Raritan.
A brand of legrand

Server Technology.
A brand of legrand

The iPDU Handbook

A DEFINITIVE GUIDE TO INTELLIGENT RACK POWER DISTRIBUTION

THIRD EDITION

The iPDU Handbook

THIRD EDITION

A DEFINITIVE GUIDE TO INTELLIGENT RACK POWER DISTRIBUTION

Contents

1.0 Introduction

The intelligent rack power distribution unit (rack PDU) has revolutionized data centers. As the last link of the elaborate data center power chain, the traditional role of the rack PDU is to deliver stable and reliable power to devices in the rack. It provides the electrical heartbeat to the systems running the critical applications that support the operation of a business (or that, in some cases, is the business). In years past, rack PDUs were often considered a simple commodity, just power strips to plug devices into. Typically, IT staff merely told facilities how much power was needed at the rack based on device nameplate specifications. The thought was that this information would provide plenty of headroom and ensure minimal risk of downtime. Little thought was given to the added value the rack PDU could provide until recently.

Over the past few years, system availability has become a given, and data center management attention is now focused on managing operational costs, efficiency improvements, sustainability, and resource optimization. With the annual expenditure for powering the average data center surpassing the cost to buy the IT equipment itself, the efficient use (and waste) of energy is now prioritized. Beyond the actual cost to power a data center, related issues affect current operations and future expansion—e.g., physical space limitations, utility power availability, carbon footprint impacts, capacity, and potential government regulation. Because most of the power delivered from the utility to the data center is consumed 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 PDU is now pivotal to a data center's management.

Not surprisingly, many of the major strategies to address the above issues and improve overall data center efficiency now depend on the capabilities of an intelligent rack PDU, which are not available in commodity off-the-shelf power strips.

Let's consider a few capabilities that an appropriately sized rack PDU can provide:

- To maximize data center space and power resources, trends are to deploy racks densely packed with 1U (1.75 in/44.45 mm tall) servers or power-hungry blade servers. Today's rack PDUs typically handle loads of 5-10 kW (17-22 kW for enterprise data centers) with 36 outlets. There are now rack PDUs designed to support 50+ kW and 54+ outlets.
- Many data centers are employing lights-out and/or remote operations to increase IT staff productivity and conserve energy. Some rack PDUs provide real-time monitoring, reporting, alerting, and secure and reliable outlet switching.
- To identify ghost (no definable function), underutilized, or grossly inefficient servers for elimination, replacement, consolidation, virtualization, or containerization, some rack PDUs provide individual outlet monitoring.
- To create individual awareness, accountability, and/or chargeback reports for power usage and carbon footprint impacts, some rack PDUs are now equipped with highly accurate, real-time power measurement capabilities at the PDU and outlet levels.
- To optimize IT workloads and make informed decisions for infrastructure capacity planning, the interoperability of a rack PDU with data center infrastructure management (DCIM) software will continually collect data on power consumption, analyze trends, and correlate with IT workload data.
- To enable a "follow-the-sun" model of operation used by facilities run directly from solar Photovoltaic (PV) arrays, , some rack PDUs are programmable to shut down connected equipment and resume operation based on energy availability.

Why Has the Selection of Rack PDUs Become So Important?

Various rack PDU configurations are available based on the number of phases, voltage, total amps, branch circuits, number of outlets, socket type, plug type, rack units (RU) consumed, and physical dimensions. Beyond the essential functions of the rack PDU, added capabilities are available in rack PDU categories that we call monitored/metered, smart, and switched. As power demands within a single rack vary, the job of a rack PDU to distribute the right power to the right equipment becomes more complex. Suppose you cannot find an off-the-shelf rack PDU matching your requirements. In that case, some vendors will assemble or even design and engineer a custom rack PDU—also called BTO (built-to-order), CTO (configured-to-order), or ETO (engineered-to-order) models—to fit your unique needs best. However, the decision on which rack PDU to deploy becomes apparent upon closer inspection. The specifications of your rack PDU will directly influence your data center's reliability, flexibility, and scalability.

This handbook will discuss the basic concepts, considerations, and approaches in designing, selecting, and deploying an appropriately sized intelligent rack PDU for typical data center applications. We will describe the fundamental ingredients for delivering reliable power to the rack and the factors and best practices contributing to a stable, operationally efficient, and environmentally sound data center for today and the future.

Throughout this handbook, the content will stay focused on rack PDUs—the PDU that sits at the end of the power chain that supplies power directly to IT equipment in a rack. The larger end-of-row data center floor PDUs used earlier in the power chain, chain, typically consisting of panelboards mounted on walls or are free-standing pedestals, will not be discussed. Unless otherwise specified, any reference to "PDU" for the rest of this handbook means "rack PDU".

2.0 **Fundamentals and Principles**

PDUs come in many configurations for the number and type of receptacles, voltage, load capacity, and physical mounting (horizontal or vertical) options. A PDU may provide power to the devices plugged into it or offer added functions such as turning the power off and on remotely down to the PDU's outlet level, monitoring power consumption, and sensing the environmental conditions in and around a rack.

2.1 Overview and Class of Devices

A rack PDU is commonly called a rackmount power strip. Not to be confused with a standard off-the-shelf power strip, a rack PDU distributes stable, reliable power and is designed for industrial-grade use. It is meant to protect the devices it is powering and to monitor, control, and manage electrical power at the rack. Available in several configurations, the typical characteristics of PDUs are described below.

2.1.1 Types of Rack PDUs

PDUs are generally divided into two main categories: Non-Intelligent PDUs and Intelligent PDUs, commonly called iPDUs. The following definitions are classifications of PDUs most used in the industry, which could differ from vendor to vendor.

Non-Intelligent PDUs

- **Basic PDUs** provide a fixed number of receptacles that distribute reliable voltage and current to power IT equipment and may or may not have branch circuit protection, such as circuit breakers or fuses. It is considered non-intelligent because it cannot be accessed remotely as the device has no network connectivity capabilities, nor does it provide insight into equipment power usage. Basic PDUs are ideal for small server rooms that are physically monitored by IT staff. However, they would not be recommended for larger, mission-critical server rooms or data centers, given their inability to be monitored or managed remotely.
- **Monitored/Metered PDUs** are built on the features in Basic PDUs and have added functionality to display current and other electrical information locally. However, this information cannot be accessed remotely as these PDUs have no network connectivity capabilities. Monitored/Metered PDUs are recommended for use in highly secure data center environments that are required to keep their power infrastructure separate from local networks.

Intelligent PDUs

- **Smart PDUs** have the added value of offering PDU-level input power measurement capabilities and branch circuit protection, allowing you to measure, monitor, and report power at the rack or PDU. It is available in various voltages and amperages and serves as a central connection point for environmental monitoring, asset location, physical access, and other monitoring and security sensors. Smart PDUs are recommended for high-density data centers, extensive colocation facilities that cannot generate manual power usage reports for billing purposes, and cloud providers needing to support varying power loads without compromising energy efficiency.
- **Switched PDUs** have the added value of offering PDU-level input power measurement capabilities and allow you to monitor and control power at the device level. It also features outlet lockout and control, power-up sequencing to reduce power-up inrushes, remotely rebooting critical equipment, and optional smart load shedding. Switched PDUs are recommended for any IT deployment in any data center.

- **Smart PDUs with Per Outlet Power Sensing** builds on the features offered in Smart PDUs with the added value of providing device-level output measurement capabilities, allowing you to measure, monitor, and report power down to the outlet/device level. Email alerts provide automated updates on the power and environmental conditions.
- **Switched PDUs with Per Outlet Power Sensing** have the added value of offering both outlet control and device-level output measurement capabilities, allowing you to measure, monitor, control, and report power down to the outlet/device level. It combines highly accurate, outlet-level power measurement technology for data center device-level power monitoring, alerting, and on/off/reboot control. With outlet control, gain features like power-up sequencing and smart load shedding.

	Non-Intelligent PDUs		Intelligent PDUs			
Core Features	**Basic**	**Monitored/ Metered**	**Smart**	**Switched**	**Smart with Per Outlet Power Sensing**	**Switched with Per Outlet Power Sensing**
Power Distribution	●	●	●	●	●	●
Input Metering	○	●	●	●	●	●
Outlet Metering	○	○	○	○	●	●
Network Connectivity	○	○	●	●	●	●
Outlet 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: Legrand, Inc.)

2.2 Electrical Terminology

Reviewing the following electrical terms before discussing PDUs in more detail will be helpful.

Alternating Current (AC): Is energy delivered in a form that can travel the long distances necessary from generating plants to homes and businesses. The term AC reflects that the voltage and current constantly change in value or alternate between a positive and negative threshold over a centerline.

Direct Current (DC): Is the energy that does not alternate over a fixed period but has a steady value with reference to zero. DC does not travel great distances well.

Voltage (Volt or V): The standard measure of electrical potential. Higher voltages allow more energy to flow within a given time for a given wire size. Electromotive force (E), or difference in electrical potential, is measured in volts and equal to the current (I) times the resistance (R). E = I*R.

Current (Ampere or A or Amp): The flow rate of electric charge. Current is measured in amperes.

Active Power (Watt or W): The work done by an IT device and is the measurement by which power companies bill for its use. Also referred to as real power.

Reactive Power (VAR): Is power flowing back and forth between the power company and electrical devices due to capacitance and inductance, the ability to store energy, within the IT device. Reactive power does no actual useful work and is not billed by power companies.

Apparent Power (Volt-Ampere or VA): The product of the voltage multiplied by current. It is the sum of the active and reactive powers. .

Power Factor: The ratio of active (real) power to apparent power. A power factor of 100 percent indicates perfect power, while lower values indicate wasted power due to reactive power or distortions. Power companies may charge a surcharge if the power factor is below a threshold.

Energy Measurement (kilowatt or 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 moment, 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. For example, a 1000W light bulb left on for one hour or a 100W light bulb left on for ten hours consumes 1kWHr (kilowatt-hour) of energy.

Line: An electrical conductor that 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. Lines are labeled as L1, L2, and L3, or X, Y, and Z.

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, with potential taken as zero, and to which an electric circuit can be connected. The purpose of a ground wire is to safely direct accidental currents to the ground rather than allowing them to pass through someone contacting this current.

4-wire and 5-wire Systems: A 4-wire 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 (see Figure 3). The voltage of two lines is available as line to line, i.e., L1-L2. A 5-wire system is the same as a 4-wire system, but with the addition of a neutral wire (see Three-Phase Wye below). The voltage of two lines is available as line to line, and the voltage of one line (line to neutral) can be supplied as well.

Three-Phase Delta "Δ": This configuration gets the name Delta because a schematic drawing has three transformers, forming a triangle or the Greek letter Delta. The three lines connect to the triangle's three corners (see Figure 3).

Three-Phase Wye "Y": This configuration gets the name Wye because a schematic drawing 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 (see Figure 3).

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 in an electrical receptacle mounted in a junction box near the rack. Depending on a data center's layout, branch circuit wiring can be overhead, underneath a raised floor, or both.

2.3.2 Branch Circuit Load Capacity

The power a branch circuit can deliver depends on the electrical characteristics of the circuit. A critical factor in providing power to a rack is whether the power is single-phase or three-phase. The amount of electricity delivered to a rack is often called load capacity. Load capacity 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 a branch circuit can deliver 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

A branch circuit's rated voltage specifies its magnitude (volts) and the number of phase conductors. Single-phase wiring is straightforward and consists of two wires (plus safety ground) where the AC voltage is a single sinusoidal wave measured across the two wires.

Three-phase wiring is more complicated and consists of three (three-phase conductors) or four (three-phase and one neutral) wires, plus the safety ground. Three-phase branch circuits deliver more power and require a PDU specially designed for three-phase branch circuits. Internally, a three-phase PDU divides the 3 or 4 branch circuit wires into pairs of single-phase circuits—and these single-phase circuits are wired to the 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 a neutral. Like the single-phase power described above, connecting across one 120V hotline and the neutral provides 120V AC. Connecting across any two 120V hotlines, 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 PDUs use a neutral wire to provide three circuits of both 120V and 208V. Regardless of the number of wires, a three-phase PDU is rated at the voltage between two phases, e.g., L1 and L2, which in the example here is 208V.

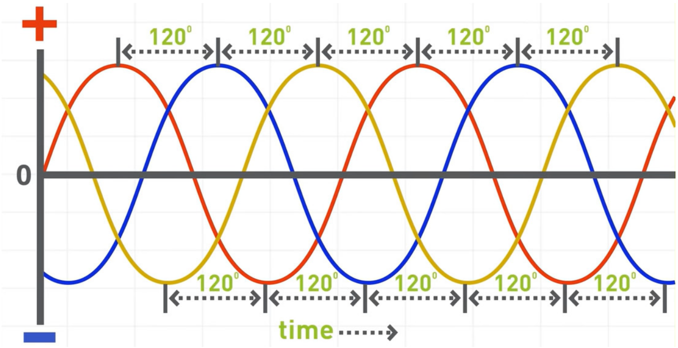

Figure 2: Three-Phase Sinusoidal Waves

(Source: Legrand, Inc.)

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

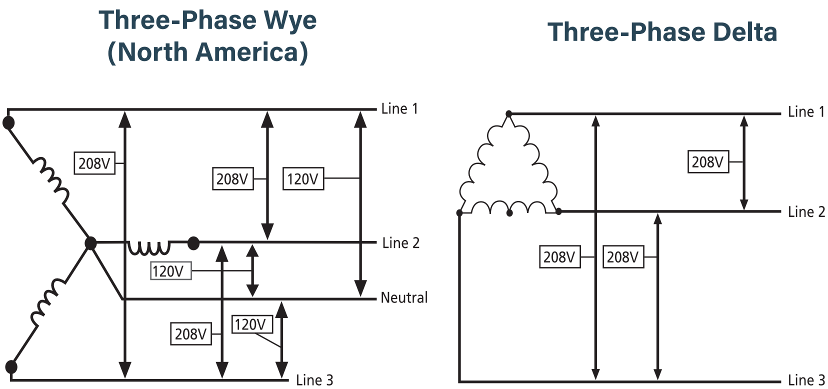

Figure 3: Three-Phase Wiring Diagram

(Source: Legrand, Inc.)

When looking at three-phase PDU specifications, you will often see the terms Wye and Delta or the Greek letters Y and Δ. These terms/letters were chosen because the electrical configuration diagram of a Delta transformer looks like a Δ. The electrical configuration diagram of a Wye transformer looks like a Y. A PDU which does not allow for a higher input voltage, e.g., 208V or 400V, to a lower output voltage reference, e.g., 120V or 230V indicates a Delta transformer is present in the power system. 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 PDU that does permit a higher input voltage to a lower output voltage reference indicates a Wye transformer is in the power system. 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 where the neutral wire is attached. Connecting any two of the three-line connections, 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

3ø indicates three-phase.

Figure 4: Branch Circuit Rated Voltage and Wire Requirements
(Source: Legrand, Inc.)

2.3.4 Branch Circuits: Rated Current

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

Devices using electricity in data centers are classified as IT equipment, and electrical safety standards apply to such devices. Internationally, these standards are IEC 62368-1. In North America, UL (Underwriter's Laboratory) has adopted similar standards as UL 62368-1.

North American safety standards for IT equipment do not allow inlets (plugs) or receptacles to operate at more than 80 percent of their rated current, per the UL 62368-1 standard. For example, a PDU rated for 24 amperes must use an inlet plug rated for 30A not to exceed the limit. This 80 percent limitation is referred to as the derated current or using the device at less than its rated maximum capability.

The following table summarizes the power available for various branch circuits.

Location	Rated Voltage	Nominal Current	Rated/Derated 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
	415V 3Ø			17.2kW
International	230V	32A	32A	7.4kW
	400V 3Ø			22.2kW
North America	208V	60A	48A	10.0kW
	208V 3Ø			17.3kW
	415V 3Ø			34.5kW
International	400V 3Ø	63A	63A	43.6kW
Japan	100V	15A	15A	1.5kW
	100V	16A	16A	1.6kW
	200V			3.2kW
	100V	30A	30A	3.0kW
	200V			6.0kW
	200V 3Ø			10.4kW

3ø indicates three-phase.

Figure 5: Branch Circuit Available Power

(Source: Legrand, Inc.)

2.4 Plugs, Outlets, and Cords

PDUs are available with several plug types. Its receptacles/outlets are designed so that only the proper PDU plug will fit into its respective feeder circuit outlet, and only the appropriate device plug will fit into its respective PDU receptacle. This is done to protect equipment and people, e.g., a device designed for 120V only is not accidentally plugged into a 208V circuit. Also, for safety reasons, for example, a server that draws 30A does not overload a circuit designed to handle only a maximum of 16A. Understanding your data center space and IT equipment needing to be powered before selecting a rack PDU's receptacle type, mounting, and power quality needs can help save costs, time, and space.

