Mirazon®

UPS and Power Management

Brent Earls, Mirazon

November 2017

Contents

Introduction

Electricity: that wonderful thing that we all take for granted that we'll always have when we work in IT. It's so ubiquitous in what we do that we don't even consider it often, just plug things in and go. However, if you've been in IT for very long, you know that taking it for granted can quickly come back to haunt you in some very painful (emotionally and physically) ways.

Our quick disclaimer that starts everything: I'm not an electrical engineer, an electrician, or anyone else who has a deep physics layer understanding of electricity. This is real world experience and research that is presented in this document.

To further this disclaimer, we're going to levelset everything at a high level. Electrical power comes in MANY forms, and you'll see devices rated for many different power measures. This article focuses on North American power, which operates most commonly at 120 volts (V) and a frequency of 60 cycles or Hertz (Hz). We're using 120 V as a generic statement for 110 V, 115 V, and 120 V ratings, which you'll see on different devices. It's really splitting hairs at this point, as they're all the same TYPE of power, just deviations due to attenuation and fluctuations; it all should be 120 V at the wall.

We're going to be talking about 208 V for our higher voltage connections, though, which is particularly distinctive from 220 V, 230 V or 240 V. While 220 V, 230 V and 240 V are TYPICALLY single-phase power, 208 V is three-phase power. Additionally, 220/230/240 V is normally two 110/115/120 V connections in the same receptacle to be used by the appliance, whereas 208 (obviously) is more complicated than that. Most datacenters and large power draw locations rely on three-phase because it's more efficient, so we'll talk about that. All that being said, most UPS/servers have a variability on input that works with 208 or 220/230/240 V, so check the spec sheet. Again, specific details of this are beyond the scope of this article, so from here on out we're just going to talk about 120 V and 208 V.

There are key things to keep in mind when you're planning power for a wiring closet, server room, or full-sized datacenter. The basic premises remain throughout, it's just the scale that really changes. This document focuses on a few things: receptacles, plugs, cables, power distribution (PDU), uninterruptible power supplies (UPS), cabling, and how to size all of it. The document has both high-level key concepts and ideas as well as detailed information about how to properly size things, but we'll start simple, with cable types, plugs, etc. and then work into the more difficult concepts, like what volts, amps, wattage, and volt-amps are and what it all means for sizing your environment.

Keep reading for a journey into all things power.

Power Definitions

Let's start with some basic definitions of the various terms we're going to use throughout the rest of this paper.

Voltage

You can think of voltage as how wide the path is for power to a device, similar to how many lanes there are on a highway.

For normal uses, the two voltages to concern yourself with are 120 and 208, 120 being the more common within American homes and 208 more common in datacenters and for the higher power draw equipment. However, 120 V is still available for datacenters. Most modern devices function equally well on 120 V or 208 V. Almost every device's power supply accepts both types of power.

Amperage

Amperage is how much electricity is available to flow to a device. This is why you can plug a .5 amp phone charger into a 15 amp electric outlet. It only pulls what it needs, but it has thirty times as much power available to it as it wants, so you can unplug that phone charger and plug in a hair dryer, which uses 13 amps. You can think of amperage as how quickly things are flowing through the path to a device, like the speed limit on a highway.

Volt-Amp (VA)

VA is the key metric to be concerned with in regards to power draw. VA is amperage multiplied by voltage. For example, three amps on a 120 V outlet is 360 VA, whereas three amps on a 208 V outlet is 624 VA. So, if a server needs 500 VA with 120 V, that's 4.2 amps, but with 208 V it's only 2.4 amps. This highlights the clear difference in how you get power to a device. You can have a narrower path (120 V) and run it at a higher flow (amps), or you can have a wider path (208 V) and run it at a lower flow to get the same volt-amps in the end. This is why a 20 amp circuit on 120 V can supply (in theory) 2,400 VA, but a 20 amp circuit on 208 can supply (in theory) 4,160 VA. That theory is, of course, not accounting for the aforementioned U.S. electric code that prohibits circuits from exceeding 80 percent, but we'll go with it for an example.

