longer sufficient. Efficiency improvement is an information-driven activity. In order to formulate and drive the most effective decisions, you will need to collect IT device CPU utilization and their corresponding actual power usage. More energy efficiency will be gained if such planning is based upon the trends observed from the actual data over time. Furthermore, the actual data collected at the rack level can be integrated with the overall data center infrastructure management (DCIM) systems and data center energy management systems for complete data center and power chain visualization, modeling, and planning, which can lead to further improvements in the data center ecosystem, e.g. computing carbon emissions generated by IT devices to report on and take steps to lower your carbon footprint.

Efficiency can also be gained from software that offers policy-based power control to automatically turn servers on and off based on granular power consumption data and a set of pre-established static or even dynamic rules. These power saving applications are finding early adopters in development labs, web server farms and cloud computing environments. Eventually they may find their way into mainstream data centers and their success will be based to some extent on the deployment of intelligent PDUs to enable their functions.

Creating energy-efficient behavior throughout your organization is a key factor in reducing waste and costs; and the essential ingredient to affect behavior is individual awareness and accountability for energy usage. Of course, to be effective, any such energy reporting or charge-back system must be based on credible, comprehensive, and coherent usage data, so PDU vendors will be expected to deliver the highest accuracy for energy usage at every level of the organization.

5.3 Integration with Higher Level Data Center Management Systems

In recent years, a variety of software products have been introduced to help both IT and Facilities people better manage the data center. While the category name may differ—Physical Infrastructure Resource Management (PIRM), Data Center Infrastructure Management (DCIM), Data Center Service Management (DCSM)—these applications provide most of the following major functions: database of all physical data center assets with detailed data for IT, power and HVAC equipment, physical data center layout, and cable connections; change management; 2D or 3D visualization of the data center building with drill down to lowest level data element; and capacity planning based on availability of floor and rack space, power, cooling, etc.

The data required to manage data center infrastructure and energy effectively are collected from power devices along the entire power chain, from IT devices themselves, environmental sensors, data center layout maps, cable plans, and cooling system design documents. The more data collected, the more complete the information will be, and the better equipped data center personnel will be to support critical IT operations reliably, efficiently, and cost-effectively.

The following is a simplified view of data measurement, collection, compilation, analysis and correlation and decision support:

- Intelligent rack PDUs measure essential power data at a predefined frequency and store such data in memory.
- The data collection services from the rack PDU management (or power management) system polls the intelligent PDUs through industry standard management protocols such as SNMP.
- The data collection service can be part of the intelligent PDU vendor's rack PDU management system, or it can be part of a DCIM or energy management system. For scalability reasons, data collection is typically delegated to specific PDU vendor's PDU management system, which is deployed along with intelligent PDUs to administer, maintain and troubleshoot the PDUs; as well as to collect power statistics from these intelligent PDUs.
- The rack PDU management system can use the collected data to perform the first level of analysis. This will help to reveal power trends and pinpoint any potential issues. The amassed data can be used by the DCIM or another energy management system for further analysis.
- The energy management or DCIM system provides visibility beyond the scope of the PDU management system including: poll information from upstream smart power devices, physical layout, cable plan, HVAC deployment information, etc. making it more suitable for analysis that must take into consideration many more factors beyond the intelligent PDUs.

With the advanced analysis conducted by a DCIM or energy management system, data center management staff can make their day-to-day operational decisions as well as longer term strategic planning, to provide reliable power for business applications while reducing waste in data center energy consumption.