Which input plugs and receptacles/outlets are most frequently used in data centers? NEMA (National Electrical Manufacturers Association) and IEC (International Electrotechnical Commission) designate two major data center plugs and receptacles. 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 many families of plugs and receptacles are used in data centers around the world.

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 which could lead to an electrical fire. Note: For North America, the smaller the AWG (wire gauge) number, the greater the diameter of the conductor.

The conductors are surrounded by insulating material and jackets in power cords and cables, which may have special properties. For example, its jacketing may be designed to resist oil exposure damage. 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 (more commonly found in general consumer applications and not widely used in data centers).
- **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.

For a data center where the power is already deployed to the rack, one of the best and easiest ways to determine the PDU's required input voltage is to know the receptacle into which the PDU will be plugged. Knowing the plug indicates the voltage, phase, phase configuration, and amperage.

Here are some of the most common examples of input configurations:

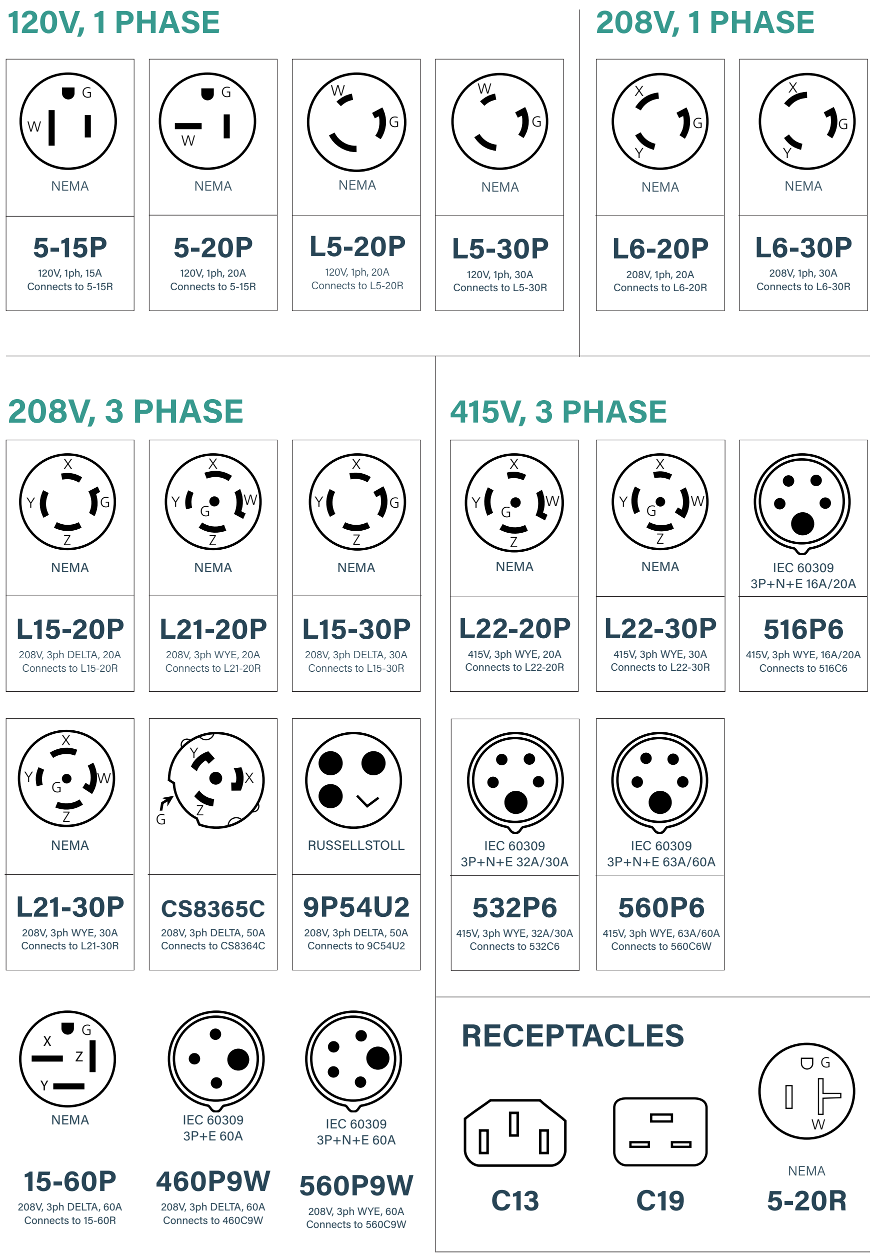

Figure 6: Common IEC/NEMA Input Plugs and Receptacles Used in Data Centers

(Source: Legrand, 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	STANDARD CONTAINER QUANTITY
				INCHES	mm	INCHES	mm			
02763	2	18	16/30	0.030	0.76	0.345	8.76	10	65	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	94	250'
02722	2	16	26/30	0.030	0.76	0.370	9.40	13	77	250'
02765	3	16	26/30	0.030	0.76	0.390	9.91	13	94	250'
02766	4	16	26/30	0.030	0.76	0.420	10.67	10	114	250'
02723	2	14	41/30	0.045	1.14	0.510	12.95	18	154	250'
02762	3	14	41/30	0.045	1.14	0.535	13.59	18	171	250'
02768	4	14	41/30	0.045	1.14	0.575	14.61	15	209	250'
02724	2	12	65/30	0.045	1.14	0.570	14.48	25	168	250'
02725	3	12	65/30	0.045	1.14	0.595	15.11	25	223	250'
02726	4	12	65/30	0.045	1.14	0.650	16.51	20	276	250'
02767	2	10	104/30	0.045	1.14	0.620	15.75	30	230	250'
02728	3	10	104/30	0.045	1.14	0.660	16.76	30	289	250'
02727	4	10	104/30	0.045	1.14	0.715	18.16	25	351	250'
16063	3	8	133/29	0.060	1.52	0.840	21.33	40	450	250'
16064	4	8	133/29	0.060	1.52	0.945	24.00	35	580	250'
16065	5	8	133/29	0.060	1.52	1.030	26.16	28	700	250'
16073	3	6	133/27	0.060	1.52	0.980	24.89	55	637	250'
16074	4	6	133/27	0.060	1.52	1.080	27.43	45	830	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	926	250'
16084	4	4	133/25	0.060	1.52	1.260	32.00	60	1145	250'
16085	5	4	133/25	0.060	1.52	1.365	34.67	48	1419	250'
16093	3	2	133/23	0.060	1.52	1.330	33.78	95	1367	250'
16094	4	2	133/23	0.060	1.52	1.460	37.08	80	1699	250'
16095	5	2	133/23	0.060	1.52	1.580	40.13	64	2066	250'

AWG-TO-METRIC CONVERSION CHART			
SIZE (AWG)	mm^2	SIZE (AWG or kcmil)	mm^2
18	0.82	**1/0**	53.5
16	1.31	**2/0**	64.4
14	2.08	**3/0**	85.0
12	3.31	**4/0**	107.0
10	5.26	**250**	127.0
9	6.63	**300**	152.0
8	8.37	**350**	177.0
6	13.30	**500**	253.0
4	21.15	**600**	304.0
2	33.62	**750**	380.0
1	42.40	**1000**	507.0

Figure 7: Cord Specifications and AWG-to-Metric Conversion Chart

(Source: The Prysmian Group)

2.5 Ratings and Safety

PDUs, like other electrical equipment, are subject to many general and specific safety standards. General industry terms and conventions should be understood to ensure a reliable and safe data center. These are discussed 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 power a device could draw. This information is found on a label near the device's electrical power input.

Sample Rack Mounted Drawer

Configuration	Condition								
	Typical Heat Release (@ 110V)	Airflow				Weight		Overall System Dimensions (W x D x H)	
		Nominal		Maximum at 35°C					
	watts	cfm	(m^3/hr)	cfm	(m^3/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

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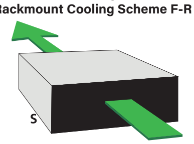	**Description**	**Model**
		Minimum	1-way, 1.5 GHz processor 15 GB memory
		Full	2-way, 1.65 GHz processor maximum memory
		Typical	1-way, 1.65 GHz processor 15 GB memory

Figure 8: Nameplate Data

(Source: HP)

2.5.2 Power Rating vs. Load Capacity

There can be confusion about power capacities and load capacities. 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 or higher. 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 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 must be loaded to no more than 80 percent 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 percent value, i.e., 16A for a 20A circuit, is often referred to as the derated value or the load capacity. A PDU vendor's specifications sheet in North America may include 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 three-phase Wye.
- **Rated current:** 24A (30A derated to 80 percent).
- **Maximum current draw:** 30A Wye delivers 30A x 3 phases x 0.8 = 72A.

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

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, on a single-phase 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

To meet applicable local and national electric codes, PDUs must operate safely and, for example, not emit electromagnetic radiation as one part of the required agency approvals. Manufacturers contract recognized approval agencies 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 must provide you with the listing number and a copy of the testing report upon request.. You can also send 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 IT system to be grounded. This means all IT equipment within a rack and the metal rack itself must be grounded.

The inlet plug of a PDU has a ground pin. When this plug is connected to a correctly wired receptacle, the PDU becomes the grounding point for the equipment plugged into the PDU. The PDU can also be used to ground the metal rack, and most PDUs have specially threaded holes for this purpose. A grounding wire is typically connected to the rack and the PDU using screws. There are also special grounding screws with teeth under the head of the screw to ensure a good ground connection.

Approval	Description	Standard	Comment
NRTL (e.g., TUV, UL, CSA, ETL)	Safety	UL 62368-1	Required in USA
cUL/CSA	Safety	CAN/CSA-C22.2 No. 62368-1	Required for Canada
IECEE CB Scheme	Safety	IEC 62368-1 IEC 60950-1	Common replacement for UL, CSA & CE in countries that accept CB
CE	Electromagnetic Interference (EMC)	EN55032 CISPR 32 EN55024 CISPR 24 EN55035 CISPR 35 EN61000-3-2 IEC 61000-3-2 EN61000-3-3 IEC 61000-3-3	Europe
CE	Safety	EN 62368-1	Europe
FCC	Electromagnetic Interference (EMC)	FCC Part 15 Subpart B Sections 15.107 & 15.109, Class A	USA
ICES-003	Electromagnetic Interference (EMC)	CAN ICES-003, Class A	Canada
UKCA	Electromagnetic Interference (EMC) & Safety	BS EN 55032 BS EN 55035 BS EN 61000-3-2 BS EN 61000-3-3 BS EN 62368-1	Great Britian
PSE	Safety	DENAN Technical Requirements & Enforcement Regulations Appendix 4, Appendix 10	Japan
KCC	Electromagnetic Interference (EMC)	KN32 KN35	Korea
KCC	Safety	KN 60950-1 KN 62368-1	Korea

Figure 9: Safety and Electromagnetic Approval Agencies

(Source: Legrand, Inc.)

2.6 Overload Protection

North American and International safety standards for IT equipment require overcurrent protection (circuit breakers or fuses) for PDUs when the PDU's rating exceeds its receptacles' rating. For example, a 24A PDU requires internal overcurrent protectors to protect its receptacles. These overcurrent protectors must be branch circuit-rated devices.

The UL 62368-1 standard applies to IT equipment safety and requires branch circuit overcurrent protection for PDUs greater than 20A. Typically, PDUs greater than 20A and certified after April 2003 must have built-in UL 489 circuit breakers or fuses suitable for branch circuit protection. UL 62368-1 permits products at a maximum current of 15A and 20A without circuit breakers or fuses since the building's 15 or 20A circuit breakers are considered sufficient to protect the PDU; however, supplementary protection in the PDU provides added protection.

Newly certified PDUs at more than 20A must use overcurrent protection that meets branch circuit protection requirements per the National Electrical Code (NEC) ANSI/NFPA 70. This means these products must 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 in 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 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 various ways and are typically mounted in panelboards (also referred to as electrical panels) and PDUs to protect branch circuit wiring. Circuit breakers are also built into equipment to protect components and systems. Interrupting a short circuit—current flow limited only by the wiring resistance—is a severe circuit breaker test. If the interrupting capacity of the breaker is not adequate, the device can overheat and fail.

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 breakers to clear a short-circuit condition but become safely destroyed in the process. UL 489 breakers can interrupt short circuits of 5,000A or more. Typically, UL 1077 breakers can interrupt fault currents of 1,000A.

Overloads can be short-term or long-term. The protective device must not trip with a momentary or short-term overcurrent event that is normal for the protected equipment. Servers, for example, may create inrush currents as internal power supplies 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. The delay curve gives the breaker with a threshold at which an overcurrent becomes damaging.

Summary

- PDUs come in many configurations with respect to the number and type of receptacles, voltage, load capacity, and physical mounting options (horizontal or vertical).
- PDUs are grouped into two main categories: non-intelligent and intelligent.
- Power is distributed to the rack over one or more electrical branch circuits.
- Single-phase wiring is straightforward and consists of two wires where the AC voltage is a single sinusoidal wave measured across the two wires.
- When looking at three-phase 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 continuous current that can flow in a circuit is determined by its wire's size (thickness) or cross-sectional area and terminating receptacle.
- PDUs are designed to fit only the proper PDU plug into their respective feeder circuit outlet. Only the appropriate device plug will fit into its respective PDU receptacle.
- In North America, circuits must be loaded to no more than 80 percent 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 NEC requires an IT system's exposed non-current-carrying metal parts to be grounded.
- Underwriters Laboratories require branch circuit overcurrent protection for PDUs greater than 20A.
- Overloads can be short-term or long-term. A protective device must not trip with a momentary or short-term overcurrent event that is normal for the protected equipment.

3.0 **Elements of the Power Distribution System**

PDUs are the final endpoint of power supplied to IT equipment from incoming building feeds through a chain of equipment, including an Uninterruptible Power Supply (UPS), transformers, larger floor PDUs, and circuit panels. IT and facilities management are increasingly viewing PDUs not merely as a collection of power outlets for IT equipment in a rack but as a network of critical devices that significantly affect the overall efficiency and effectiveness of the data center. PDUs need to be effectively managed like the IT equipment they power. This is driving the trend for deploying more iPDUs in data centers that also power a connected system of environmental sensors (see section 6.1) and integrate with higher-level data center management systems.

This section describes the components of the physical PDU and the PDU management system that uses the intelligence in PDUs for operational improvements and energy use reduction. Further, a PDU's management system can interface with a larger enterprise IT and facilities management system ecosystem.

3.1 Rack PDU

With the adoption of High Performance Compute (HPC) and GPU-based computing, the average power consumption per server continues to increase rapidly. In addition, ongoing deployment of densely packed storage, virtualization, and cloud computing results in data centers with greater watts per sq. ft. /meter requirements, causing more densely packed racks. You must deliver more power to the rack to support new, power-hungry IT equipment. Over the last decade, the typical power required at a rack has increased from 4 to 12 kilowatts and continues upward.

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

Three-phase PDUs capable of supplying multiple circuits, higher voltages, and higher currents are typically deployed to accommodate the increased power demands of racks in data centers. There are also energy efficiency gains in operating devices at 208V or 240V over 120V. In contrast, single-phase PDUs typically only distribute up to 120V of power, considered a lower-density option, which is inadequate to support increased power demands or high-density environments within data centers or mission-critical facilities.

For example, the amount of power available for use is referred to as apparent power. Apparent power is calculated as Volts x Amps and is measured in volt-amps or VA. A 120V 24A circuit has an apparent power of 2,880VA or 2.8kVA. A 208V 24A circuit has an apparent power of 4,992VA or 4.9kVA. Thus, one 208V circuit provides 1.73 times as much power as one 120V circuit, assuming the current (amperage) remains the same. With three 208V circuits, a substantial amount of power can be deployed in one three-phase PDU, allowing you to achieve higher densities and extend the life of your IT equipment.

Circuit Volts	Circuit Amps	Apparent Power	Apparent Power
120V	24A	2,880VA	2.8kVA
208V	24A	4,992VA	4.9kVA
280V 3Ø	24A	8,646VA	8.6kVA
415/240V 3Ø	24A	17,280VA	17.3kVA

3ø indicates three-phase.

Figure 10: Apparent Power for Single and Three-Phase Systems

(Source: Legrand, Inc.)

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 deliver the same amount of power using either single-phase 120V or single-phase 208V. Running a single three-phase power cable to each three-phase PDU reduces the number of cables and the bulk of the cables. Therefore, it makes installation easier and fills less space with wires that can block necessary cooling airflow under raised floors and within racks. Three-phase power is cost-effective and provides room for future growth.