Since volt-amps are what we need to run power supplies in devices, and 208 V can get one-and-three-quarters times the VA with the same amperage, it makes a lot of sense to go with 208 V for higher VA needs. For example, if a large server needs 1,200 VA of AC power, with 120 V that means that 10 amps of power for that single server, versus 5.8 amps with 208 V. If the two circuit options are a 15 amp/120 V or a 15 amp /208 V, the single server will utilize almost the entire 120 V circuit (12 amp maximum usage), whereas the 208 V will still have over half of its amperage leftover for another server (or series of servers).

Wattage (W)

Watts are the actual power used by equipment. For example, the total power a device may use is 1,200 W. Depending on the equipment, that might mean it draws 10 amps on a 120 V circuit, or it could mean something very different.

Volt-Amp versus Watt

This is one of the places where power sizing is confusing: wattage and volt-amps. UPSes are (almost) always listed by their VA rating, while equipment is almost always listed in watt rating. The relationship between these two is confusing, so let's start with an example:

Watts are the real power drawn by equipment, whereas Volt-Amps are the apparent power, determined by the voltage applied to the equipment multiplied by current drawn. Think of it like this: sometimes you have to work harder to get the same output than other times. The difference between the two comes down to phase displacement, or lag, between the voltage and current, but let's not get into that now. What do we actually need to know about the two of these?

VA, at their basic level, are used purely to size circuits, wiring, etc. for a device. This is how many volts and amps you need at a given time, whereas wattage is what that device is actually DOING with that electricity. It's a nuanced difference, but it's important. The difference to the lay person, suffice it to say that volt-amps are always larger than watts. It's easy to convert from one to the other if you have the information you need.

Receptacles, Plugs and Cables

Now that we have some definitions, let's start with one of the most basic premises of power: connectivity. Unfortunately, not even this is as simple of a matter as you would hope, because there are three different standards organizations involved in this. The National Electrical Manufacturers Association (NEMA) is responsible for our day-to-day wall receptacles and the plugs that go in them. They also specify what the receptacles and plugs are "capable" of, such as 125 V. The International Electrotechnical Commission (IEC) is responsible for the connectors that go on electronic equipment, both the receptacle and the plug, as well as specifying what those receptacles and plugs are "capable" of. "Capable" is in quotes because there is another governing body, concerned with safety, who specifies what those receptacles and plugs are actually allowed to transfer in the real world, the Underwriters Laboratories (UL). Because of these differences in governing/certifying bodies, you'll see multiple different voltages/amperages on things. We will point out some of the differences moving forward.

NEMA Receptacles and Plugs

For people familiar with normal home appliances and electronics, they only know two plug types, the NEMA 5-15R and 5-20R. Before we get into that, let's look at the naming convention a bit:

5-15P/5-20P and 5-15R/5-20R

The 5 is the NEMA standard the connector is part of (NEMA 5, 125 V, three-wire connectors). The second number is the amperage the connector is rated for (15 amps in this case) and the letter is a P or an R, depending on whether it is a Plug or a Receptacle. Once dissected, it's pretty intuitive and makes it easy to quickly identify a plug. It's worth noting that the voltage/amperage listed here is NEMA's specification, it's not what's actually approved for regular "safe" use -- that's what the UL does.

Getting back to the receptacles that normal people are used to, we have the 5-15R and the 5-20R. These receptacles are 120 V with 15 and 20 amp receptacles, respectively. They look like this:

5-15R, courtesy of Schneider Electric

5-20R, courtesy of Cooper Industries

The receptacle to the left is a 15 amp and the one to the right is a 20 amp, the perpendicular extension in the neutral slot being the key visible differentiator. These plugs are found in homes across the (American) nation, but most people probably don't even realize these are two separate types of plugs or that there's a reason they look different. These plugs are also used in server rooms and datacenters, more often for lower power draw or less dense power needs.