Chapter 5 Summary

- Two factors influencing rack PDU development and demand are the increasing power and density of IT equipment at the rack, and the industry-wide goal of creating energy efficient data centers. These are driving the need for PDU intelligence.
- The growing popularity of 1U servers, blade servers, network-attached storage, storage area networks (SANs) and multi-gigabit, and chassis-based network communications gear, place enormous demands on rack PDUs and create both power and cooling challenges.
- More power consumption means more cooling to remove the additional heat. PDU vendors will be expected to supply environmental sensors for heat, humidity and air flow to help understand the overall environmental conditions and to identify zones that must be fine-tuned or supplemented with dedicated or specialized cooling.
- Many data centers have grown larger and more complex in recent years as the consolidation trend continues. With increasing size and complexity there is a greater need to drive intelligence to the IT equipment rack to create what industry people are beginning to think of as an "Intelligent Rack."
- Proper monitoring and metering of the IT equipment along with environmental sensors at the rack will collect the data necessary to produce the most significant overall efficiency, savings and operational improvement.
- An essential ingredient of energy efficient behavior is individual awareness/accountability for energy usage. Billing grade accuracy and billback reports can be driven by intelligent PDUs and energy management software.
- Data collected at the rack-level can be integrated with a data center infrastructure management (DCIM) system for complete data center and power chain visualization, modeling, and planning, which can boost efficiency and lower carbon footprint.

Reference

Reference

Key Terms to Know

4-wire: A 4-wire rack PDU consists of one ground wire and three lines (see Three phase Delta below), each line carrying equal voltage but each voltage sine wave is 120 degrees out of phase with the others. The voltage of two lines is available (line to line, e.g., L1-L2).

5-wire: A 5-wire system is the same as the 4-wire system but with the addition of a neutral wire (see Three phase Wye below) so that the voltage of one line can be supplied (line to neutral) as well as the voltage of two lines (line to line).

Active Directory®: is a directory service that authenticates and authorizes all users and computers in a Windows domain type network—assigning and enforcing security policies for all computers, and installing or updating software.

Active Power: The real power drawn by an IT device. This determines the actual power purchased from the utility company and the heat loading generated by the equipment since for IT equipment, 1W of electricity equals 1W of heat.

Air Pressure: Air pressure that is too high will result in both higher fan costs and greater leakage which can short circuit cooling air, while pressure that is too low can result in hot spots at the areas most distant from the cool air supply point.

Airflow: There are two primary areas for monitoring airflow in the data center: above the floor and below the floor. Differential airflow sensors are used to ensure that the pressure differential between the subfloor and the floor is sufficient to control air flowing from the subfloor to the floor above. Blockages in under-floor supply plenums can cause high pressure drops and uneven flow, resulting in cold spots in areas where cooling air is short circuiting to the return path.

Amperage: Also known as the current. It is the flow or rate of flow of electrons, ions, or holes in a conductor or medium between two points having a difference in potential, measured in amperes and equal to the ratio of the voltage to the resistance.

Apparent Power: Expressed in volt-amperes (VA). The product of the voltage applied to the IT device times the current drawn by the equipment. The VA rating is used for sizing wiring and circuit breakers. The apparent power is always equal to or larger than the active power.

ASHRAE: The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE). ASHRAE recommends temperature probes be placed at specific locations at the inlets of equipment rack to monitor temperatures.

Basic PDU: Basic PDUs are power strips that are constructed out of high-quality components for use in critical environments such as data centers. They distribute correct voltage and current to multiple outlets.

Branch Circuit: Branch circuits are power feeds that originate from a panel, switch or distribution board and terminate into an electrical receptacle mounted in a junction box near the IT equipment rack.

Circuit Breaker: An automatically operated electrical switch designed to protect an electrical circuit from damage caused by overload or short circuit. Circuit breakers are available on several PDU models.

Delay Curve: Gives a circuit breaker the ability to discriminate between normal and damaging over-currents, which are quite common when powering servers.

Delta: This configuration gets the name Delta because a schematic drawing of it has three transformers forming a triangle or the Greek letter Delta. The three lines connect to the three "corners" of the triangle.

Derated: In North American data centers must be certified to UL 60950-1. UL 60950-1 limits a device to draw no more than 80% of the rating of its input plug. This 80% limitation is commonly known as "derated" current.

Energy Management: Driven by industry-wide goal of creating energy efficient data centers and reducing carbon footprint. Requires information that leads to smarter energy usage and more energy efficient behavior.

Fuses: A type of low resistance resistor that acts as a sacrificial device to provide overcurrent protection of either the load or source circuit. Not ideal for data center energy distribution as it may increase the mean time to repair (MTTR).