In cases where power needs to be provided at 120V—for devices such as out-of-band modems, routers, hubs, and switches—as well as at 208V for demanding servers, three-phase (Wye) 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 rack allows you to deploy greater power capacity and flexibility efficiently.

3.1.2 Alternating Branch and Alternating Phase Outlet Distribution for Three-Phase PDU Input Power

Three-phase power distribution is commonplace in the data center to handle increased power demands. Traditionally, three-phase power distribution in a PDU was delivered on an alternating branch basis to provide the required kilowatts needed to support increased densities. Today, there are two options: alternating branch or alternating phase outlet distribution.

As mentioned in section 2.3.3, there are two wiring systems for three-phase PDUs: Delta and Wye. In the case of an alternating branch outlet distribution, Delta wired PDUs use a pair of powered lines and a ground to power each branch. Meanwhile, each receptacle in a Wye-wired PDU has its own phase instead of phased pairs.

Dividing the phased power into equal branches is essential. Careful planning of the IT load can help distribute it equally across branches. This is referred to as load balancing. Equalizing the load on all three phases provides the greatest efficiency in power distribution and reduces heat generation in the IT loads of the data center rack. Equalizing the loads in a three-phase Wye system also minimizes the current that exits the rack on the Neutral wire. This translates into power that is charged for but is not used by IT equipment. Without proper load balancing, heat is generated, resulting in higher cooling costs, data center inefficiencies, and an increased chance of tripping either a PDU or upstream breaker and consequently losing power.

Good practice in the data center is installing rack-mounted equipment so that the current draw is similar on each branch. This is easy if the rack is filled with only one type of device, but this is often not the case. Mixed devices such as switches, storage devices, blade servers, and different brands and types of 1U/2U/3U servers can create an unorganized mess of power cables in the back of the rack. The issues above are compounded by the move to higher-density racks that demand more kilowatts of power. Higher power distribution at 60A and 100A forces even more complex PDUs with six, nine, twelve, or more branches, and the power draw still must be evenly distributed across these branches.

A solution for these issues is to use a PDU with alternating phase outlet distribution. Alternating phase outlets alternate the phased power on a per-outlet/per-receptacle basis (A-B-C-A-B-C) instead of a per-branch basis (A-A-B-B-C-C), and outlets have unique coloring (typically black, gray, white) to differentiate their line phase pairing. This allows for shorter cords, quicker installation, and more manageable load balancing for three-phase rackmount PDUs, reducing the risk of tripping a circuit breaker. Shorter cords mean less mass, making them less likely to come unplugged during the transport of the assembled rack. Shorter cords also promote better airflow, reducing the risk of equipment overheating and saving cooling costs in the data center.

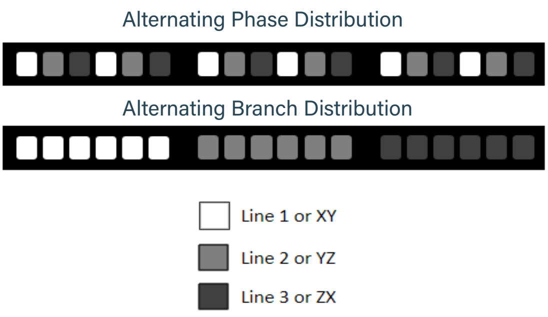

Figure 11: Power Distribution Options

(Source: Legrand, Inc.)

3.1.3 Form Factor

Horizontal PDUs are available in heights of one-rack unit (1U; 1.75 in/44.45 mm), two-rack units (2U; 3.5 in/88.9 mm), or more for mounting in a rack.

By far, Zero U (0U) PDUs that mount vertically, typically to rails or mounting hardware at the back of the rack, are most common. Zero U PDUs do not 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, it allows for the use of shorter device power cords. This results in neater cable arrangements, contributing to better airflow within the rack and improving cooling efficiency. Depending on the IT equipment in the rack, Zero U PDUs can be mounted with screws or hung into the track via toolless mounting buttons.

Higher-power PDUs will commonly be equipped with circuit breakers for branch circuit protection. These circuit breakers may cause the 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.

3.1.4 Outlet Density and Types

PDUs vary in the number of outlets supported based on the physical size (length, width, and depth), the total space available for mounting outlets and internal components, and the power handling capacity of the PDU. For example, a 1U rackmount PDU may have enough room for eight to ten IEC 60320 C13 outlets. A 2U rackmount PDU may have enough room for 20 to 24 IEC C13 outlets. On the other hand, a Zero U PDU may have 36 to 48 IEC C13 outlets or just four 208V/30A NEMA L15-30R outlets.

A typical dense "pizza box" rack deployment (1U server that appears flat like a pizza box) would include two PDUs for redundant power. Each PDU is loaded to 40 percent so that if one power feed fails, the other feed will not exceed the NEC requirement of 80 percent derated current (for North America). Typical outlets for "pizza box" servers are IEC C13 (up to 240V, 12A) or NEMA 5-20R (up to 120V, 16A rated).

In the case of high-power consumption at a rack for a few devices, each of which consumes a lot of power, the total amount of power required might be comparable to the high outlet density example above. Still, the number and type of outlets may be different. Density for devices such as blade servers depends on the number of power supplies (as few as two or a dozen more for redundancy), how the power supplies are configured (a power supply is most efficient when it operates close to its 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 larger quantity of outlets may not be needed, but outlets capable of delivering substantial power may be required.

Typical outlet types for high-demand devices, such as blade servers at 208V or 230V, are:

- IEC C13 (up to 240V, 12A in North America, 10A international).
- IEC C19 (up to 240V, 16A).
- NEMA L6-20R (up to 240V, 16A).
- NEMA L6-30R (up to 240V, 24A).

As the need for computing power increases, one recent innovation from Legrand engineered a C13 and C19 outlet into one flexible high-density outlet—the high-density Cx Outlet (shown in Figure 12) accommodates both C14 and C20 power cables in a single outlet. This new outlet option reduces complexity, increases flexibility, and simplifies the PDU selection process to stay ahead of shifting outlet requirements. The Cx outlet ensures you have access to the right outlets in the right locations.

Figure 12: High Density Cx Outlet
(Source: Legrand, Inc.)

Another challenge is the limited space in a rack for PDUs to be mounted to allow provisions for necessary cables, ventilation, cooling, and convenient access. To help address this, Legrand designed PDUs with High Density Technology (HDOT). HDOT removes unnecessary molding around C13, C19, and Cx outlets to fit more outlets in the same space, resulting in outlets 20 percent smaller than standard IEC outlets. This allows for up to 42 HDOT outlets in a 42U tall PDU.

3.1.5 Cord Length, Feed, and Retention

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

PDU power cords may exit the PDU itself from the device's front, top, or bottom. With the power cable exiting the bottom of a Zero U PDU, you must ensure enough space for the cable's bend radius. In general, a bend radius of 5.25 in/133.35 mm (3U) will be enough, but this should be confirmed as the bend radius 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 needed for heavy-duty cables.

The orientation of the PDU power cord may seem trivial. Still, it can be a potential problem depending on the rack's physical configurations 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, and is there room inside the rack to route the cable so as not to block airflow?

A significant concern in data center power distribution is unintentional power disruption by accidentally disconnecting cords. Proper PDU cord retention practices, like rack cable management, can significantly affect operational efficiency and reliability. Using some cord retention method to support, organize, and secure the many power cords 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 PDUs. Arranging and securing the power cords between the equipment and the PDU is recommended to allow maximum airflow.

Various solutions exist that lock the plug into a PDU's the receptacle to prevent the cord's separation from the receptacle. Some options to secure the plug in the receptacle include:

- Use a plug inserted into a receptacle with a built-in locking mechanism that grips the plug's body to keep it firmly locked in place, such as Legrand's RamLock mechanical locking mechanism (see Figure 13).
- Use a plug with locking tabs inserted into the receptacle, locking them together. One example is using P-Lock locking power cables (see Figure 14).
- Use wire retention clips mounted to the PDU chassis to hold the plug in the receptacle.
- Use cord sleeves that provide a snugger fit for the plug inserted in the outlet. Fins on the side allow the sleeve to be removed from the outlet.

Figure 13: Legrand's RamLock Locking Mechanism

(Source: Legrand, Inc.)

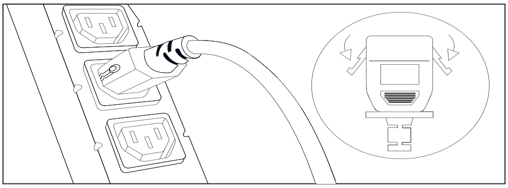

Figure 14: P-Lock Locking Power Cables

(Source: Legrand, Inc.)

In some higher-density outlet configurations, the module design incorporates high native cord retention instead of locking mechanisms.

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

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

- **No connectors:** Cannot be managed externally and may not feature a local display.
- **Local navigation:** Provides a navigable display via local buttons on the PDU itself. You can see basic data collected at the PDU or outlet level.
- **Serial RS232 Connector:** Used for local metering, like an LCD or LED monitor. It can be plugged into a terminal or console server for Telnet or SSH (Secure Shell) remote access; power is typically non-switched. Access via a menu or command-line interface using terminal emulation. Local buttons allow navigation to see unit data. No SNMP (Simple Network Management Protocol) support is available for alarms, except via a specially developed serial console server.
- **Ethernet (RJ45) and RS232 Serial (DB-9M) Connectors:** Used 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 is 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 (see section 6.1). Remote metered models typically have an LCD or LED display with buttons for navigation to see the basic unit and outlet data.

3.1.7 Branch Circuit Protection

Since April 2003, Underwriters Laboratories has required branch circuit protection—circuit breakers or fuses—for PDUs where the inlet current is greater than the outlet current, e.g., 30A (24A rated) plug with 20A (16A rated) outlets. 15A and 20A (12A and 16A rated) PDUs can be supplied without branch circuit breakers because circuit breakers in upstream panelboards can provide the necessary protection. PDUs with breakers or fuses are like mini subpanels, ensuring that the rest of the PDU continues to operate if a branch circuit goes down. For example, a 208V 30A (24A rated) three-phase PDU has three branch 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 differentiating between safe, temporary surges and prolonged overloads.
- **Magnetic circuit breakers** operate via a solenoid (a cylindrical coil of wire that carries current) and trip instantly as soon as the threshold current has been reached. This delay curve is not ideal for servers since it typically has inrush currents anywhere from 30 percent to 200 percent above its 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 standard inrush current and a solenoid actuator for fast response at higher currents. Both thermal and thermal-magnetic circuit breakers are sensitive to ambient temperature (the temperature around a piece of equipment). A magnetic circuit breaker can be combined with a hydraulic delay to tolerate current surges.
- **Hydraulic-magnetic breakers** have a two-step response curve. It supplies a delay on normal overcurrents but trips quickly on short circuits and is not affected by ambient temperature.

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

Fuses are also acceptable for PDU circuit protection. However, replacing a fuse can be time-consuming and, in rare cases, require an electrician, leading to a longer mean time to repair (MTTR). It is recommended to stock spare fuses in inventory and replace the correct fuse to ensure reliability and protection. Fuses' most significant advantage over circuit breakers is the interrupt current rating. For breakers, this is often 5 to 10 kAIC, while fuses typically have an interrupt rating of 100 kAIC.

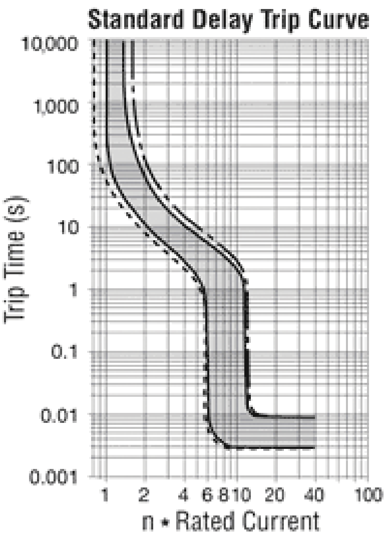

Figure 15: Standard Delay Trip Curve

(Source: Legrand, Inc.)

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

- Compliance with the latest fuse and circuit breaker standards.
- The acceptable MTTR for fuse replacement versus circuit breaker resetting.
- The proper kAIC interrupt rating for the overcurrent protection device.

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

An important consideration is the reliability and flexibility of a PDU's 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 protect two different circuits, i.e., two hot wires. Since one double-pole breaker is less expensive than two single-pole breakers, this will lower the design cost. Double-pole breakers will trip if either of the two circuits they protect is overloaded. It is less expensive than two (or three) single-pole breakers but unless the poles can be used independently in a maintenance shutdown or trip, all two or three circuits are de-energized.

For example, assume a PDU with six branch circuits is protected by circuit breakers. Some 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 it protects is 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 PDUs protect the six circuits with six single-pole circuit breakers—one breaker per circuit, which is more expensive. However, single-pole breakers are more reliable and less limiting. Thus, consider PDUs that allow only one circuit to be de-energized for improved reliability and flexibility.

3.1.9 Circuit Breaker Metering and Relays

Circuit breaker metering is a valuable feature of any PDU. It is imperative when dealing with high-power computing equipment because the consequence of tripping a breaker can be disastrous if it means losing several servers. With circuit breaker metering, you set a threshold. When that threshold is crossed, an alert is shown to reduce power demand. If not, then there is a risk of tripping a circuit breaker. Monitoring branch circuit breakers are essential since a high-power draw means a greater chance of tripping a breaker.

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

When a malfunction occurs in a device connected to a PDU, the malfunction can cause an overcurrent condition in the connected device that trips a circuit breaker. Circuit breaker trip alarming provides an electrical and, in some cases, a sound-based signal when a breaker has tripped due to a faulty condition, bringing immediate visibility to the fault event for faster remediation. iPDUs monitor for overcurrent conditions, and when a circuit breaker trips, it can alert which outlet caused the circuit breaker to trip. iPDUs can be configured to isolate the outlet causing the trip so that the rest of the devices in the rack can be powered up quickly without causing another trip.

Metering occurs at multiple locations on the PDU:

- **Metering at the PDU's Inlet**— Metering at the inlet helps determine power usage and available capacity at the rack, making it easier to provide equipment. By metering at the inlet level, you can avoid overloading the whole circuit and more easily calculate efficiency metrics like Power Usage Effectiveness (PUE).
- **Metering at the PDU's Outlet**— Like monitored/metered PDUs, some metered models help you determine power usage and available capacity at the rack. More importantly, outlet-level metering allows you to understand power consumption at the device or server level to identify ghost servers, find underutilized servers, and allocate costs to specific business units or customers.
- **Metering at the PDU's Circuit Breaker**— Metering at the PDU circuit breaker provides early warning if a circuit is becoming heavily loaded and runs the risk of tripping. Typically, you receive an alert based on a pre-existing threshold informing you when power demands need to be reduced. Branch circuit metering allows you to add new devices to the rack without worrying about tripping the circuit breaker.

Relays allow you to switch power on and off to equipment, improving power flow control and current overload prevention. Switched PDUs have relays included, allowing unused PDU receptacles to be turned off and controlling which outlets are used for devices. Bi-stable latching relays make outlet switching safer while consuming less energy than conventional alternatives. Latching relays can be configured to permanently keep their on/off state so that critical power is maintained even in the case of PDU failure.

The status of the circuit breaker and the continuous monitoring of its current load help the rapid remediation and lower the risk of failures at the rack. Data provided about the conditions can also help plan for equipment maintenance. These options lead to the benefits every data center aims for optimization, lowered costs, extended equipment life, and better failure analysis.

3.1.10 PDU Controller and Local Interface

The PDU's controller—also referred to as a network management module, network card, main control module, or network interface module—is the brain of the PDU. The controller leverages the PDU's higher processing power to coordinate the communication and data transfers occurring from other controllers and connected devices, such as sensors, located on the bus and within the PDU itself. Its interface, connectivity ports, and embedded applications allow for increased interoperability with DCIM software to manage data center operations more efficiently and at a lower cost. A PDU must have intelligence and the ability to connect to a network to provide information to a DCIM or access the PDU's web-based GUI. Linking PDUs allows input from multiple "link" PDUs to be recorded from one main "primary" PDU, saving costs and providing more efficient information. Sensor ports allow more parameters to be measured simultaneously, providing more intelligence to your data center's operations. For example, an individual webcam can be added through a USB port to add more security by taking a photo when a cabinet door is open.