The plug for 5-15P will look very familiar to you as well, as it's what's seen on most American devices:

5-15P plug, courtesy of StarTech.com

The plug for 5-20P probably won't look as familiar:

5-20P plug, courtesy of Signal Power and Delivery Systems, Inc.

It’s not uncommon to hear it referred to as “some weird European plug.” Nope, that’s all American; it’s just uncommon to see 20A, 120 V plugs outside of a datacenter, as normally when things need that kind of amperage, they move to L5-20R or 208 V (which we’ll get to later). Another noteworthy thing: the two have one direction compatibility between them. A 5-15P can fit in a 5-20R, but a 5-20P can’t fit in a 5-15R. This means the lower draw appliance can fit in the higher (and only use what it needs), but the higher draw can’t fit into a lower amperage receptacle.

Before we move to 208 V, let’s tackle even more uncommon 120 V plugs, the locking plugs. These are often found on larger power draw devices such as UPSes and heavy machinery. The plug is inserted and then twisted to lock it into place (hence the common name, Twist-lock, though that name is actually a trademark of Hubbell, Inc.).

L5-20R

First, a different type of 120 V/20 amp receptacle, the NEMA L5-20R. This adds a preceding letter to our previous label, (L for locking) but otherwise keeping the same amperage and receptacle/plug nomenclature.

L5-20R outlet, courtesy of Legrand

This is the more common type of plug to see for 20 amps/120 V. One nice thing about the locking plugs is that they include the specifications clearly on the connectors themselves, as you can see above, this is a NEMA L5-20-R, rated for 20 amps at up to 125 V. This is a very nice way to differentiate these plugs, which as we’ll soon see are very similar. The plug, as you may imagine, looks like the receptacle L5-20P.

L5-20P plug, courtesy of Green Electrical Supply

Notice each of the prongs is curved, with one being keyed to go into the L-shaped slot. This pleasantly means there is only one way to plug it in.

L5-30R

Next up is our last 120 V foray that you may run into with servers/datacenters, the L5-30R, which is -- you guessed it – 120 V at 30 amps.

L5-30R outlet, courtesy of Legrand

The L5-30P also looks very familiar:

L5-30P plug, courtesy of JobSite Products

This one is very hard to tell apart from the L5-20R at a glance. The prongs are spaced slightly differently, though. Being less than obviously different at a glance means that you'll often find yourself trying to put an L5-20P into an L5-30R, or vice versa, but they do not fit. Why not, since with the normal NEMA 5 specification, a 5-15P can fit in a 5-20R, they're the same voltage, with just extra amperage? It really comes down to limiting the maximum flow of the whole path. If each worked in the other, it would be possible to plug a 30-amp appliance into a 20 amp plug, at which point you could draw more power than the wiring in the walls could safely supply, risking fire. The main reason that the two are designed to be differently shaped is so that there's no way a plug can be modified or tweaked to fit in the wrong receptacle. There has to be a concerted effort to completely change the plug to make it work.

There are larger amperage 120 V plugs, but they're not very common.

L6-20R/L6-20P

Finally, we get into 208 V, which come in several different types. We'll start with the most common plugs for UPS devices, the NEMA L6 (locking) plugs. In this case, the nomenclature is the same as above, except for these are from NEMA standard 6, 250 volt, three-wire connectors. These are nearly identical to their L5 brethren.

The L6-20R for 208 V at 20 amps:

L6-20R outlet, courtesy of Elliott Electric Supply

The L6-20P:

L6-20P plug, courtesy of Grainger Industrial Supply

And then the L6-30R for 208 V at 30 amps:

L6-30R outlet, courtesy of Home Depot

And the L6-30P:

L6-30P plug, courtesy of Best Materials

So, let's point out the elephant in the room here: it's NEARLY impossible to tell these apart at a glance if you don't actually read what's on them. This is a VERY common first-round mistake on getting the proper plugs. People go to buy a UPS, glance at the picture on the website and think, "yeah, I have one of those round three-prong things," buy it and then have to bring in an electrician to fix it. You HAVE to read it. If you don't believe me, let's look at an L5-20R, L5-30R, L6-20R and L6-30R next to each other:

Remember which of those plugs is on the wall of your server room? Can you say for sure which one it was? This is where so many people find themselves in trouble, as it isn't that bad to go from an L6-30R to an L6-20R, but if you need to go from an L5-20R to an L6-30R, you're now changing a lot of variables and the wiring in the wall most likely won't work either. For any draw larger than 208 V/30 amps, things are normally hardwired in, so you don't have to worry about those plugs. Even though there ARE larger connectors, they're rarely used due to the difficulties of such high current that are beyond the scope of this document.