Ground: A conducting body, such as the earth or an object connected with the earth, whose potential is taken as zero and to which an electric circuit can be connected. The purpose of a ground wire is to safely direct stray currents to ground rather than allowing them to pass through someone contacting the stray currents.

Humidity: A measurement of moisture in the air. High humidity can cause condensation buildup on computer components, increasing risks of shorts. Likewise, if the humidity is too low, data centers can experience electrostatic discharge (ESD).

IEC: The International Electrotechnical Commission creates classifications for plugs and receptacles used in data centers.

Inline Meter: Device that allows users to monitor power remotely at the rack PDU inlet or circuit-level.

Intelligent PDU: PDUs that can be accessed and controlled remotely via a web browser or command line interface (CLI). They meter power at both the PDU and individual outlet-level; support alerts based on user-defined thresholds; provide security in the form of strong passwords, authentication, authorization and encryption; and incorporate rich environmental management capabilities.

Kilowatt-Hour (Kwh): A kilowatt-hour (kWh) is a unit of electrical energy or work, equal to the power supplied by one kilowatt for one hour.

Lightweight Directory Access Protocol (LDAP): An open, vendor-neutral, industry standard application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network.

Line: An electrical conductor which is a source of voltage, e.g., 120V. In a single-phase system there are one or two lines. In a three-phase system there are three lines.

Load Balancing: A procedure which attempts to evenly distribute the rack equipment's current draw among the PDU's branch circuits so that as you come closer to perfect balance, more total current can be supplied with the greatest headroom in each branch circuit.

Metered PDU: PDUs that measure the current draw (load) at the PDU level, and display the data locally. More sophisticated models also offer user-defined alarm functions and remote access to the data over a serial or network port.

Nameplate: Nameplate data is the electrical power consumption information specified by the equipment manufacturer. It is typically a conservative estimate of the maximum amount of power the device could draw.

NEC: The National Electrical Code is a regionally adoptable standard for the safe installation of electrical wiring and equipment in the United States. It has specific requirements for data centers.

NEMA: The US National Electrical Manufacturers Association is the organization that sets standards for electric outlet and receptacles found in data centers.

Neutral: An electrical conductor which provides a return path for the voltage supplied by a line. The neutral itself is not a source of voltage, e.g., 0V.

Power Distribution Unit (PDU): Large data center PDUs are used earlier in the power chain and take the form of panel boards mounted on walls or free standing pedestals. In this handbook we're discussing only the rack PDU, at the end of the chain, which supplies power to the IT equipment in the rack.

Power Factor: Defined as the ratio of the real power flowing to the load, to the apparent power in the circuit, and is a dimensionless number between -1 and 1.

Power Usage Effectiveness (PUE): A measure of how efficiently a data center uses energy; specifically, how much energy is used by the computing equipment. Defined as PUE = Total Facility Power / IT Equipment Power.

Rack PDU: A rack PDU is mounted in an IT equipment rack and provides electrical power to various IT devices such as servers, networking, and storage equipment (See Section 2.1.1).

Rated: The Rating of an electrical appliance indicates the voltage at which the appliance is designed to work and the current consumption at that voltage.

Simple Network Management Protocol (SNMP): An internet-standard protocol for managing devices on IP networks. It is used mostly in network management systems to monitor network-attached devices for conditions that warrant administrative attention.

Single-Phase: The distribution of alternating current electric power using a system in which all the voltages of the supply vary in unison.

Switched PDU: Switched PDUs offer the features of metered PDUs and also provide controlled on/off switching of individual outlets and load metering (see metered PDUs) at the PDU level. They enable authorized users to securely power cycle devices remotely; and they may also provide a power sequencing delay as well as some outlet use management.

Temperature: In the data center, temperatures must be kept at predetermined set points in order to keep sensitive IT equipment from overheating, and prevent costly overcooling.

Three-Phase: Three-phase electric power systems have at least three conductors carrying alternating current voltages that are offset in time by one-third of the period.

Underwriters Laboratories (UL): A safety consulting and certification company that drafts several standards that affect electricity in the data center.

Voltage: Electromotive force, or difference in electrical potential, measured in volts and equal to the current times the resistance.