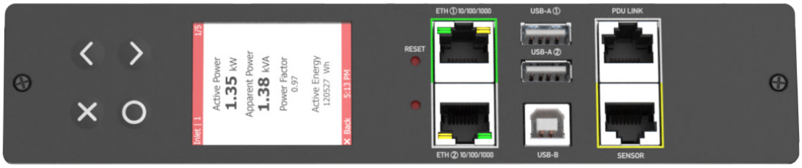

Figure 16: Legrand iX™ PDU Controller
(Source: Legrand, Inc.)

Virtually all iPDUs designed for data center use have built-in displays (typically LCDs or LEDs) to show the current draw or other information like alarms 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 helpful when working at the rack. The local typically allows you to toggle between current draw and voltage. On PDUs that monitor individual outlets, the display can sequence through the outlets to determine the current drawn by each IT device. Some switched and intelligent PDU models may have LED indicators next to each outlet to display status, whether on/off, booting, executing a firmware upgrade, or experiencing a fault.

3.1.11 Remote User Interface

For remotely accessible PDUs (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 PDU over an IP network. The most common is a web-based GUI 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) to allow access to the unit remotely using SNMP or CLI.

One factor to consider is integrating central directory services for user authentication and access control. This becomes especially important when the PDU can 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 LCD/LED and its associated panel navigation buttons.

3.2 System Connectivity

3.2.1 Physical Topology

Like many data center management functions, best practices for the physical topology and, by extension, remote management of PDUs are evolving. The best approach is connecting remotely accessible iPDUs to the management network (separate from the production network). Then collect periodic power meter readings, get immediate notifications of any faults or potential problems, and enable remote power cycling of IT equipment (depending on the PDU's level of intelligence).

3.2.2 Communication Protocols

The communications protocols are typically TCP/IP when PDUs are Ethernet connected, and proprietary protocols for PDUs serially connected to a console server, which connects to the TCP/IP network via Ethernet. Most often, SNMP protocol or open RESTful APIs are used for management. LDAP (Lightweight Directory Access Protocol) and Microsoft Active Directory® protocols are used for authentication, authorization, and user access control. Sometimes with SSL encryption methods, SSH and Telnet may be used for command-line management and HTTP/HTTPS for web-based access.

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

3.2.3 Managing the Rack PDU

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 Distributed Denial of Service (DDoS) or other attacks affecting this critical function. In mission-critical facilities, there are often two connections to each PDU equipped with remote communications: one for syslogging, SNMP traps, access via Web browser, and kWh logging; another for critical functions like remote power cycling, status checks of circuit breakers, and load monitoring (at the unit or circuit level). Some administrative tasks like PDU configuration are performed via command line scripting through a secondary interface such as a serial port. At the same time, Ethernet remains the primary interface for all other functions.

Some essential PDU management functions are listed below:

1. **Audit logging:** Tracks activity like switching 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® BTO, IBM Tivoli®, and others. SNMPv2 is still the most used, but SNMPv3, with its built-in security, is recommended for outlet control applications.
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 is needed for Ethernet-enabled PDUs. A central tool is essential to manage large numbers of PDUs, simplify management, and reduce the cost of ownership.
5. **Alerts:** Normally via SMTP (Simple Mail Transfer Protocol) messages or SNMP traps.

Some or all of the above capabilities are required to effectively manage a data center's rack power. Check your application requirements, service level agreements, and security thresholds to choose the best PDU type for your application.

3.3 Rack PDU Management System

A PDU management system is a software application (sometimes delivered as a software appliance) that consolidates all communications with your PDUs, making them equipped for remote communication. This management system is why intelligence has been engineered into PDUs—it puts the "i" in iPDU. Its main functions are data collection, reporting, power control, element management, and fault management. The system collects and converts raw power data into helpful information and provides a central point for secure access and control across multiple PDUs with operation validation and an audit log. Consequently, it simplifies the management of PDUs and alerts you to potential incidents. We include it in this section because it is a must-have for larger data centers to realize several of the benefits of iPDUs like improved energy efficiency, increased uptime, and lower operational costs.

3.3.1 Data Collection

Data collection is fundamental to any PDU management system and enables all reporting and other rack power 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 iPDUs provide varying levels of individual outlet data and more. It is essential to understand what data you will want or need to analyze when selecting PDUs.

Typical data elements collected, as available from the PDUs, include:

- Total unit active and apparent power.
- Line current, unit, or branch and capacity.
- Outlet-level current and active power.
- Environmental information collected from linked sensors.
- Real-time kWh metering data.
- Door lock status.
- Power quality measurements at the infeed and outlets.

Next, you will want to determine the granularity of the data you need. Your management system should offer a user-configurable data polling interval. Most applications have a 5-minute polling interval, meaning the system will collect data points every five minutes. However, if you require greater granularity, your PDU may need to be able to store data readings so the network is not overloaded with polling traffic. You will want to use a roll-up algorithm to collect data for extended periods without causing the database to balloon and affect performance. For most energy management applications, actual data readings are rolled up to a minimum, maximum, and average, and data is collected hourly, daily, and monthly.

Advanced polling options enable a customer to minimize network traffic while allowing granular data collection. Advanced polling requires a PDU with the memory capacity to record and store readings called samples. The PDU management system should offer the ability to configure optional sample rates for each PDU and set optional polling intervals for the management system itself to collect the stored samples at each PDU not previously collected. For example, a PDU can be configured to record and store samples every minute. The PDU management system can be configured to poll the PDU once an hour. Each poll will pull the 60 one-minute samples from the last survey, with the intelligence to know the final reading recorded on the previous poll. This process is key to preserving data integrity in the case of a network failure or if some other problem occurs and ensures that you can still access the required data.

3.3.2 Reporting and Analytics for Power Monitoring and Measurement

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

Reports on maximums and minimums for current, ambient temperature/humidity, and active power simplify key tasks to ensure that you are not in danger of exceeding circuit breaker ratings, overcooling, or undercooling. For instance, environmental information can give you the confidence to raise temperature setpoints 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 your data center 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. Still, you will need to understand the resulting power demand of the host servers. You can also establish objectives and usage reports and implement changes for both physical components of the data center—floor, room, row, rack, and IT device; logical groupings like by customer, department, application, organization, and device type as examples. This level of detail creates visibility and accountability for energy usage. Some organizations even use this collected data to issue energy bill-back reports to users/owners of the IT equipment.

3.4 Advanced Power Metrics

Rising power demands and the increased capabilities of iPDUs, combined with the global initiative of reducing carbon footprints, make measuring all aspects of rack power critical in providing high-quality power reliability. Beyond the actual power used by the rack, these details can be used to troubleshoot sources of power issues such as leaks, distortions, and variations to increase energy efficiencies, improve availability, and manage the existing capacity of the entire data center.

Two crucial factors in providing high-quality power are determining if you have power quality issues and what type of power quality problem exists. In a data center, different power quality problems can arise throughout the power chain and are described as glitches, spikes, disturbances, flickers, blinks, etc. However, while they all show that "something" occurred, the IEEE has defined them as:

- **Transients:** Spikes, other short-term events that raise the voltage and/or current.
- **Interruptions:** Complete loss of supply voltage categorized as instantaneous, momentary, temporary, or sustained.
- **Sags:** Reduced AC voltage duration of half-cycles to 1 second.
- **Undervoltage:** Reduced AC voltage of several seconds or longer (i.e., "brownout" of utility power).
- **Swell/Overvoltage:** Inverse of a sag, increased AC voltage.
- **Waveform Distortion:** Any imperfections of a sine wave: harmonic distortion, voltage or current clipping, or spiking.
- **Voltage Fluctuations:** Unstable random or repeated voltage changes: short or long term.
- **Frequency Variations:** Any change from the nominal frequency.

Once it is determined that a power quality problem has occurred, a proper plan can be derived to mitigate it. Using an iPDU that includes built-in advanced power quality metrics at the rack, both at the PDU inlet or what is often called the PDU infeed, and at the PDU's outlets is a start. Without an iPDU that already has the proper built-in meters, you would have to install external meters throughout your power chain to provide continuous power quality monitoring to detect power quality problems. These power quality meters should include capturing and viewing waveforms, detecting disturbances like voltage sags and voltage swells, measuring harmonic power flow, and providing alarms when measurements are outside a set tolerance range. They will need connectivity to your management systems to track and monitor events. If you choose an iPDU with embedded power quality measurement capabilities, your audit tracking and performance analytics will be optimized.

Also, note that power quality can mean different things to different operations. For example, mission-critical and data center facilities housing sensitive electronics and IT equipment have small tolerances for deviations in power. These facilities are called upon to operate 24/7 and rely on high-quality power to keep the systems running normally. IT equipment is more sensitive than ever toward minor and transient variations in voltage. Furthermore, every minute of downtime carries a greater penalty than ever before.

In short, the downtime costs have increased dramatically and so has the likelihood of failures caused by power quality problems. The following sections review some advanced power metrics that help monitor power at the rack level.

3.4.1 Neutral Monitoring

Neutral monitoring is essential in three-phase Wye wired PDUs and monitors both neutral current and neutral voltage. Neutral monitoring often happens at the PDU inlet.

In Wye-wired three-phase PDUs, all three phase lines return their current flow on a common shared neutral. Because the three phase lines are 120 degrees out of phase with each other, the sum of the three phase line currents returned in the neutral tends to cancel out. It can be shown that under "normal" situations, the neutral current never exceeds the current in any one of the phase lines. For this reason, the neutral wire size can be, and is, the same size as the three phase conductor wires. However, when there is excessive harmonic distortion (see section 3.4.3) on the power waveform, the phase line currents do not cancel in abnormal situations. They add up, which can lead to wire overheating. iPDUs can monitor neutral current and alert when an overcurrent condition on the neutral occurs.

In Wye-wired three-phase circuits, the common neutral wire is attached to the ground within the data center infrastructure, which means the neutral voltage at the rack is never more than a few volts. However, if a break in the neutral wire occurs, the neutral voltage will increase dramatically. Because the neutral line has no protection on the line itself and the neutral voltage is not monitored by a circuit breaker, there will not be a trip when this occurs. Therefore, a neutral overload can cause the voltages to be distributed to rack server loads that exceed device ratings and can cause equipment failures. Some iPDUs monitor neutral voltage and can alert when a neutral overvoltage condition occurs.

3.4.2 Residual Current Monitoring

Typically, all current flowing in a circuit is contained in the phase and neutral lines. When a device malfunctions, an abnormal condition can occur where current leaks out of the circuit and flows in the chassis and ground wires. This leakage, called residual current, is a safety hazard that can cause electrical shock and fires. Some iPDUs can be purchased with embedded residual current monitors that can detect excessive leakage. There are typically three Residual Current Monitoring (RCM) options to choose from to equip your PDUs: RCM Type A, RCM Type B Single Channel, and RCM Type B Three Channel.

RCM solutions are usually used at the branch circuit level, at the outlet, or inside the panel protecting the branches. RCM solutions integrated into a PDU measure residual current at the rack level, per phase, which allows for reporting of leaks at an early stage and with accurate location information. When embedded in a PDU, it will help you maintain uptime by providing residual current monitoring data into your management system. By adding RCM measurements to other power quality readings, you can identify which server power supply is at risk and could eventually fail. The PDU can then be easily configured to send the appropriate alerts to your monitoring system to act preemptively.

Key benefits of a PDU's RCM measurement technology:

- Higher operational reliability through early detection of possible critical facility conditions.
- Significant reduction of operating costs due to optimized maintenance.
- Easy monitoring of available energy distribution facilities and early warning of system errors.
- Localization of single faulty outlets, less effort for troubleshooting.
- Critical residual currents are detected at an early stage, thus increasing fire safety.
- Fulfillment of the security criterion "RCM fault current monitoring" in the data center.

3.4.3 Total Harmonic Distortion

Total Harmonic Distortion measures how much your electrical load distorts your utility's power. Harmonics are always present in current and voltage. Still, too much variation can accumulate and disrupt power quality, leading to increased electrical usage, power quality fluctuations, or over-heated wires that can cause damage to IT equipment.

Today's iPDUs have monitoring devices built into them to conduct continuous harmonics monitoring and alerting by measuring harmonics, voltage dips and swells, crest factor, and more on the power being fed directly to the PDU or at the PDU's outlets or load. These rack-level measurements elevate your rack-level power quality monitoring and prevent problems before they occur, keeping your data center and equipment optimized for years to come.

3.4.4 Waveform Capture

Waveform capture is commonly known as the process where an oscilloscope captures voltage and current waveforms after a threshold is crossed or some unexpected power event occurs. In data centers, waveform capture can help identify power quality problems. For example, when a power supply blows, you can capture the current and voltage values via a waveform capture in a concise period to gather the cause of the issue, most often its current surge. Some iPDUs have advanced power monitoring systems that can measure down to the waveform level. Additionally, some monitoring systems use waveform capture to send alerts of unusual voltage readings.

3.5 Security and Access

Security has always been a concern in the data center. Still, with increasing high-scale security breaches and building intrusions, data centers are doing more than ever to maximize physical and cyber security to prevent breaches and hacks. The building and the equipment in the building need to be secure from events like natural disasters, robberies, and attacks, to name a few. Compounding multiple security measures offers higher security, as various barriers prevent hackers from gaining access to sensitive information.

Since the PDU is typically housed within the rack, this handbook will focus solely on rack-level security measures. Below, some common forms of data center security measures are outlined.

3.5.1 Physical Rack-Level Security

Security inside the data center floor protects the equipment storing data, including PDUs. Since the server rack is the last point of data vulnerability in the data center, where your data infrastructure is powered, it makes sense to consider implementing the same level of sophisticated physical security and access control monitoring at the rack that is already established at every other level of entry in the data center. While physical damage inside a data center is uncommon but possible, it is essential to put measures in place in case a security event happens.

To fulfill your rack-level compliance requirements with the utmost confidence and efficiency, you must make intelligent decisions for the near and long term. Implementing control and audit solutions that readily fit into your existing building and DCIM technology that uses your current security systems is a start. Over time, you can add new and required rack-level control and monitoring capabilities cost-effectively without adding counter-productive friction to your staff's daily tasks.

Escalating regulatory requirements across industries now requires sensitive systems and data to be subject to specific protections. Ensuring that only authorized staff enters the data center is no longer enough. You must track and monitor each person's access to specific sensitive systems and ensure they are properly authorized for a particular area. And you must be able to provide an extensive audit trail on who touched those systems when—and what they did each time.

Some measurements to secure data center hardware from a physical attack include:

- A single-entry point into the data center floor.
- Security guards at the point of entry into the server room monitor who enters and exits the room.
- Security cameras with video retention at the aisles, racks, and front and back of racks to monitor any equipment tampering or other causes of equipment not working.
- Fire alarms, smoke detectors, and sprinklers will protect the hardware in case of a fire. Similarly, sensors monitoring water leakage and temperature can help prevent floods or any cause for overheating/overcooling.
- Real-time alerting/alarming that notifies appropriate parties of problematic events requiring immediate attention.
- Door locks with a key or card authorization allow authorized users to access different racks (see Figure 17).
- Ability to lock outlets, individually or in groups, preventing accidental disconnection of equipment.
- Integration with DCIM software and/or other access control systems facilitates a single point-of-control and easy consolidation of security/compliance-related audit trails.
- PDU integration with the rack ensures continuity of security and compliance even in a power outage.
- Covered wires and rodent repellent to prevent unwanted living creatures from damaging wires which may cause infrastructure damage.

Figure 17: Legrand's SmartLock Enclosure Door Handle

(Source: Legrand, Inc.)

Physical security is the first step to keeping any data center safe. It is worth noting that rack-level compliance requirements will continue evolving as your data center and regulators become increasingly concerned about the potential social and economic impact of security events. It is wise to take a long-term view of your rack-level physical security needs—rather than focusing only on what current regulations require. The primary goal of compliance is similar: Ensure that your most sensitive systems are protected against inappropriate access and that your compliance with regulatory mandates is accurately documented.

3.5.2 Password Protection

The standard form of cyber security for data is strong passwords. Passwords are ubiquitous in the online world and remain one of the critical components of security. Some measures to ensure secure passwords include:

- Passwords with 8 or more characters including at least 1 of each type of character: lower case, upper case, numeral, and special characters.
- Changing passwords frequently, including the first time after using a default password to log in. That way, hackers cannot use a system default password to gain access to sensitive information. Changing passwords frequently also ensures that if a hacker gains access to past data, they cannot use a past password to access PDUs.
- Using unique passwords ensures if other sectors of the company face a security breach, hackers cannot guess their way into information from every sector within the company.

Password management may seem tedious, and more of a nuisance at this point, but passwords are still a crucial layer to keep your information and data being collected by the PDU safe from a cyberattack.