L14-30R/L14-30P

There's one more NEMA plug we need to concern ourselves with, sometimes used as an output for a UPS: the L14-30R.

L14-30R outlet, courtesy of Grainger Industrial Supply

The interesting thing about this plug is that it has four wires: two hot, a ground and a neutral. That means this same plug can be used to power two 120 V circuits, or it can be used to power 208 V circuits depending on how it's configured and plugged in. Similarly, a L14-30P connects to it:

L14-30P plug, courtesy of Grainger Industrial Supply

You can see on the plug itself that it points out it can be 30 amps 125 V/250 V, so it can satisfy either need. There are some very interesting uses for this (a PDU with two 120 V circuits in it with a single plug) or a UPS receptacle that is user configurable on its power output.

Alas, we are not done yet. We still have to accommodate the actual equipment plugs, and just in case you were finally comfortable with the NEMA standards, don't worry, these use a completely DIFFERENT set of standards (just for spite, I'm pretty sure). It's time for us to venture into the International Electrotechnical Commission (IEC) 60320 standards.

IEC Receptacles and plugs

IEC works differently to NEMA with its plug names. Rather than a long, actually quite descriptive name like NEMA uses (L5-20R), IEC simply uses the letter C (for coupler) and then a number to specify which type of coupler it is. The nice thing about the IEC is that the male plugs are always one number below the female receptacle. A C13 is the male version of a C14. This can still be tricky to remember, because which one is the male versus which is female? Do you base it on the prongs, or on the overall plug?

The answer is the overall connector. Try to remember that the one with visible prongs is the higher number (and female connector).

For example, the most common male connector you see is a C13:

C13 plug, courtesy of Kenable

It’s used to plug in home PCs, servers, switches, routers, firewalls, storage, almost everything uses this cable, which is nice. Further, this same plug can do 120 V and 208 V … even better! This plug is rated by the IEC for up to 10 amps, but the UL certifies it up to 15 amps, presumably due to North America using 120 V instead of 240 V. In reality, power supplies on 120 V don’t exceed 1,100 W on a C13, and power supplies on 208 V don’t exceed 1,600 W on a C13.

The most common receptacle that is seen on the back of almost all equipment is the C14:

C14 receptacle, courtesy of Markertek

This pairs up with C13.

Occasionally (mostly on networking equipment) you'll run into a C15 plug, which a C13 but it's keyed (again, just for spite):

C15 plug, courtesy of StarTech

Which mates with a C16 receptacle…

C16 receptacle, courtesy of Conrad

Lastly, on high amperage power supplies (2000 W server power supplies for example) you'll commonly see a C19:

C19 plug, courtesy of StarTech

Which mates to a C20:

C20 receptacle, courtesy of Cable Leader

These are rated for up to 20 amps by the UL (only 16A by the IEC).

Yay, we made it! It only took three standards organizations, 10 plugs, and 10 receptacles. Huzzah!

Let's look at a few cables and then call this whole "plugging it in" thing done, shall we?

Common Cables

There are certain combinations of these standards that occur rather frequently. It's often a combination of all three standards organizations that make these cables possible. Often there will be an IEC connector at one end and a NEMA connector at the other with UL certifying the whole cable not to burn down the server room.

