Web Exclusive

Understanding Power Supplies for POWER9

August 2018 | by Jaqui Lynch

In 2015, I wrote an article on planning for power that looked into the requirements for power connections and the rules about power distribution units(PDUs) when implementing IBM servers. IBM publishes maximum power usage for each of their servers in the technical overview manuals in the operating environment section. Since then the servers coming out require even more power and I'm running into power issues even more often.

A great deal of the confusion about electrical power is centered around what is actually being measured, as well as questions about why IBM requires so many PDUs when the measurements show less power is actually being used. In order to clear this up we will first cover some terminology and the discuss PDUs, power cords and other items to think about when planning for power.

One thing I have been running into is hosting sites – many of them only allow up to 10kVA per rack, which is basically 2 x 30 amp PDUs (4.99 usable kVA each). The Lenovo 3950 HANA server with 1400w power supplies can draw up to 7.010 kVA, each E980 node can draw up to 5.57 kVA and the E950 has a maximum draw of 3.9 kVA. As you can see you can rapidly run out of power on a pair of 30 amp PDUs which is why power planning is critical. Further down I'll talk about why you need double the PDUs—basically, it's to support redundant power requirements.

Terminology

The IBM technical overview and introduction Redbooks provide the maximum power consumption for the various servers. Measurements will show you're using less than that, but you still need to plan for the maximum in case you end up using all the resources at some point. Measurements they provide typically consist of watts, kVA (kilo volt amps) and voltage. Depending on the server these may be per server or per node. The same applies to the xSeries servers.

Volts is a measure of the strength of an electrical source at a given amperage (similar to water pressure in a pipe). Amps (amperes) is a measure of the total electrical current that can flow (similar to water current). kVA (kilo volt amps is the key number typically provided by vendors for power planning. It is 1000 volt amps. A volt amp is the unit used for apparent power in an electrical circuit. Volt amps are used with AC (alternating current) power. Other calculations are available for DC (direct current) power.

I normally plan using kVA but if you want to plan using Amps then the following calculations can be used:

Single phase Amps = ((kVA x 1000)/voltage) or (Watts / (voltage x power factor)) Typically, we divide by 208 for voltage. Calculations for both single and 3 phase can be found at rapid tables and in many other places.

In addition to power, you'll also need to plan for air conditioning and UPS systems. BTU (British thermal unit) is an energy unit that is used for air conditioning and watts or kVA are typically used for the UPS systems. All of these are published in various IBM Redbooks. I recommend putting together a spreadsheet for each rack in your data center that includes kVA, amps, watts and BTUs—that way you can ensure you have the information you need at all times. It also helps to keep a map of how server power supplies are connected to the various PDUs.

PDUs and Line Cords

The first key thing to understand the impact on power of the National Electrical Code. Because of past issues with overloading and fires, etc, PDUs in the USA are derated. This means that a single phase 208v 60 amp PDU is only allowed to have up to 48 amps (9.98kVA) of equipment connected to it and a single phase 208v 30 amp PDU is derated to 24 amps (4.99 kVA). A three phase 30 amp PDU is allowed to use up to 8.65 kVA. (see Figure 1).

For the IBM PDUs, it's actually the line cord that determines the amperage as the current PDUs from IBM are able to support up to 60 amps (48 derated) depending on the line cord. IBM currently provides multiple PDUs—the 9188/7188, the 5889/7109, the 7196+ and the new intelligent PDUs (they all start with EP). The first four PDUs (9188/7188/5889/7109) have six pairs of IEC320-C13 outlets rated at 200-240 volts. Each outlet is rated at 10 amps and each pair of outlets is connected to a 20 amp circuit breaker in the PDU with each circuit breaker derated to 16 amps. The full amperage rating for the PDU (derated) is 16 amps, 24 amps or 48 amps, depending on the line cord chosen. The difference between the PDUs is that the 5889/7109 combination provides power monitoring capabilities whereas the 9188/7188 does not. The 7196+ PDU has a fixed IEC60309 3phase 60 amp power cord only. It has 6 outlets that are IEC 320-C19, each of which is derated to 16 amps. There is also an HVDC 240 to 380 volt DC 90 amp PDU, but I've never seen this used in the U.S.

The newer PDUs from IBM are as follows:

1. EPTG, EPTJ – these are 16 to 63 amp 200-240v single or three phase PDUs that include 9 x IEC320-C19 and 3 x IEC320-C13 plugs. The C13 plugs are difficult to get to.
2. EPTK, EPTL – these are 208v 60 amp 3 phase PDUs with 9 x IEC320-C19 and 3 x IEC320-C13 plugs. Again, the C13 plugs are difficult to get to.
3. EPTM, EPTN – these are the most common and provide for 16 to 63 amps 200-240v single or three phase. These have 12 x IEC320-C13 receptacles.
4. EPTP, EPTQ – this is the same as the EPTM/EPTN but is 60 amp 208v 3 phase only.

Some of these PDUs can be vertically mounted and some can only mount horizontally. Additionally, some of the high-end servers require horizontal mount PDUs so attention needs to be paid to the requirements for PDUs.

The two most common line cords to the PDUs that we see in the USA are:

```
6492            14' 200/240v/48A  UTG0247 IEC309 60A
6654            14' 200/240v/24A  UTG0247 L6-30P 30A
```

The 6492 uses a plug type 262P6W which means the electrician needs to provide a 60 amp circuit with a 363P6W receptacle. The 6654 uses a NEMA L6-30P plug which requires a 30 amp circuit and a NEMA L6-30R receptacle. There are multiple other line cord options so it is important to discuss these with your electrician and data center personnel to ensure the correct ones are ordered. If you are using non-IBM PDUs then you will need to know the voltage and derated amperage of the PDUs in order to plan for your power needs. You may also need to get horizontal PDUs for some servers. Details and pictures can be found at the line cord reference and in Figure 1 below.