3.5.3 Firewalls

Another standard form of online security used in PDUs is firewalling. Since iPDUs are accessed for several reasons, ranging from data collection to power control, it is critical to identify which users are authorized and which are unauthorized to access this information. One type of firewall used is a Role Based Access Control (RBAC) firewall, which allows access based on the roles of individual users. Another type of firewall is an IP Access Control List (IP ACL), which provides access based on the IP address of the host sending or receiving the request.

3.5.4 Encryption

As PDUs are connected to management networks and sometimes even production networks, all data sent or received by the PDUs must be encrypted. Some examples of encryptions include HTTPS (with SSL or TLS), SNMP, and SSH.

HTTPS connections use TLS 1.2/1.3 with AES 128/256-bit ciphers supporting a wide range of browsers. SSH connections use Public Key authentication where password authentication is not adequate or feasible, like in scripts. SNMPv3 connections are encrypted with MD5 or SHA authentication protocols and DES or AES privacy protocols.

3.5.5 Other Security Measures

In addition to the security measures listed above, there are added measures that can be taken to improve PDU security. This list is non-exhaustive but will give some thought to how to keep PDUs secure from breaches.

- Automatically logging out after a period of inactivity to prevent unauthorized access.
- Limiting the number of clients that can use the same login credentials.
- Blocking access after repeated failed login attempts to defend against potential DDoS attacks and logging the source of the attempts.
- Enforcing Restricted Service Agreement warnings and requiring users to accept them to log in.
- Encryption of session ID cookies to identify each unique user.
- Annual infiltration testing by a third party.

Keeping your PDUs updated on security measures is critical to protecting your information and preventing a significant breach. The more data there is, the more complex and layered security methods should work together to support a secure network.

Summary

- PDUs are the final endpoint of power supplied to IT equipment from incoming building feeds. As a result, IT data centers increasingly view iPDUs as integral to managing their network of critical devices.
- Data centers are deploying PDUs that can supply multiple circuits, higher voltages, and higher currents, and energy efficiency gains are seen in operating devices at 208V or 240V over 120V.
- Power distribution of the PDUs on a per-branch or per-outlet basis changes the power distribution and must be considered for load-balancing purposes.
- PDUs are available in various rack form factors, leading to neater cable arrangements, better airflow, and more efficient device powering.
- Basic PDUs are falling out of favor as iPDUs can collect power quality and environmental information and warn of potential problems.
- High-density Cx outlets accommodate both C14 and C20 power cables in a single outlet.
- PDU power cords may exit the PDU from the front, the top, or the bottom of the device.
- Circuit breakers in PDUs are typically thermal-magnetic or hydraulic-magnetic with delay curves that allow for reasonable inrush currents while protecting devices from excessive fault currents.
- Circuit breaker metering is essential, mainly when dealing with high power, because tripping a breaker can be disastrous if it means losing several servers.
- Virtually all PDUs for data center use have built-in displays (typically LCDs or LEDs) to show the current draw for the entire PDU unit.
- To remotely access a PDU, there are typically two choices: a graphical user interface (GUI) or a Command Line Interface (CLI).
- Integration with central directory services is essential when the PDU can remotely turn on/off/recycle individual outlets or groups of outlets.
- The management system for data center power is often run on a management network separate from the production network.
- A PDU management system is a software application that consolidates all communications with your PDUs. Its main functions are data collection, reporting, power control, element management, and fault management.
- Advanced power metrics allow the management of specific parts of the PDU to troubleshoot issues quickly and improve uptime.

4.0 Considerations for Selecting Rack PDUs

Several approaches to deploying power to racks affect PDU selection and configuration. Some approaches provide degrees of redundancy and higher reliability/availability than others but may not be appropriate for certain types of IT equipment. Redundancy and higher availability require power resources, so if you have limited power resources, you must decide what IT equipment justifies redundant power and what equipment does not.

4.1 Power Available and Distributed to Racks

4.1.1 Single Feed to Single PDU — Low Availability Power Deployment

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

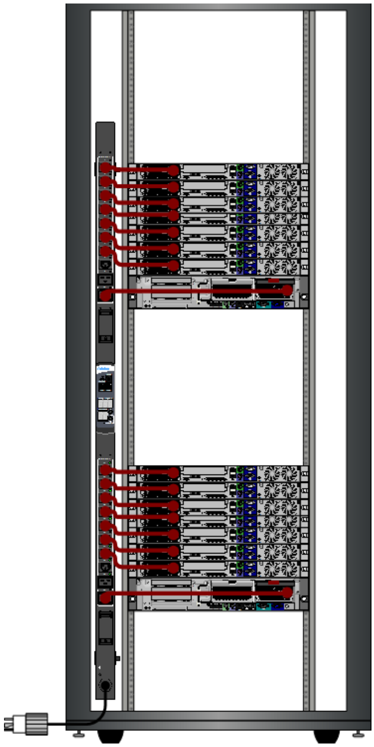

Figure 18: Single Feed to Single Rack PDU

(Source: Legrand, Inc.)

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

4.1.2 Dual Feed to Single PDU with Rack Transfer Switch—Slightly Better Availability Power Deployment

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

There are two main types of transfer switches: static transfer switch (STS) and automatic transfer switch (ATS). An 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 electromechanical relay technology, which results in slower transfer times.

This arrangement still loads the PDU to 80 percent of the maximum. But the electrical power capacity required has doubled—one feed is operational, and the second feed is a backup. Also doubled is the amount of upstream equipment necessary to supply the added feed.

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

The third type of rack transfer switch to consider, patented by Legrand, is a hybrid rack transfer switch. Hybrid rack transfer switches prevent downtime in racks with single power supply devices. It offers electromechanical relays and silicon-controlled rectifiers to overcome the limitations of a traditional ATS. It provides load transfer times comparable to an STS, is more energy-efficient, and is available at a far more accessible price point. With inlet, outlet, and branch circuit-level power metering and outlet-level switching, you can realize better remote power control, blazing-fast transfer times, uncover hidden capacity, and ensure power resources are used efficiently.

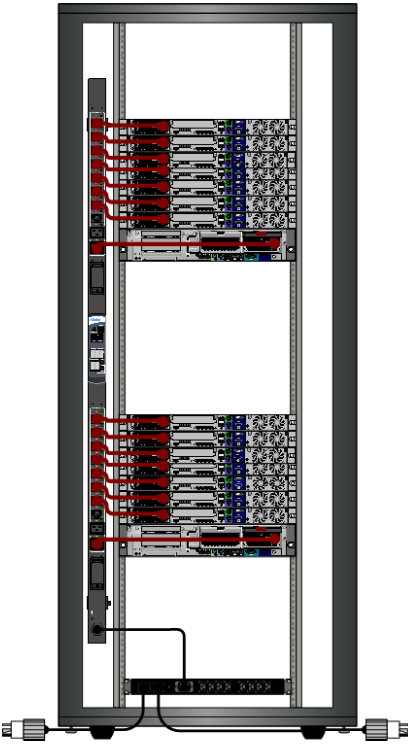

Figure 19: Dual Feed to Single Rack PDU with Rack Transfer Switch

(Source: Legrand, Inc.)

4.1.3 Dual Feed to Dual Rack PDUs—Improved Availability Power Deployment

Many data center components like servers, network devices, storage systems, 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 six power supplies. The most reliable power deployment is to use two power feeds to two PDUs. With this configuration, if one PDU or power feed fails, a second one is available to support power to the IT equipment in the rack. A common practice using dual feeds is using PDUs with color-coded chassis, such as red and blue. A PDU's color-coded chassis enables visual control for the installation of or changes to its 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 color-coded chassis, typically provided by vendors in a palate of colors and sometimes by using color labels placed on the PDU chassis, helps lessen the confusion about which PDU is fed by circuits A or B and lowers the risk of unplanned downtime.

It is important to remember the requirement that each circuit be loaded to no more than 40 percent. If the two circuits feeding the rack are both loaded to 80 percent, 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 percent to 160 percent, and the circuit breaker for that feed would trip, causing the second circuit to the rack to also lose power. To prevent this event, both feeds should be loaded to no more than 40 percent so that if one fails, the remaining circuit will not be loaded to more than 80 percent. Compared to the case with ATS (see section 4.1.2), where one feed is the backup, both feed power IT equipment in this configuration.

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

4.1.4 Multiple Power Supplies—High Availability Power Deployment Depending on the Configuration

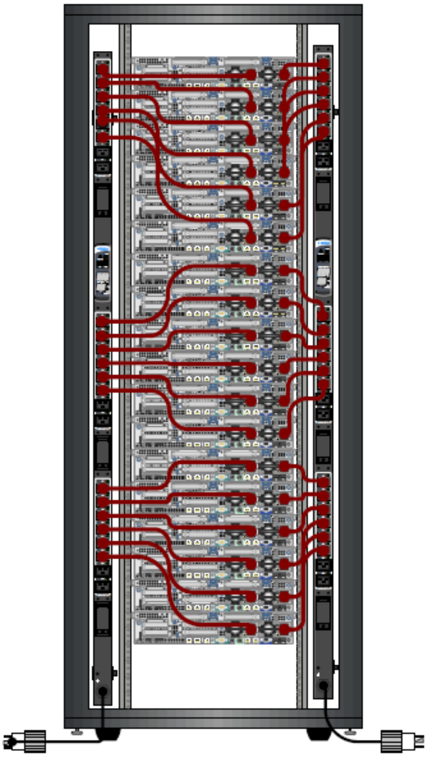

Figure 20: Dual Feed to Dual Rack PDUs

(Source: Legrand, Inc.)

IT equipment with two or more power supplies can vary in how power is delivered to the equipment. Some devices have a primary and backup power supply, some alternate between the power supplies, and some 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 if any primary power supply fails. 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's specifications to understand how the power supplies work to optimize balanced load configurations on the PDU, especially those with branch circuits and three-phase models.

4.1.5 Load Balancing

Load balancing is a procedure that tries to evenly distribute the rack equipment's current draw among the PDU's branch circuits; as you come closer to a perfect balance, more total current can be supplied with the greatest headroom in each branch circuit.

In three-phase PDUs, for example, as the load comes closer to perfect balance, the current draw is more evenly distributed among the three phase lines (more upstream headroom), and the 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 is 18A, 18A, and 0A, respectively, and the total current is 36A. The PDU has 7.6A (18.0A – 10.4A) more headroom when the load is balanced across all three lines.

Load balancing can be tricky because many IT devices draw varying amounts of power based on the computational load. To load circuits evenly 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. This is true both within a rack and across multiple racks. Devices with dual power supplies should be plugged into different PDU circuits on a rotating basis. A typical deployment would be the scenario of the dual feed to dual 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 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 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 you want to try to balance the load across all circuits and lines and cannot be sure that each of the four-blade servers will be performing tasks that equally load the circuits, you should stagger the second blade server power supplies. Therefore, 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 into 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 extremely helpful for (re)balancing loads in the rack.

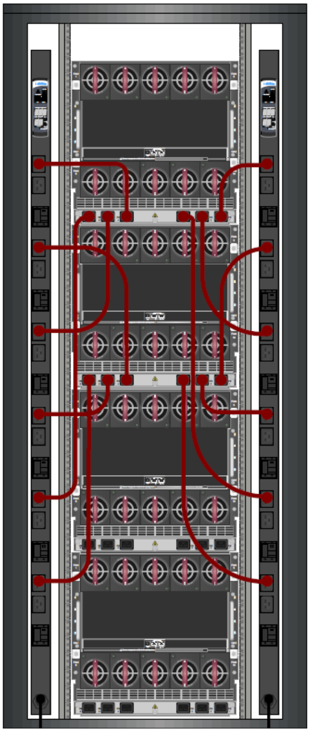

Figure 21: Multiple Power Supply Configuration

(Source: Legrand, Inc.)

4.1.6 Inrush Current

Server power supplies draw more current when first turned on. This is known as inrush current. As discussed in section 2.6, PDUs with circuit breakers are designed not to trip during noticeably short periods of high currents. However, it is easier on upstream circuits if the sudden surge when equipment is turned on is minimized. For this reason, some iPDUs provide outlet sequencing and allow you to configure the sequence and the delay time in which the outlets are turned on. Some iPDUs even will enable the programming of outlet groups and the sequencing of outlet groups.

4.2 Power Requirements of Equipment at Rack

Section 4.1 describes ways to deploy electrical power to a rack. This section helps determine 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), specifying a voltage and current (amps) higher than usually seen during actual deployment. As a result, there is a convention of using a percentage of this value, let's say 70 percent when computing the maximum PDU load capacity required: PDU load capacity = ∑ (device nameplate (V x A)) x 70%) for all devices in the rack. For example, 208V x 2.4A x 70% x 14 servers = 4.9kVA.

For the example above, if you supply 208V power to the rack, you need a 30A (5kVA) PDU since you will load it to 80 percent 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 PDU and load both PDUs up to 40 percent. You will need to specify the appropriate number of outlets. It is a good idea to have spare outlets for other devices, even if the PDU is 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 added power capacity to power more equipment. Current best practices are standardizing a PDU with IEC C13 and/or C19 outlets at 208V. Most servers and data center devices can run at 208V (even up to 240V). Many are also considering Legrand's Cx Outlet, which combines a C13 and C19 into one flexible high-density outlet that accommodates a C14 or a C20 plug in a single outlet, standardizing a single PDU for multiple applications (see section 3.1.4).

The nameplate power derating factor of 70 percent mentioned above was just an estimate. Often, the power drawn from the server is 20-30 percent below the nameplate rating. Therefore, you should measure the actual power consumption to ensure power is accurately deployed to prevent unused power.

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

It is also important to note that three-phase Wye 208V PDUs may support both 120V and 208V in the same PDU, depending on the model. This is handy when IT various 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, the total power consumption is 42 x 200W = 8.4kW. To allow for the NEC requirement of 80 percent limitation, the rack needs 10.5kVA (8.4kW / 0.8). Two PDUs that provide 10.5kVA each are required to allow redundant power feeds. 208V single-phase at 60A (48A rated) can deliver 10.0kVA. This could suffice, particularly if the 200W per server estimate is high. Another alternative is 208V three-phase at 50A (40A rated), which can deliver 14.4kVA. The 208V three-phase alternative provides headroom to add higher power demand servers in the future and can handle the existing servers, even if its average power consumption increases from 200W to 220W.

Three-phase power enables one whip or PDU to deliver three circuits instead of just one circuit. The whip or input power cord on the PDU will be larger (in thickness) 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 (as shown in Figure 3). 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. When graphed, the power level 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 you need to include a factor of $\sqrt{3}$ (1.732). The apparent power formula for three-phase is: V x Derated A x 1.732 = VA. As an example, three-phase 208V with 50A (40A derated) is 208V x 40A x 1.732 = 14.4kVA. A three-phase Delta deployment provides three separate circuits and over 70 percent 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 or by combining a line with the neutral. As an example, a three-phase 208V Wye wired 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 or utility 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 line (L1, L2, L3) to the line and line to neutral can provide power to both 208V and 120V outlets. Whether the three-phase wiring is Delta or Wye, the voltage is always referenced line-to-line, not line-to-neutral. This is even true in the 400V or 415V example in 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 wired PDUs that support 120V/208V or Delta-wired PDUs that support only 208V changing the data center wiring. For example, a Delta wired 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, as it allows the 5-wire branch circuit to be used for a 120V receptacle without any safety issues. For example, a data center could deploy Delta wired PDUs to racks where there is only a need for 208V and Wye wired PDUs to racks where there is a need for both 120V and 208V.

Three-phase cables may be slightly larger than single-phase cables. However, it is essential 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 400V and 415V Three-Phase Rack PDUs

The more common method of delivering substantial power to densely packed racks is via 415V (typically North America) and 400V (typically Europe, Asia/Pacific, Australia) three-phase Wye wired PDUs. 400/415V power distribution from panels to racks is now an accepted practice without stepping down the voltage. A data center designer could specify 400/415V Wye whips to 400/415V Wye PDUs. Since most data center equipment can safely operate from 100V to 240V, the 400/415V Wye PDU can provide three circuits—L1+N, L2+N, L3+N—each supplying 230V (400V / 1.732) or 240V (415V / 1.732).

4.3 Rack PDU Selection

4.3.1 Basic Steps to Selecting a PDU

As you can imagine, various PDU configurations are available based on the number of phases, voltage, total amps, branch circuits, number and type of outlets, input plug type, rack units consumed, and physical dimensions. Here, we will discuss some essential topics you should follow to help select and deploy appropriately sized PDUs in a data center.