Your standard 120 V/15 amp cable that will plug a server into your house, or some server rooms:

5-15P to C13 plug, courtesy of Cable Leader

NEMA 5-15P to IEC C13

A 120 V/20 amp cable that will plug a server into a high amperage outlet. Note that the cable still doesn't support more than 15 amps:

5-15P to C15 plug, courtesy of Cable Leader

NEMA 5-20P to IEC C13

A keyed 120 V/15 amp spiteful cable that you can plug into a normal outlet:

5-20P to C13 plug, courtesy of PDU Whips

NEMA 5-15P to IEC C15

The 20 amp version of that cable are VERY uncommon, and you'll probably never see one, but it does exist. The cable still doesn't support more than 15 amps.

5-20P to C15 plug, courtesy of PDU Whips

NEMA 5-20P to IEC C15

Moving on to 208 V with the most common 208 V/15 amp plug in the US datacenter:

5-15P to C15 receptacle, courtesy of PDU Whips

IEC C13 to IEC C14

These are often referred to as "extension" cables, and while they can service that need, it's not the intention. They're made to work with PDU that have a C13 coupler on them. Be careful when buying these cables as there are often multiple thicknesses of wire (gauge). The cables labeled as "extension" cables are normally smaller wires (we'll talk about that later) made for residential use, which can cause problems if you use them for high amperage uses.

The keyed C15 to C14 also exists for plugging a spiteful device into a 208 V PDU at up to 15 amps:

C15 plug to C14 receptacle, courtesy of Data Comms Direct

Finally, there is the high amperage (20 amp) 208 V connector:

High amperage C19 plug to C20 receptacle, courtesy of Dell EMC

We made it. Now we can plug anything into anything, but we're not quite there yet … there is still another choice when it comes to the cable; not all cables are created equal, even if they have the same ends on them.

Wire Gauges and the AWG

Wires come in different gauges, which are standardized by the American Wire Gauge (AWG). The lower the AWG, the thicker the copper wires inside the sheathing. It's important to pay attention to that point -- the numbering is counter intuitive, smaller gauge wire is larger than higher gauge wire. Why do we care? Thicker wires can carry more amperage than thinner wires.

Normally it doesn't matter. Your personal desktop doesn't pull enough power for it to actually make a difference, but if you're looking to power a big server with 1,600 W power supplies, having a thicker cable is the difference between being able to power the server and possibly melting your cable and starting a fire. Normally power cables are seen with 18, 16 or 14 AWG wires. The cost difference is normally minimal, and there isn't a penalty for having the larger wires, so you should probably go with the lower AWG (bigger wire) unless you have a specific reason not to. It allows the cables to be used for any server anywhere, rather than having only certain use cases.

So … what do we plug into what, again? Let's move on…

Power Distribution Units (PDUs)

Within server rooms and datacenters, it's important to be able to get the power where you need it. UPS units are normally located centrally, or at the bottom of a rack, and normally don't have enough, or any receptacles on the back of them to plug devices into. This means we need something like a power strip to be able to plug in more things. This is where a PDU comes into play. A PDU is like a specialized, higher-end power strip, similar to what you may put in your office or behind your entertainment center. A PDU, however, is made for larger power draw and to sustain that high power draw all the time. PDUs are created with built-in circuit breakers to prevent overpowering and some provide advanced features such as a screen to monitor the power in/out and remote monitoring over IP. Like cables, this isn't a simple and straightforward decision -- there are 120 V and 208 V PDUs, PDUs with multiple different types of outlets, and different mounting options.

Mounting Options

Let's start with the simplest decision: how do we want our PDU to be mounted? There are primarily two different methods of mounting a PDU, horizontally and vertically. A horizontal mount will simply take 1U (rack unit) and therefore will fit in pretty much every rack on the planet, including two-post racks. In fact, horizontal PDUs are pretty much the only choice for two-post racks and some older, smaller racks as well.

APC 120 V horizontal PDU

APC 208 V horizontal PDU

This makes the horizontal PDU a compelling choice, but there are two downsides. It takes valuable rack space that could be used for the real equipment (a few rack Us per rack add up into substantial loss over several cabinets). Secondly, the cable management from horizontal PDU to the devices can quickly get pretty nasty and confusing. Confusing cabling often leads to mistakes, such as unplugging the wrong device.