Volt-Amp (VA): The product of the voltage applied to the IT device times the current drawn by the equipment. The VA rating is used for sizing wiring and circuit breakers. The apparent power is always equal to or larger than the active power.

Watt: The real power drawn by an IT device. This determines the actual power purchased from the utility company and the heat loading generated by the equipment since for IT equipment, 1W of electricity equals 1W of heat.

Watt-Hour (Wh): A unit of energy equivalent to one watt of power expended for one hour of time.

Wye: An electrical configuration that derives its name from its schematic drawing of three transformers meeting in the center forming the letter "Y". The three lines connect to the three "branches" of the "Y" and the neutral connects to the center of the "Y."

References and Further Reading

Alger, Douglas. *"Build the Best Data Center Facility for Your Business: A comprehensive guide to designing and operating reliable server environments,"* Indianapolis (IN): Cisco Press, 2005.

ASHRAE., *"Thermal Guidelines for Data Processing Environments,"* 2nd ed. Atlanta (GA): ASHRAE, 2009.

ASHRAE,. *"ASHRAE Workshops on Improving Data Center Energy Efficiency and Best Practices,"* NYSERDA sponsored workshop; New York (NY) ASHRAE, November 6, 2009.

ASHRAE., *"2008 ASHRAE Environmental Guidelines for Datacom Equipment,"*. Atlanta (GA): ASHRAE, 2008.

ASHRAE., *"Best Practices for Datacom Facility Energy Efficiency,"* 2nd ed. Atlanta (GA): ASHRAE, 2009.

ASHRAE,. *"High Density Data Centers: Case studies and best practices,"*. Atlanta (GA): ASHRAE, 2008.

Haas, Jon, Monroe M, Pflueger J, Pouchet J, Snelling P, Rawson A, Rawson F. et al,. *"Proxy Proposals for Measuring Data Center Productivity,"*. White Paper #14,. Beaverton, (OR): The Green Grid, 2009.

Raritan Inc., *"Data Center Power Overload Protection: Circuit breakers and branch circuit protection for data centers,"*. White Paper,. Somerset (NJ): Raritan Inc., Somerset, NJ, 2009a.

Available at http://www.raritan.com/high-power-white-paper. Accessed February, 2014.

Raritan Inc., "*Data Center Power Distribution and Capacity Planning: Understanding what you know – —and don't know – —about power usage in your data center.,"* White Paper,. Somerset (NJ): Raritan Inc., Somerset, NJ,; 2009b.

Available at http://www.raritan.com/power-distribution-and-capacity. Accessed May 22, 2014.

Raritan Inc., *"Power Distribution Units (PDUs): Power Monitoring and Environmental Monitoring to Improve Uptime and Capacity Planning,"*. White Paper,. Somerset (NJ): Raritan Inc., Somerset, NJ,; 2009c.

Available at http://www.raritan.com/power-and-environmental-monitoring. Accessed May 22, 2014.

Raritan Inc., *"Deploying High Power to IT Equipment Racks,"*. White Paper #V1156,. Somerset, (NJ): Raritan Inc., Somerset, NJ, 2012.

Verdun, Gary. Editor. *"The Green Grid Metrics: Data Center Infrastructure Efficiency (DCiE) Detailed Analysis,"*. White Paper #14,. Beaverton, (OR): The Green Grid, 2008.

UL 60950-1, *"Information Technology Equipment – Safety – Part 1: General Requirements"*, Underwriters Laboratories, Inc. 2013.

To learn more about Raritan PDUs, visit: www.raritan.com/ipdus

Index

Raritan
A brand of legrand

Understanding Rack PDUs with Ease

The Intelligent Power Distribution Unit (iPDU) Handbook is the definitive guide to rack PDUs and understanding how they can change the face of your data center. Written with the rack PDU novice in mind, this simple but elegant illustrated guide will cover the fundamental principles of powering your IT equipment using single or three-phase power while leveraging the advanced capabilities of different rack PDU models in order to increase uptime and availability while improving power efficiency.

Call 1.800.724.8090 or visit www.raritan.com