- Find out what devices will be in the rack.
- Establish the rack kilowatt budget, adding room for future growth.
- Confirm the facility power system that will be used.
- Determine the circuits, phase, and amperage for the rack.
- Ensure the loads are balanced, not just across phases but across branches.
- Learn the PDU installation options, noting that sometimes this depends on the location of the whip and the type of rack being used.
- Figure out if advanced /intelligent PDU features are needed.

4.3.2 Rack PDU Selection and Special Application Requirements

Other factors in selecting a PDU include data center location, application requirements, IT equipment requirements, available power, space in the rack, energy management, and efficiency objectives. All these factors combined will decide what type of PDU should be used. Some of the considerations below will guide you in selecting the feature set and type of PDU you will need to satisfy your requirements.

Understanding the type of equipment and how many devices are going into the racks, 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, will help define the physical configuration, i.e., the number and type of outlets, and the power capacity of your PDU(s). Average rack power requirements have risen from 10.5kW in 2016 to 11.7kW in 2019 and 12.02kW in 2021; it is not unusual to see racks wired to provide 17kVA or more than 57kVA. It is recommended to consider your needs for today and also to make considerations for future expansion.

Decision criteria for 24/7 staffed sites will differ from those of remote management of lights-out facilities. If you need remote or lights-out management of a facility, you will need a switched PDU, requiring more security and user access management. Remote applications may also call for SNMP management features.

Integration with directory services, like LDAP or Microsoft Active Directory, is increasingly required to control access to resources in place of a separate access control system. This capability applies to all applications requiring local or remote central authentication. For many data center applications, encryption and configurable password support (see section 3.5.2) are necessary for remote access.

The PDU must supply uninterrupted power to each device plugged into it. You will want to prevent or mitigate any events that can cause circuit breakers on the PDU or any upstream point to trip. Outlet sequencing is a valuable feature to prevent inrush current (see section 4.1.6) from tripping a circuit breaker by establishing a sequence and appropriate delay for powering multiple devices. Outlet sequencing prevents the undesired tripping of a circuit breaker and lets you 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 combined with the outlet grouping capability.

For some applications and equipment, you may need a customizable alarm threshold for each outlet that can switch off an outlet should it exceed a specific power draw. This would prevent a temperature or other sensor (see section 6.1) from causing a broader shutdown that impacts many servers. An advanced application is the control of the HVAC system using the temperature reported by a PDU's connected temperature sensors.

In critical environments, devices often have multiple power supplies fed from different feeds or circuits for failover and redundancy. In this case, the managed device power supplies must be controlled as "one single" device, and all outlets must be handled simultaneously. This capability is available through the outlet grouping features of PDUs 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, suppose a device in a remote location does not respond and the WAN is not operational. In that case, there are two options: an expensive time-wasting truck roll to restart, or a PDU with the intelligence to trigger a restart of a malfunctioning device (cycling power to the device after it has not responded for 20 minutes).

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

Adding color to a PDU may seem like little more than a vendor gimmick, but the reality is that there are several advantages to using color in your data center. You can use full-color PDU chassis or vendor-supplied color labels/stickers placed around the PDU's plate label, controller/network card, and/or power cord to keep track of the complex system of pipes and wires to equipment that is otherwise difficult to manage. When you have dozens, if not hundreds, of identical PDUs and cables, it becomes difficult to see which feed supplies power to your devices. By using a PDU with color options, you can easily find the feed that supplies power to your rack and troubleshoot any problems you are experiencing more quickly, thereby lessening the likelihood of downtime.

In general, PDU color options help to:

- Easily find redundant power feeds to IT equipment.
- Make it easy for technicians working in an IT rack.
- Clearly show the power chain.
- Identify different voltages.
- Reduce lighting requirements.

4.3.3 Benefits of an Intelligent Rack PDU

As IT buyers struggle with smaller budgets, a lower price point is often the deciding factor when selecting a PDU. Since basic PDUs retail for less, it may appear to be a "better buy". Others may not believe that the advanced features of iPDUs provide additional value or that their organizations have the time or resources to benefit from the features and value of an iPDU. As a result, many buyers choose basic PDUs even though iPDUs provide greater value and cost savings in the long term.

Over the past decade, data centers of all sizes have moved towards more sophisticated, manageable systems. This fact is nowhere more evident than the trend toward using iPDUs.

An intelligent PDU will provide, at minimum, PDU-level power monitoring, rack temperature and humidity monitoring, and may also offer outlet-level power monitoring and/or remote outlet switching. For top-tier data centers, the deployment of intelligent PDUs can make a significant difference in the ability to improve uptime and staff productivity, efficiently use power resources, make informed capacity planning decisions, and save money. In doing so, it will operate greener data centers. Suppose your data center has dozens of racks. In that case, the greatest benefits will be realized by using a PDU management system to consolidate data acquisition, reporting, PDU administration, and control, all supported by the intelligence in an iPDU.

Let us examine the benefits of deploying iPDUs in more detail.

iPDUs Improve Uptime and Staff Productivity

- Monitoring power at a PDU and its outlets, with user-defined thresholds and alerts via email or SNMP, provide awareness of potential issues before they occur. If a problem occurs, the PDU can provide alerts so that it may be remedied sooner.
- Remote reboot of servers and IT equipment, based on the outlet it is plugged into, from anywhere in the world via a web browser reduces downtime and personnel costs.

iPDUs Use Power Resources Safely

- User-configurable outlet-level delays for power sequencing prevent circuits from tripping due to IT equipment inrush currents.
- Control of outlet provisioning prevents accidentally plugging IT equipment into already heavily loaded circuits and is at risk of tripping circuit breakers.

iPDUs Enable Informed Power Capacity Planning Decisions

- Outlet-level monitoring may offer opportunities for 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. If critical cooling is lost, an alert can be sent, or if the equipment is in colocation, this monitoring can confirm if service level agreements are being met.
- Many data centers grow in a chaotic and unplanned way. Technicians often plug new equipment into the first available outlet with little understanding of how much electrical capacity is available. This can result in blown fuses and cause downtime.

iPDUs Save Power and Money

- Monitoring power at the outlet level combined with trend analysis can find ghost or underutilized servers that are candidates for virtualization, containerization, or decommissioning.
- Remote power cycling enables you to quickly reboot hung or crashed IT equipment without incurring the cost of site visits, reducing MTTR.
- Temperature and humidity sensors help to optimize air conditioning and humidity settings, avoiding the customary practice of overcooling and related waste of energy (see section 6.1).

4.4 Power Efficiency

4.4.1 Introduction to Power Usage Effectiveness (PUE)

It is essential to understand the role of the PDU in providing accurate data that you will need to improve efficiency. We all know The Green Grid's phrase, "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. To manage power efficiently, IT equipment aims to maximize the amount of valuable computational work for the total energy consumed by the equipment and the infrastructure that supports the equipment. The most used metric of this is The Green Grid's Power Usage Effectiveness or PUE: PUE = Total Facility Power ÷ IT Equipment Power.

Total Facility Power represents all the power required to operate the entire data center. Total Facility Power is made up of IT equipment items: servers, storage, network equipment, etc.; and support infrastructure items: CRAC units (Computer Room Air Conditioning), fans, condensers, UPS, and lighting. IT Equipment Power is 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 IT equipment consumes 50 percent of data center power. Another commonly used metric is The Green Grid's Data Center Infrastructure Efficiency (DCiE), which is the inverse of PUE. Since it is derived from the same data, there is no substantive difference in the measurement or its usage. The average PUE for the data center has decreased in recent years from increased measures in data center efficiency.

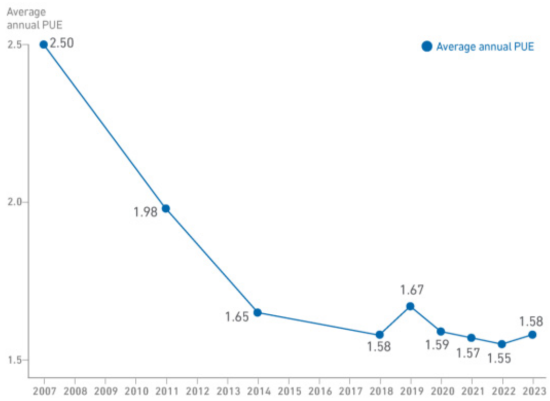

Figure 22: Average Annual Data Center PUE

(Source: Uptime Institute)

4.4.2 PUE Levels

The Green Grid defines three levels of PUE: Level 1, Level 2, and Level 3 (see Figure 23). Many industry analysts recommend measuring IT power consumption at Level 2, i.e., the PDU level. While it is true that PDU-level power consumption will provide the denominator needed to calculate PUE, this information alone is 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 time for typical power usage to ensure the peaks and valleys have been captured to establish a baseline and to track your improvements.

	PUE Level 1 Basic	PUE Level 2 Intermediate	PUE Level 3 Advanced
IT equipment power measurement from	UPS outputs	PDU outputs	IT Equipment inputs
Total facility power measurement from	Utility inputs	Utility inputs, UPS inputs/outputs, Mechanical inputs	Utility inputs, PDU outputs, UPS inputs/outputs, Mechanical inputs
Measurement interval	Monthly/weekly	Daily/Hourly	Continuous

Figure 23: Definition of The Green Grid PUE Levels

(Source: The Green Grid)

4.4.3 Why Advanced, Level 3 PUE?

A lower PUE may be misleading since it can result from inefficiencies in power consumed by IT equipment, which merely increases the denominator. A lower PUE is better than a higher one, but it is possible to implement measures that reduce data center energy consumption but increase your PUE. For example, if you replace older, less efficient servers with more efficient ones, eliminate ghost servers, turn off idle servers during the night, or employ server virtualization, the net result would be power reduction with an increased PUE. The detailed IT load data from Level 3 provides the granularity of information to reduce energy consumption and not just improve the PUE metric. PUE (and its inverse, DCiE) becomes a more helpful 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. You can then attack the numerator and squeeze inefficiencies out of the infrastructure.

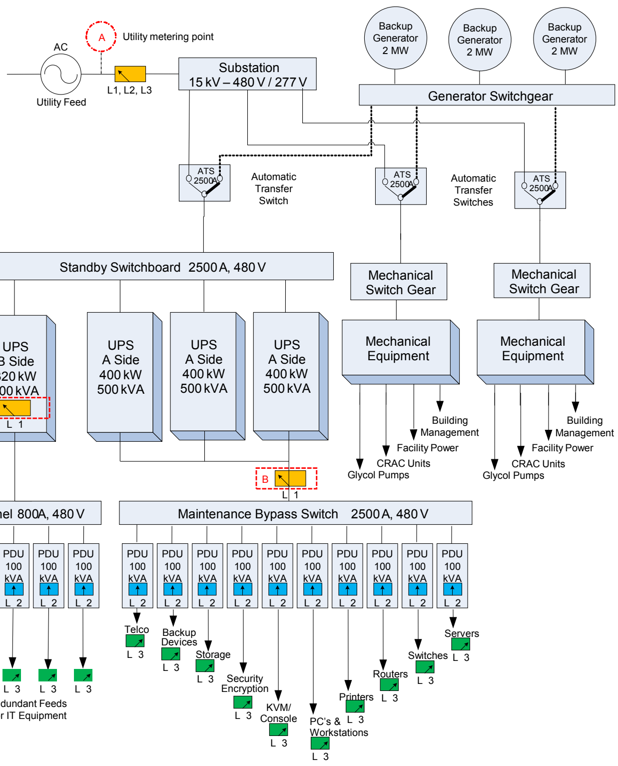

Figure 24: Measuring at the Three PUE Levels

(Source: The Green Grid)

4.4.4 The Advantages of High Voltage

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 PDU at 40A (32A rated) circuit provides 11.5kVA. A three-phase 415V PDU at 20A (16A rated) circuit provides 11.5kVA.

Running higher voltages at lower currents means utilizing smaller cables that use less copper, weigh less, take up less space, and, most importantly, cost less. Running three-phase power instead of single-phase power means fewer cables, simplifying deployment and resulting in smaller cables. Plugs and receptacles are less expensive at higher voltages and lower current ratings. A plug/receptacle combination of a 30A 415V Wye is typically $75 (USD), compared to a 60A 208V Delta at $350 (USD)—prices listed are estimates and may vary by quantity purchased, region, and product availability.

There are other benefits to higher voltages. A 400/415V power circuit will eliminate voltage transformations and reduce energy costs by 2-3 percent relative to 208V distribution and 4-5 percent relative to 120V distribution.

Consolidating data centers will reduce total power consumption and may create opportunities for using high-density racks and high-power PDUs. For example, a 42U rack filled with 1U servers consuming 250W each draws 10.5kW, which requires 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 requires two three-phase 208V, 80A or two three-phase Wye 400V, 50A PDUs. These examples allow enough headroom should one of the feeds fail. It also supports the North American requirement for 80 percent 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.

Summary

- The simplest power deployment to a rack is a single appropriately sized power feed to a single PDU.
- The next step up in availability is still a single feed to a single PDU, but with the addition of a rack transfer switch that typically has two feeds from the same or different building feeds.
- The most reliable deployment is two power feeds to two PDUs, so if one PDU or power feed fails, a second one is 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.
- Deciding how much power should be deployed to a rack requires knowing the device voltage and, ideally, the current draw.
- Three-phase power enables one whip or PDU to deliver three circuits instead of just one. One slightly thicker three-phase cable will be significantly smaller and weigh less than three single-phase cables for the same voltage and amperage.
- Many factors are involved in selecting a PDU: data center location, application requirements, IT equipment requirements, available power, space in the rack, energy management, and efficiency goals.
- Deployment of intelligent PDUs can make a significant difference in your ability to improve uptime and staff productivity, efficiently use power resources, make informed capacity planning decisions, and save money.
- The Green Grid's Power Usage Effectiveness, or PUE, is the most used data center power efficiency metric. The best practice is to gather data on typical power usage over time to ensure the peaks and valleys have been captured in calculating your PUE to establish a baseline and track your improvements.
- Running higher voltages at lower currents means smaller cables that use less copper, weigh less, take up less space, and cost less. High-density racks can be effectively deployed in small, medium, or large data centers.
- Load balancing should be used to evenly distribute the rack equipment's current draw among the PDU's branch circuits and phases, ensuring the most headroom for each branch/phase while mitigating upstream power issues like overloads and breaker trips.
- 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**

5.1 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 related primary forces are influencing PDU development and innovation trends. First is the demand for increased power and density of IT equipment at the rack or compute density per U of rack space. Second is the industry-wide mission to create energy-efficient or "green" data centers, including carbon footprint reduction. Both trends challenge PDU vendors to improve hardware and software design; the second requires IT and facilities organizations to better understand how data center power is consumed and to ensure measures are taken to reduce and optimize it.

The above trends are underscored in Omdia's "The World Market for Rack Power Distribution Units—2022 Edition" report, which says that the Compound Annual Growth Rate (CAGR) for PDUs between 2020-2026 is predicted to be:

- <5kW is expected to have a unit growth of -4.84% (-3.49% for North America).
- 5kW – 10kW are expected to have a unit growth of 5.8% (7.48% for North America).
- >10kW is expected to have a unit growth of 9.82% (11.47% for North America).

In the same report, when comparing the growth rate of single-phase PDUs to three-phase PDUs, analysts predict:

- Single-phase PDU units are expected to grow by 0.15% (2.67% for North America).
- Three-phase PDU units are expected to grow by 8.45% (9.79% for North America).

Omdia continues by reporting that when comparing the growth rate of non-intelligent PDUs to intelligent PDUs, analysts predict:

- Non-intelligent PDUs are expected to have a unit growth of 0.79% (-0.88% for North America).
- Intelligent PDUs are expected to have a unit growth of 9.92% (10.01% for North America).

These analyst predictions help support the continuing trends driving intelligent PDU market demand the need to support increased power consumption and higher densities at the rack.

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 places enormous power and usage demands on PDUs. For example, four or more blade chassis in a single rack could draw over 20kW of power, creating both power and cooling challenges. From a power perspective, racks will require three-phase power with 60A, 80A, or even 100A of service. Some data centers are bringing 400/415V three-phase service to the rack to accommodate the power demand while increasing efficiency from reduced voltage step-downs. Similarly, others are packing dozens of 1U servers into a single rack and pressing 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. A rack containing several such servers could experience much more power consumption, driving the need for additional visibility into server power loads in the rack to manage power capacity.

More power consumption means more cooling to remove the added heat. PDU manufacturers will be expected to supply basic environmental sensors for heat, humidity, and airflow. These metrics help better understand overall environmental conditions and identify zones that must be fine-tuned or supplemented with dedicated or specialized cooling (see section 6.1).