This is why most datacenters and server rooms are starting to go with vertical-mount PDUs. These PDUs mount vertically from the top to the bottom of the cabinet, normally to the inside, just beside the rails used for mounting equipment.

This means that these PDUs do not consume any rack units, and they also put the power plugs close to where they're actually needed all the way up and down the rack. Due to that, short power cables can often be used to keep cable management clean and tidy.

However, there are some downsides to these PDU as well. First, they aren't as universally mountable as are the horizontal PDUs, and different PDUs have different mounting mechanisms. Two post racks are very hard to mount vertical PDUs to, for example. Similarly, some smaller cabinets may not have room for equipment to slide all the way in when there are vertical PDUs in place. It's best to try to get a PDU from the same vendor as your rack if you want the most probable compatibility, and pay careful attention to WHERE it's going to mount vertically.

Voltage and Amperage

Past the mounting type of the PDU, the next important thing to figure out is what voltage, amperage and how many circuits you will need. To do this you first need to calculate the draw that will be on each PDU, but we'll discuss that a little later in this paper. Each rack should have at least two separate circuits fed from two different UPS circuits or separate UPS units. Some PDUs are designed to get two separate circuits and have the plugs divided between them, while others only have a single circuit and therefore two would be needed in a rack. Sometimes, depending on power choices, more than two PDUs will be needed in a rack to satisfy all the power needs. With PDUs that have two circuits, it's very important to distribute the load across the two circuits evenly. Essentially they need to be treated as two separate PDUs when it comes to load balancing to make sure no circuit is overloaded.

It's critical to keep track of redundancy while designing PDUs within a rack. Remember that a circuit failure has to be tolerated. A common design is to have two 208 V PDUs, each of which with 30 amp circuits, which gives "30 amps" redundant, even though only 24 amps are usable (80 percent).

Just as a side note: why do we only get 80 percent of our amperage? Well, there are regulations in the US (National Electric Code) that say no circuit can exceed 80 percent of its rated capacity sustained. There can be some spikes, but sustained draw must be at or below 80 percent. This is often listed on manufacturer's websites as "Regulatory Derated Input Current" and it is the number you need to pay attention to, not the rated amps. If there are two 20 amp PDUs in a rack, that means that the total USABLE amperage in that rack is only 15. Despite the fact that there are 40 amps going into the rack, no more than 75 percent of a circuit should be used, and must stay up and functional even if one circuit fails.

Voltage of the PDU is tied in with the amperage decision. What kind of voltage you need depends exclusively on what's supplied to you by your source UPS. What kind of amperage is then dictated by some math around how much total volt-amps you need to push through the PDU. We'll address that in detail in upcoming sections.

Circuits

Amperage is the commodity that you are working with on individual circuits. Whether they are from the building's AC power, on the UPS, or on the PDU, the amperage is the key thing that is being divided up when calculating needs. If there are 20 servers each pulling two amps (either at 120 or 208 V), you need at least 100 amps of circuits. Forty amps are actually needed to run the load, but we found out earlier, you can't use more than 80 percent of a circuit, which means we need 50 amps at a minimum for our load. Then on top of that, we need redundancy, leading us to 100 amps of circuits.

PDU Plugs

At the PDU level, the only plugs we normally have to worry about are:

110 V – 5-15R and 5-20R: Pretty simple. Normally there aren't even any 20 amp plugs in the datacenter, so you only really have to worry about the 5-15P.

208 V – C13, and C19: These are pretty straightforward. The biggest thing to keep in mind is how many C19 receptacles you need if there are a lot of blade chassis or other high-draw plugs. Sometimes additional PDUs just for those high-draw devices are needed, in addition to the PDUs used by the rest of the environment. Also remember just a few of those high-draw devices can use a whole circuit's worth of amperage and not leave anything for the devices with lower needs.