Designing for Power

Much of the confusion comes from the way redundant power supplies work and the requirements this puts in place for additional PDUs. As an example, the E950 has 4 power supplies set up as 2 redundant pairs. This means that 2 cords should go to one PDU and 2 should go to another so that if a circuit or PDU fails the server can continue to run. For the most part, two of the power cords may never draw power, however we don't get to choose which ones will draw power. This means that redundant power supplies can hide the problem of accidentally exceeding the 4.99 kVa loading on a 30 amp PDU in normal operation, since we will never see one of the PDUs in the pair being used. However, if one PDU in a

redundant pair happens to fail, for example, if a PDU cord is unplugged or an input circuit fails, the entire load of all the devices will be forced to the remaining PDU where it is likely to fail due to the overloading. At that point one of several things will happen – the circuit breaker will trip and power everything off, the equipment on that PDU may power off or you could have a fire if the circuit breakers do not do their job. I've seen overloaded PDUs several times and the result is never good.

The first thing you need to determine is the amperage (and derated amperage) of your circuits. POWER8, POWER9 and x3850/3950 HANA servers draw plenty of power, so wherever possible it is recommended you go with 60 amp circuits. If you remain with 30 amp circuits then you will need to provide a lot more PDUs. Once you determine the amperage and voltage, then it is a simple case of planning connections such that you don't exceed the rules for power usage. Here's an example for a system consisting of an E950, a 2 node E980 and an HMC and monitor.

Using the maximum use numbers from the Technical Overview redbooks: Each 980 node is 5.57 kVA and 5422 watts, the E950 is 3.9 kVA and 3850 watts and the 7063-CR1 HMC is about 0.33 kVA and 300 watts. The monitor requirements are negligible (.022 kVA and 22 watts for a TF4). So for planning purposes we need:

	kVA	Power Plugs
2 node E980	11.14	8 (2 pair per node)
E950	3.90	4 (2 pair)
HMC	0.33	2 (HMC)
7316-TF4 monitor	0.02	1
TOTAL	15.39 kVA	

If we're using 30 amp (derated to 24 amp) PDUs then we can only support up to 4.99 kVA on a PDU so the configuration above would require 6 PDUs (3 to support the primary needs and 3 for failover). With 60 amp PDUs we can support up to 9.98kVA so we would only need 4 PDUs (2 PDUs for primary and 12for failover). If you switch to three phase PDUs then you can reduce a little further. And this does not include any storage or other devices you may need to plan for.

Summary

IBM has taken great care to ensure their support statements adhere to the electrical codes. Even though you can use the energy estimator to determine numbers closer to what you would really use it is still important to plan for enough power to support the servers if they were to have all the resources activated or in use. I normally plan for power using the energy estimator as well as the published maximums and I use the larger of the two numbers. I do not expect to draw that amount of power, but in order to be supported and in compliance with electrical codes I have to plan for the potential additional power. Plus I also ensure that the requirements for redundancy are taken into account and that IBM's rules for maximum servers per PDU are followed. This ensures that the servers are fully supported and that PDU failure does not cause servers to go down due to lack of power or, worse, does not cause fires.

These redundancy requirements along with the regulations mean that you need to plan for what you could use, not just what you are currently drawing when all is normal. This is why it's important that any time a server or I/O drawer or similar is added a review should be done of the PDU layouts and usage. This is also why it is important to have a preinstall planning meeting before servers or upgrades are ordered.

It may also be time for you to review your current environment to determine if you are adhering to these standards. I see a lot of locations that were planned out correctly, but over time additional equipment has crept into the rack and no one realized the power draw that some of this equipment has, especially the new x3850 and x3950 HANA servers.

USA	PDUs			Per Pair
Circuit			NEMA	Usable
Amps	Volts	Phase	Receptacle	KVA
20	120	single	L5-20R	1.92
30	120	single	L5-30R	2.88
20	208	single	L6-20R	3.33
30	208	single	L6-30R	4.99
50	208	single	L6-50R	8.32
60	208	single	L6-60R	9.98
20	208	three	L15/L21-20R	5.76
30	208	three	L15/L21-30R	8.65
50	208	three	L15/L21-50R	14.41
60	208	three	L15/L21-60R	17.29

References

< Return to main article

Print Email

For more information on electrical power, check out the following reading materials:

Useful Power Formulas
http://www.rapidtables.com/electric/ampere.htm
https://www.rapidtables.com/calc/electric/index.html

Supported PDUs and PDU Power Cords (IBM)
https://www.ibm.com/support/knowledgecenter/en/POWER9/p9had/p9had_pdupowercords.htm
https://www.ibm.com/support/knowledgecenter/en/POWER8/p8had/p8had_specsheetpdu.htm

Lenovo and IBM Intel Power Requirements
http://systemx.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.systemx.common.nav.doc%2Foverview_rack_servers.html&cp=0_1

E980 Technical Overview and Introduction
http://www.redbooks.ibm.com/Redbooks.nsf/redbookabstracts/redp5510.html

E950 Technical Overview and Introduction
http://www.redbooks.ibm.com/abstracts/redp5509.html

< Return to main article