5.1.1 Customizing IT Equipment for Power Efficiency

One trend to watch, especially for the largest data centers, is designing and deploying customized equipment to maximize power usage efficiency. For example, Meta and The Open Compute Project have each begun to deploy unique hyperscale solutions, taking AC power in and distributing it out in DC to devices. Other solutions include 400/415V solutions, including 480V three-phase Wye power where each line is wired to the neutral, so the outlets deliver 277V to increase efficiency. This approach is very efficient yet highly specialized since today's IT equipment is not typically built with power supplies that support 277V. Furthermore, common data center receptacles are IEC C13 and C19, which do not support 277V.

The savings and efficiencies (1-2 percent over 415/240V three-phase systems) are enough for Meta and The Open Compute Project to justify building custom triplet racks, custom servers with custom power supplies, custom battery/UPS, and 480V/277V PDUs with custom TE Connectivity 3-pin MATE-N-LOK receptacles. The key is creating more efficient power delivery systems while maintaining uptime and service level agreements.

5.2 Increased Intelligence at the Rack to Support Efficiency Initiatives

Many data centers have grown larger and more complex as consolidation continues (shrinking a data center's footprint by moving workloads to the cloud). With increasing size and complexity, there is a greater need to drive intelligence to the IT equipment at the rack to create what the industry refers to as the "intelligent" or "smart" 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 most of the power is or should be, consumed. Proper monitoring and metering of the IT equipment and environmental sensors at the rack will collect the data necessary to produce the most significant overall efficiency, savings, and operational improvements. 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 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.

In general, capacity planning allows you to identify the racks into which you can install new devices. It is not just knowing which ones are available "today" but knowing that there is enough capacity to do so in the future. At the rack level, capacity is managing power consumption. If you look at the whole concept of redundancy, if you do not know the overall load of the PDU, then you cannot load it to the proper capacity.

Capacity planning based on equipment nameplate data is no longer enough. Efficiency improvement is an information-driven activity. To formulate and drive the most effective decisions, you must collect IT device CPU utilization data and its corresponding actual power usage. More energy efficiency will be gained if such planning is based on the trends observed from the actual data over time. Furthermore, the data collected at the rack level can be integrated with DCIM software and energy management systems for complete data center and power chain visualization, modeling, and planning. This data collected can lead to further improvements in the data center ecosystem, e.g., computing carbon emissions generated by IT equipment to report on and taking steps to lower your carbon footprint.

You can gain efficiency from policy-based power control software. For example, one can automatically turn servers on and off based on granular power consumption data and a set of pre-established static or even dynamic rules.

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

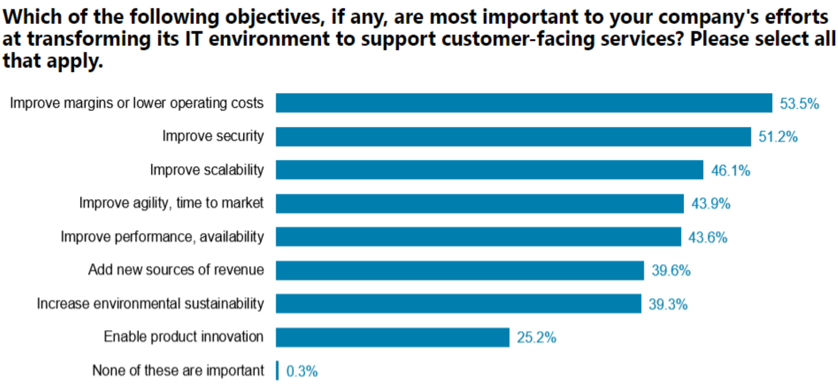

Figure 25: Survey on Objectives to Transform an IT Environment

(Source: 451 Research)

5.3 Integration with Higher-Level Data Center Management Systems

Various software products exist to help you better manage the data center, such as Artificial Intelligence Operations (AI Ops). This application uses artificial intelligence to enhance IT operations and has grown in popularity in recent years. AI Ops measures and analyzes data for failure analysis and preventative maintenance. For example, if you notice your power factor decreasing over time, the data center management software that factors in AI Ops will help to detect and suggest remedies.

The data needed to manage data center infrastructure and energy effectively is collected from power devices along the entire power chain, IT equipment, 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 you 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:

- iPDUs measure essential power data at a predefined frequency and store such data in memory.
- The data collection services from the PDU management (or power management) system of the iPDUs through industry-standard management protocols such as SNMP.
- The data collection service can be part of the energy management system. For scalability reasons, data collection is typically delegated to a specific PDU vendor's PDU management system, which is deployed along with iPDUs to administer, maintain, and troubleshoot the PDUs and collect power statistics.
- The 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. DCIM or another energy management system can use the amassed data for further analysis.
- The energy management 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 factors beyond iPDUs.

With the advanced analysis conducted by an AI Ops tool or energy management system, you can make day-to-day operational and long-term strategic planning and preventative maintenance decisions. Your staff can then provide reliable power for business applications while reducing data center energy consumption waste.

Summary

- Two factors influence the development and demand for iPDUs: the industry-wide goal of creating more energy-efficient data centers and the increasing power and density of IT equipment at the rack. These are also driving the need for PDU intelligence.
- The growing popularity of 1U servers, blade servers, network-attached storage, storage area networks, and multi-gigabit chassis-based network communications gear place enormous power and usage demands on PDUs creating power and cooling challenges.
- Many data centers have grown larger and more complex as the consolidation trend continues. With increasing size and complexity, there is a greater need to drive intelligence to the rack.
- Proper monitoring and metering of the rack's IT equipment and environmental sensors will collect the data necessary to produce the most significant overall efficiency, savings, and operational improvements.
- An essential ingredient of energy-efficient behavior is individual awareness and accountability for energy usage.

6.0 Environmental Data Center Management and Monitoring

With an increasing focus on improving data center efficiency, more PDU manufacturers are offering integrations with a suite of sensors to measure and monitor environmental conditions around IT equipment. There are sensors for monitoring temperature and humidity, the difference in air pressure between two locations, the rate of airflow and the presence of particles within it, the presence of water, whether an enclosure door is open or closed, and tracking IT assets.

Sensors can be deployed as a separate overlay network or as part of an existing network. Deploying sensors as an independent overlay network requires an intelligent managing device/controller with dedicated network connections. That means if you have deployed basic or monitored/metered PDUs, a separate managing device with an intelligent controller would need to be deployed. The extra hardware, work to deploy network drops, and necessary U space to house the additional hardware may make an overlay approach to deploying sensors cost-prohibitive.

On the other hand, suppose sensors are deployed as plug-and-play options to iPDUs. In that case, there is no need for a separate managing device's intelligent controller because one will already be available on the iPDU. With plug-and-play sensors, no dedicated network needs to be set up, eliminating the work and cost of cable installation and reducing costs.

In either sensor deployment approach, the data collected by a sensor is communicated to the intelligent controller, which can then be sent to data center management software for further monitoring and management. If the intelligent controller is connected to a web-based GUI, you can report and plot sensor values over time and define custom thresholds and alerts in real time.

6.1 Environmental Monitoring Sensors

Sensors that monitor temperature, humidity, airflow, differential air pressure, dust/particles, and water/leaks are called environmental monitoring sensors. A popular deployment of environmental sensors in a data center's white space is detailed below.

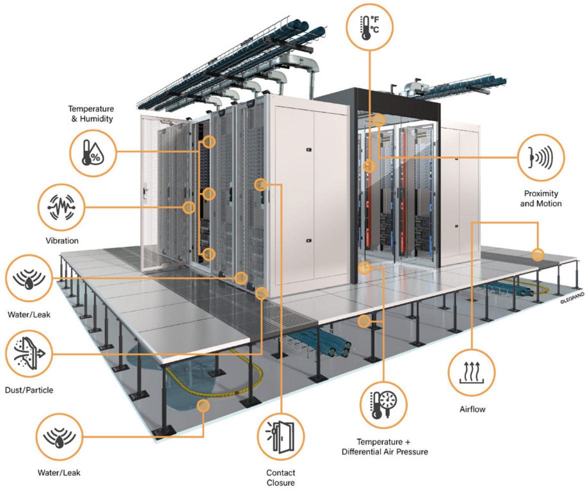

Figure 26: Popular Sensor Deployment in a Data Center

(Source: Legrand, Inc.)

6.1.1 Temperature Sensors

Temperature sensors monitor temperature ranges in data centers. Since IT equipment generates considerable heat, device manufacturers specify a range of acceptable temperatures for proper operation. As the power consumed at the rack has dramatically increased, so has the heat generated by IT equipment. Accurate placement of temperature sensors on a rack means temperatures can be monitored to ensure that IT devices do not overheat and that facility energy is not wasted by overcooling. From a data center plant perspective, the cost of cooling and moving air is a large infrastructure expense, and maintaining IT inlet air temperatures colder than necessary wastes energy and money.

Temperature sensors at the rack can also provide early warning about temperature extremes, hotspots, or cold spots and help identify when an HVAC (Heating, Cooling, and Ventilation) system has become unbalanced. To ensure that IT equipment gets enough cool air, the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) has recommended that temperature probes/sensors be placed at specific locations along equipment racks (see Figure 28). A sensor-capable iPDU should allow thresholds to be customized, for example, by sending automatic alerts about temperature fluctuations.

6.1.2 Humidity Sensors

Understanding the basics of humidity and how it affects your data center or server room can affect how long your computer equipment lasts and how much your electricity bill costs. Humidity is a measurement of moisture in the air, and high levels can cause condensation 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 a safe range. ASHRAE also recommends humidity ranges for data centers, as outlined in Figure 27.

Classes	Equipment Environmental Specifications						
	Product Operations					Product Power Off	
	Dry-Bulb Temperature (°C)	Humidity Range, non-Condensing	Maximum Dew Point (°C)	Maximum Elevation (m)	Maximum Rate of Change (°C/hr)	Dry-Bulb Temperature (°C)	Relative Humidity (%)
Recommended (Applies to all A classes; individual data centers can choose to expand this range based upon the analysis described in this document)							
A1 to A4	18 to 27	-9°C DP to 15°C DP and 70% RH or 50% RH					
Allowable							
A1	15 to 32	-12°C DP and 8% RH to 17°C DP and 80% RH	17	3050	5/20	5 to 45	8 to 80
A2	10 to 35	-12°C DP and 8% RH to 21°C DP and 80% RH	21	3050	5/20	5 to 45	8 to 80
A3	5 to 40	-12°C DP and 8% RH to 24°C DP and 85% RH	24	3050	5/20	5 to 45	8 to 80
A4	5 to 45	-12°C DP and 8% RH to 24°C DP and 90% RH	24	3050	5/20	5 to 45	8 to 80

Figure 27: Recommended and Allowable Humidity and Temperature Ranges

(Source: ASHRAE. See standard for specific details.)

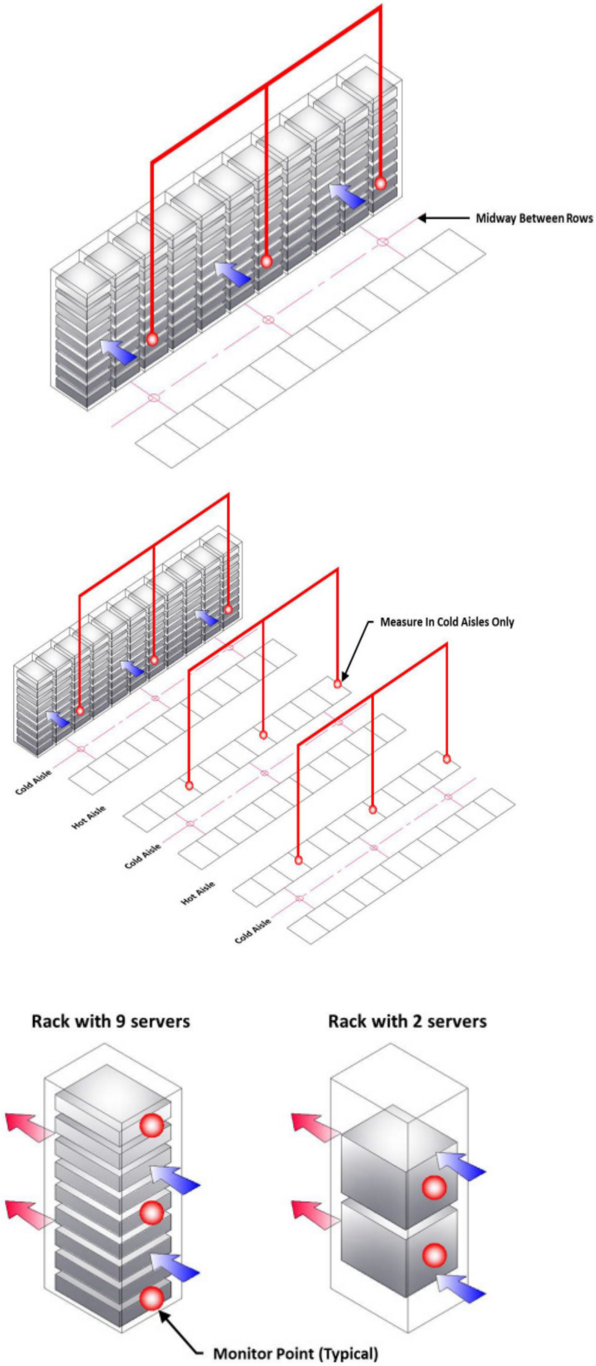

Figure 28: Recommended Temperature Sensor Placement

(Source: ASHRAE)

6.1.3 Airflow Sensors

Airflow sensors detect a reduction of air movement that might create the potential for overheating, which can damage IT equipment. There are two primary areas for monitoring airflow in the data center—above the floor (monitored at several points) and below (monitored at select points). Like other environmental sensors, airflow sensors should have set thresholds and alarms to ensure that you are alerted when conditions are less than optimal for efficient cooling.

6.1.4 Dust/Particle Sensors

Airborne particle and dust deposits—such as organic dust, concrete dust, ferrous metal particles, electrostatic dust, and haze due to forest fires—are an invisible threat to data centers, server rooms, and other areas where IT equipment is installed. Although not naturally harmful to IT equipment, these ultrafine particles may build up over time, interfere with a device's normal functioning levels, and even lead to its short circuit. Industry guidelines, such as ISO 14644 and ASHRAE TC9 outline air cleanliness measures to help ensure the air quality around the equipment is maintained at optimal levels to prevent the possibility of overheating.

6.1.5 Differential Air Pressure Sensors

Differential airflow sensors help to ensure that the pressure differential between a subfloor and floor is enough to control air flowing from the subfloor to the floor above. Underfloor air pressure sensors provide feedback to CRAHs (Computer Room Air Handlers), CRACs, or the building management system that controls fan speeds as necessary to meet the underfloor pressure set point. Blockages in under-floor air-supply plenums can cause high-pressure drops and uneven flow, resulting in cold spots where cooling air is short-circuiting to the return path. This can lead to poor efficiency fixes to correct the problem, such as lowering the supply air temperature or overcooling the entire space to address a few hot spots. Maintaining proper room pressure prevents airborne particulates from entering the data center and prevent thermal leaks.

6.1.6 Water/Leak Detection Sensors

Water/Leak detection sensors can be used under or in racks to detect leaks. There are individual sensors and "rope" or "cable" sensors. A rope sensor can be laid under a row of racks and detect water anywhere along its length, including where water was detected. A rope sensor can also be wrapped around pipes or equipment to detect leaks. Water sensors can also detect condensation and the presence of a glycol mixture, a water-coolant mixture that is frequently used in heat transfer and cooling applications. It is important to detect if there is a risk of a leak or an actual leak within the data center's cooling system. The presence of water or moisture can rapidly increase the current in a circuit, which in most IT equipment will cause a short circuit and blow a fuse, damaging IT equipment.

6.2 Physical Access and Security Sensors

It was sufficient to regulate general access to the data center's entry points at one time. You would be okay if you could ensure that no unauthorized person had access to your sensitive digital infrastructure and prove those reasonable measures to auditors. Times have changed. Escalating regulatory requirements across industries now require sensitive systems and data to be subject to specific protections. Ensuring that only authorized staff enters the data center is no longer enough. Now, you must track and monitor each person's access to specific sensitive systems and ensure they are authorized for a particular area, rack, or even a specific device within the rack. And you must be able to provide an extensive audit trail about who touched those systems when and what they did each time.

The primary goal of compliance standards across industries is similar: ensure that your most sensitive systems and data are especially protected against inappropriate access—and that your compliance with regulatory mandates is accurately documented.