Features of PDUs

There are also various different featuresets that can come with your PDU. The need from these can vary from environment to environment, and are sometimes simply a matter of what you're willing to pay, examples include remote monitoring, environmental monitoring, and switching capability. Full network monitoring allows you to be able to see the status of the PDU remotely, even down to individual port draw. Coupled with environmental monitoring, the temperature and humidity of the environment can also be monitored directly from the PDU, allowing for each different rack to be tracked independently. Switching capability allows for individual power ports to be turned off remotely. This can be very useful if there are troublesome devices that need frequent reboots, or a last-ditch recovery mechanism on an unresponsive server or switch.

Transfer Switch PDU

Another type of PDU is for devices in your environment that do not have dual power supplies. Sometimes even very important devices (routers/firewalls/appliances) don't have dual power cords to be able to connect to redundant circuits. In order to get past this there are rack-mountable transfer switches. These switches take two different input sources (like from the rack's two PDUs) and then provide failover to downstream devices. They can be 110 or 208 V and can have various different supply and output ports. Sometimes they just have one output plug as an L6, sometimes up to ten C13 and everything in between.

These types of switches are critical to supplying a redundant environment to all of your equipment, even the less redundant pieces. Don't take these as a replacement for dual power supplies though; the power supply itself is still a single point of failure, this just prevents a single circuit from taking down the equipment.

Rack mounted transfer switch with two L6-30P supplying 16 C13 and two C19 with redundant circuit power, courtesy of Newegg

Step-down Transformers

There is also a type of PDU called a step-down transformer. If the UPS supplies 208 V power, but there are devices that can use no more than 120 V, a step-down transformer can be the answer. This unit takes in 208 V and supplies out 120 V. The process of converting power uses some of the electricity, and it is given off as heat. That means that there is an inherit inefficiency to using a step-down transformer. Additionally, these aren't cheap, small or lightweight devices, so they should only be used when explicitly necessary.

Determining What Is Needed

Determining what is needed from a power standpoint is a very challenging task. It's easy to look at a server and think, "Okay, these are 1,100 W power supplies so I need 1,100 W," but that's not necessarily the case. Just because the power supplies in a server CAN do 1,100 W doesn't mean that they are going to be pulling that regularly. Power supplies of 1,100 W can be in a single-socket server with a low power CPU, two hard drives and two sticks of RAM, or it could be in a dual socket, 305 W system with 24 DIMMs, 32 hard drives and dual 300 W video cards. Both systems have the same power supplies, but only one is going to come close to actually utilizing that power supply. This also holds true for things like PoE switches, too. A PoE switch with 48 phones and wireless APs plugged in compared to a PoE switch with just some generic gigabit desktops will have a drastically different power draw.

Figuring out the draw of servers, and the relative usage of those servers is a very difficult task. If the servers in question are already running and have reasonably smart management (iDRAC, iLO, etc.) then information about average power draw can be pulled from there. Some managed switches also have information about their power usage in the CLI. Obviously, at scale, this is a very frustrating process. Many different companies have software that will go and poll your equipment for power use. For example, if you're a Dell shop, there is a piece of software called OpenManage Power Center. This software can pull the information for up to 2,000 servers, give averages, peaks, historical and individual rack consumption.

However, if software of that nature isn't available for your environment, a lot of time a sampling and then extrapolation effort works the best. This happens through a few steps:

1) Figure out the different types of servers and switches that are in the environment, like ESXi hosts, physical Oracle servers, VDI servers, backup servers, etc.
2) Validate several servers and switches manually within each category, make sure that you have picked switches and servers that are representative of each different category and measure their usage.
3) Multiply by how many switches/servers there are of each type.
4) Pad the number with a percentage to represent high load situations, this will be different from environment to environment, but you should have a reasonable idea of what your highest load times are.
5) Cross reference those results with which rack the equipment is in and figure out how much wattage each rack is going to need to function.
6) Using the above math, figure out how much amperage is needed to supply the VA needed at the supplied voltage.
7) Figure out how many circuits are going to be needed based on that amperage.