It is recommended that when considering a physical access sensor deployment, you need to:

- Ensure that enclosure locks can be remotely administered so that appropriate permissions can be mapped between the right people and the right systems using enterprise security policies and/or ad hoc administration.
- Implement proximity card authentication to make it easy for authorized personnel to quickly gain access to the enclosures for which they are authorized.
- Install in-rack cameras that capture live video and photos that are automatically tagged with relevant data (time, date, user ID, system data, actions, etc.) for audit documentation and forensics.
- Track proximity and motion detection around IT equipment and racks to support audit documentation of such events.
- Ensure interoperability with DCIM and/or other access control systems to facilitate a single point-of-control and easy consolidation of all security/compliance-related audit trails.
- Implement encryption and detection safeguards to ensure the integrity of rack-level security protections and audit systems.
- Define custom real-time alerting and alarming that notifies appropriate parties of problematic events requiring immediate attention.
- Implement incident management processes and procedures, including specifying courses of action, procedures for notification, escalation, mitigation, and documentation.
- Ensure business continuity of security and compliance even in case of a power outage.

It is worth noting that rack-level physical security compliance requirements will continue to evolve as customers and regulators become increasingly concerned about data breaches' potential social and economic impact. It is wise to take a long-term view of your rack-level physical access and security needs—rather than focusing only on what current regulations require.

6.3 Asset Management Tags and Sensors

You need to track many pieces of information for each of the thousands of physical assets in your data center. You must know what this physical asset is, where it is, what it is doing and how often, how many ports it has, what they are connected to, how much power it draws from where, and much more.

Whether complying with regulations or organizational procedures, IT equipment audits are required annually and usually more often. By deploying asset management tags and sensors, you can create an accurate, automated, real-time inventory of your IT assets and their locations down to the 1U level. These tags and sensors typically integrate with DCIM software to easily track assets, determine capacity, and manage adds, moves, or changes.

Summary

- With the IT industry's increasing focus on improving data center efficiency, more PDU manufacturers are offering environmental sensors.
- There are sensors for determining temperature and humidity, the difference in air pressure between two locations, the rate of airflow and the presence of particles within it, the presence of water, whether an enclosure door is open or closed, and for tracking IT assets.
- Sensors can be deployed as a separate overlay network or as part of an existing network. Deploying sensors as an independent overlay network requires an intelligent controller with dedicated network connections.
- Temperature sensors at the rack provide early warning about temperature extremes, hotspots, or cold spots and help identify when an HVAC system is becoming unbalanced.
- It is a good security practice to control physical access to critical IT resources in the data center. By implementing door and rack-level security systems and sensors, you can mitigate unwanted and unauthorized modifications to infrastructure.
- By deploying asset management tags and sensors, you can create an accurate, automated, real-time inventory of your IT assets and their locations down to the 1U level.

Reference

Key Terms to Know

4-wire: A 4-wire PDU consists of one ground wire and three lines, 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 as line to line, i.e., L1-L2.

5-wire: A 5-wire system is the same as a 4-wire system but with the addition of a neutral wire. The voltage of two lines is available as line-to-line, and the voltage of one line (line-to-neutral) can be supplied as well.

Active Directory: 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 (W or Watt): The work done by an IT device and is the measurement by which power companies bill for its use. Also referred to as real power.

Air Pressure: Force exerted by the air on any surface in contact with it. Air pressure that is too high will result in higher fan costs and greater leakage, which can short circuit cooling air, while too low pressure can result in hot spots at the area most distant from the cool air supply point.

Airflow: Is the movement of air, or the measurement of the amount of air per unit of time flowing through an object.

Alternating Current (AC): Is energy delivered in a form that can travel the long distances necessary from generating plants to homes and businesses. The term AC reflects that the voltage and current constantly change in value or alternate between a positive and negative threshold over a centerline.

Amperage (Amp): 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 (VA): The product of the voltage multiplied by current. It is the sum of the active and reactive powers.

ASHRAE (The American Society of Heating, Refrigerating, and Air-Conditioning Engineers): Is a professional association that advances heating, ventilation, air conditioning, and refrigeration systems design and construction.

Basic PDU: Power strips that are built out of high-quality components for use in critical IT environments, such as data centers. It distributes reliable voltage and current to multiple outlets but has no metering capability and no network interface.

Branch Circuit: Power feeds originate from a panel, switch or distribution board and terminate into an electrical receptacle mounted in a junction box near the rack.

Circuit Breaker: An automatically operated electrical switch designed to protect an electrical circuit from damage caused by overload or a short circuit.

CRAC (Computer Room Air Conditioning): A device that monitors and maintains the temperature, air distribution, and humidity in a network room or data center.

CRAH (Computer Room Air Handler): A device that monitors and maintains the temperature, air distribution, and humidity in a network room or data center.

Current (A): Is the flow rate of electric charge. Current is measured in amperes.

Delay Curve: A characteristic of a circuit breaker giving the ability to discriminate between normal and damaging overcurrents, which are quite common when powering servers.

Delta: A power delivery topology using 4 wires—3 hot conductors (circuits) and a ground. This configuration gets the name Delta because a schematic drawing of it has three transformers forming a triangle like the Greek letter Delta. The three lines connect to the triangle's three corners.

Derated: In North America, data centers must be certified to UL 62368-1. UL 62368-1 limits a device to draw no more than 80 percent of the rating of its input plug. This 80 percent limitation is commonly known as derated current.

Direct Current (DC): Is the energy that does not alternate over a fixed period but has a steady value with reference to zero. DC does not travel great distances well.

Energy Management (kWh): Driven by the 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.

Floor (or Building) Power Distribution Unit (PDU): Large data center PDUs used earlier in the power chain that takes the form of panelboards or free-standing pedestals.

Fuses: A type of low resistance resistor that acts as a sacrificial device to provide overcurrent protection of either the load or source circuit.

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

The Green Grid: A nonprofit industry consortium of end-users, policy-makers, technology providers, facility architects, and utility companies collaborating to improve the resource efficiency of data centers.

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

HVAC (Heating, Ventilation, and Air Conditioning): HVAC systems control the ambient environment such as temperature, humidity, airflow, and air filtering in the overall data center.

IEC (International Electrotechnical Commission): Is a non-profit, non-governmental organization that publishes classifications for plugs and receptacles used in data centers.

Intelligent PDU (iPDU): PDUs having an embedded network interface that can be accessed and controlled remotely via a web browser or command-line interface.

kAIC (Kilo Ampere Interrupting Capacity): Refers to measurements of a circuit breaker's ability to withstand a short circuit or overload.

Kilowatt-Hour (kWh): A kilowatt-hour 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 application protocol used for directory services authentication. It provides the communication language that applications use to communicate with other directory services servers.

Line: An electrical conductor that 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. Lines are labeled as L1, L2, and L3, or X, Y, and Z.

Load: Is the amount of amps being used by devices. Amps are the standard measurement of electrical current for how much electricity is moving through a wire at a set time. The amp draw is calculated by the electrical requirements of the devices plugged in and is regulated by a circuit breaker or fuse.

Load Balancing: A procedure that tries to evenly distribute the rack equipment's current draw among the phases of a three-phase transformer and/or PDU.

Monitored/Metered PDU: PDUs that measure the aggregate current draw (load) at the PDU level and display the data locally. However, this information cannot be accessed remotely as these PDUs have no network connectivity capabilities. The terms "Metered" and "Monitored" are sometimes used interchangeably.

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

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

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.

Power Factor: The ratio of active (real) power to apparent power. A Power Factor of 100 percent indicates perfect power, while lower values indicate wasted power or distortions. Power companies may charge a surcharge if the power factor is below a threshold.

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.

Provisioning: The process of setting up IT infrastructure.

Rack PDU: A rack PDU is a PDU mounted in a rack or cabinet that provides electrical power to various IT devices such as servers, networking, and storage equipment.

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

Reactive Power (VAR): Is power flowing back and forth between the power company and electrical devices due to capacitance and inductance, the ability to store energy, within the IT device. Reactive power does no actual useful work and is not billed by power companies.

Remote Power Panel (RPP): A remote power panel directly connects power distribution extensions from floor PDUs or other power sources to server racks.

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 call for 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.

Smart PDUs: PDUs that have the added value of offering PDU-level input power measurement capabilities and branch circuit protection, allowing you to measure, monitor, and report power at the rack or PDU. It is available in various voltages and amperages and serve as a central connection point for environmental monitoring, asset location, physical access, and other monitoring and security sensors.

Smart PDUs with Per Outlet Power Sensing: PDUs that build on the features offered in Smart PDUs with the added value of offering device-level output measurement capabilities, allowing you to measure, monitor, and report power down to the outlet/device level. Email alerts provide automated updates on power and environmental conditions.

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

Switched PDUs with Per Outlet Power Sensing: PDUs that have the added value of offering both outlet control and device-level output measurement capabilities, allowing you to measure, monitor, control, and report power down to the outlet/device level. It combines highly accurate, outlet-level power measurement technology for data center device-level power monitoring, alerting, and on/off/reboot control. With outlet control, gain features like power-up sequencing and smart load shedding.

Temperature: The degree of heat present in an object. Temperature must be kept at predetermined set points in the data center to prevent sensitive IT equipment from overheating and prevent costly overcooling.

Three-Phase: An electric power system containing at least three conductors carrying alternating current voltages that are offset in time by one-third of the period. Three-phase service is most often found in either a Wye or Delta configuration.

Truck Roll: The need to dispatch a technician to a remote location for service.

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

Voltage (Volt or V): The standard measure of electrical potential. Higher voltages allow more energy to flow within a given time for a given wire size. Electromotive force (E), or difference in electrical potential, is measured in volts and equal to the current (I) times the resistance (R). E = I x R.

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 for IT equipment, 1W of electricity equals 1W of heat.

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

Whip: Power cables from a floor or building PDU or overhead busway.

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

451 Research, part of S&P Global Market Intelligence, Voice of the Service Provider: Workloads, Key Projects & Infrastructure Evolution 2021 - Charts & Figures.

Alger, D. (2005). *Build the Best Data Center Facility for Your Business: A Comprehensive Guide to Designing and Operating Reliable Server Environments*. Indianapolis (IN): Cisco.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers Technical Committee 9.9. (n.d.). *Index | ASHRAE 9.9 Mission Critical Facilities, Data Centers, Technology Spaces and Electronic Equipment*. Retrieved October 21, 2021, from https://tc0909.ashraetcs.org.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2008). *2008 ASHRAE Environmental Guidelines for Datacom Equipment*. Atlanta (GA): ASHRAE.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2008). *High Density Data Centers: Case Studies and Best Practices*. Atlanta (GA): ASHRAE.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2009). *Best Practices for Datacom Facility Energy Efficiency (2nd ed.)*. Atlanta (GA): ASHRAE.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2009, November 6). *ASHRAE Workshops on Improving Data Center Energy Efficiency and Best Practices*. [NYSERDA sponsored workshop]. New York (NY): ASHRAE.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2018). *ASHRAE IT Equipment Power Trends (3rd ed.)*. Atlanta (GA): ASHRAE.

American Society of Heating, Refrigerating, and Air-Conditioning Engineers. (2021). *Thermal Guidelines for Data Processing Environments (5th ed.)*. Atlanta (GA): ASHRAE.

Gartner. (2005). *Use Best Practices to Design Data Center Facilities*. Research ID Number G00127434, Publication Date: 22 April 2005. Stamford (CT): Gartner, Inc.

The Green Grid. (2014). *PUE: A Comprehensive Examination of the Metric*. Atlanta (GA): ASHRAE.

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

Hewlett Packard Enterprise. (2022). HPE ProLiant DL380 Gen10 Server Specifications. Spring, TX: HPE.

Informa (2022). Omdia Rack PDU Annual Tracker – 2022. [Microsoft Excel file]. London (UK): Informa.

International Electrotechnical Commission. (2018). *Audio/Video, Information and Communication Technology Equipment*. Newark (NY): IEC.

International Organization for Standardization. (2015). I*SO 14644-1:2015(en) Cleanrooms and Associated Controlled Environments—Part 1: Classification of Air Cleanliness by Particle Concentration*. ISO. Retrieved 21 October 2021, from https://www.iso.org/obp/ui/#iso:std:iso:14644:-1:ed-2:v1:en.

Legrand, 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]. https://www.raritan.com/resources/white-papers

Legrand, Inc. *Data Center Power Overload Protection: Circuit Breakers and Branch Circuit Protection for Data Centers*. [White paper]. https://www.raritan.com/resources/white-papers.

Legrand, Inc. *Deploying High Power to IT Equipment Racks*. [White paper]. https://www.raritan.com/resources/white-papers.

Legrand, Inc. *Environmental Data Center Management and Monitoring*. [White paper]. https://www.raritan.com/resources/white-papers.

Legrand, Inc. *Embracing the Digital Data Center: How Environmental Monitoring and Visibility Lead to Operational Insights*. [White paper]. https://www.raritan.com/resources/white-papers.

Legrand, Inc. Three-Phase Power Explained. [Video]. YouTube. https://www.youtube.com/watch?v=iMn7dq7B1oo.

Legrand, Inc. *Rack-Level Security and Compliance: Five Ways to Fulfill Tougher Requirements with Limited Resources*. [White paper]. https://www.raritan.com/resources/white-papers.

Legrand, Inc. *Alternating Phase: A New Look at 3-Phase Power Distribution*. [White paper]. https://www.servertech.com/support/technical-references/white-papers.

National Fire Protection Association. (2019). *NFPA 70: National Electrical Code 2020*. Quincy (MA): NFPA.

Neudorfer, Julius. (2020). Data Center Frontier Special Report: Understanding Importance of Power Quality in the Data Center. [White paper]. Lawrenceville (NJ): Data Center Frontier. https://www.datacenterfrontier.com/energy/whitepaper/11431799/understanding-the-importance-of-power-quality-in-the-data-center.

The Prysmian Group. (2017). *Cord & Cordset Products for Industrial, Commercial, and Specialty Applications*. Highland Heights (KY): The Prysmian Group (formerly General Cable). https://general-cable.dcatalog.com/v/Cord-Cordset-Products?page=1.

S&P Global Market Intelligence. (2021). *Voice of the Service Provider Workloads, Key Projects & Infrastructure Evolution 2021 – Charts & Figures*. New York (NY): 451 Research.

Underwriters Laboratories, Inc. (2013). *UL 60950-1, Information Technology Equipment—Safety—Part 1: General Requirements*. Northbrook (IL): UL, Inc.

Underwriters Laboratories, Inc. (2015). *UL 1077, Standard for Supplementary Protectors for Use in Electrical Equipment*. Northbrook (IL): UL, Inc.

Underwriters Laboratories, Inc. (2016). *UL 489, Molded-Case Circuit Breakers, Molded-Case Switches, and Circuit-Breaker Enclosures*. Northbrook (IL): UL, Inc.

Underwriters Laboratories, Inc. (2019). *UL 248-5, Standard for Low-Voltage Fuses—Part 5: Class G Fuses*. Northbrook (IL): UL, Inc.

Underwriters Laboratories, Inc. (2019). *UL 62368-1, Standard for Audio/Video, Information and Communication Technology Equipment—Part 1: Safety Requirements*. Northbrook (IL): UL, Inc.

Uptime Institute Research. (2023). Uptime Institute's Global Data Center Survey Results 2023. [Survey Summary Report]. Uptime Institute. https://uptimeinstitute.com/resources/research-and-reports/uptime-institute-global-data-center-survey-results-2023.

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

Some general electrical term descriptions used in this handbook were cited from Wikipedia at https://en.wikipedia.org.

Index

Notes

Notes

WHY SERVER TECHNOLOGY

Server Technology's power strategy experts have provided power solutions for labs, data centers, branch offices, and telecommunications operations for 30 years. Over 60,000 customers around the world rely on our cabinet power distribution units and award-winning power management solutions to reduce downtime, facilitate capacity planning, improve energy utilization, and drive efficiency. With the best quality, best technical support, and most patents, Server Technology products provide uncompromising reliability, innovation, and value for the data center.

Only with Server Technology will customers Stay Powered, Be Supported and Get Ahead.

WHY RARITAN

Raritan, a brand of Legrand, delivers more than 30 years of innovation and experience from the data center market to your IoT and smart city network operations centers. Our products remotely monitor every critical aspect of your gateways and edge data centers:

- Remote Access & Control to Reduce On-site Maintenance
- AC & DC Power Management
- Technology Extensions to Virtually Any Internal Applications with our Xerus Technology Platform

See how Raritan can help you reliably deliver power, access, and control for your smart city applications.

WHY LEGRAND

At Legrand, we deliver access to power, light, and data to millions of spaces around the world. Learn more at www.legrand.us.

Raritan Americas, Inc.
sales@raritan.com
800.724.8090
rartian.com

Server Technology, Inc
sales@servertech.com
800.835.1515
servertech.com