Figure out which PDUs are necessary to supply those circuits, at that wattage, and with the appropriate receptacles. Also consider what would happen if you have to turn everything on at the same time. If the entire environment goes down for some reason, the surge of turning everything on at the same time will be much higher than the sustained running of devices. Normally this shouldn't be a problem, as the whole goal is for everything to always stay online. Also keep in mind any future power savings that are going to be attributed to new, more efficient hardware.

Device-specific Voltages

Lest we forget during all of this, pay attention to the plugs that your device has. Sometimes devices are only capable of running on 208 V due to the amount of power that they need to draw. For example, look at a Dell FX2 chassis with 2,400 W power supplies. It comes with a C20 plug. We can tell by that plug choice (as well as the wattage), that it wants a higher wattage than we can supply over a normal 120 V plug. While 20 amps at 120 V is technically 2,400 W, we know that we can't run the circuit any higher than 1,920 W. That means that we couldn't supply that device on 120 V. Also, we know that the C20 is specifically made for higher amperage 208 V. So now we know that we need to be at 208 V, and we also know that we need a PDU that has C19 plugs on it.

Also, keep in mind that not ALL devices can operate at 208 V. Most can, but look at each device's input power ratings before you just plug it in or you may have smoke on your hands. Most have a rating something like 100-240 V, but sometimes that range is smaller. On older devices the rear of the power supply may have a little switch on it to switch from 110 V to 220 V. If that switch isn't flipped before the device is plugged into 208, again, smoke.

Uninterruptable Power Supplies (UPSes)

Uninterruptable Power Supplies are key to any switch closet, server room, or datacenter. They help equipment survive temporary power disturbances or serve as the short-term holdover unit before a generator takes over. Additionally, UPSes "clean up" the power that they are receiving before giving to the end devices. If the power coming in is over or under voltage, the UPS can take out the peaks and valleys and provide a predictable and steady supply on the output side. In order to do all of that, of course, they have to be able to sustain the load of their connected equipment for both short durations and for a length of time, depending on the type of disruption. This doesn't simply mean making sure that they have enough batteries to stay online for a given period of time, but also enough VA/wattage, and that the legs of output are of appropriate amperage. This makes sizing them very important, because if they are sized or configured incorrectly, in the event of a power disruption, they won't help.

Types

There are a few basic types of UPSes that all operate differently and provide different pros/cons. We'll briefly go over the high-level variations. There are a few things to keep in mind with regards to the different types of UPS, the voltage conditioning, efficiency, transfer times, and of course, costs.

Voltage conditioning is, as it sounds, how well the output voltage of the unit is conditioned before it is provided to end devices. This is basically how well the sine wave is maintained and how well spikes/brown outs are handled. Efficiency is how much power is lost in the process of getting from input to output power. For example, if 1,000 W come in, and only 800 go out, the unit is 80 percent efficient. This efficiency shows up in extra heat generated by the UPS, and (obviously) lower output loads. Transfer times refer to when the unit has to transfer from one type of power to another, and how long that takes. If a UPS has to switch frequently, and each switch takes several milliseconds, it can cause disruption to downstream devices. Costs should be self-explanatory; but in short, the nicer the power conditioning, and the better the line conditioning, the more expensive the UPS.

Standby UPS

This UPS does not in any way use the batteries unless the AC power fails. There are two distinct paths for the electricity, one that charges the battery and one that is surge suppressed, filtered and then sent to the output. There is a transfer switch at the output that chooses either the conditioned AC line or the battery power that is inverted to AC. The simplicity of this design makes it much more affordable, and the standby nature makes it very efficient as well (the only lost power is to charge the batteries).

The biggest problem with this design is the time it takes to switch from AC to battery power in the event of a disruption. Additionally, since the only two states for the UPS is on AC power or on battery, if there is a low power situation (brown out) the unit has to fully switch to battery, rather than using the battery to help shore up the low power. This makes the UPS more likely to run out of battery power in that situation than a UPS that can simply subsidize AC power with batteries. Because of these concerns, this type of UPS is often used for client computers and less sensitive equipment. A simple block diagram of this is below.

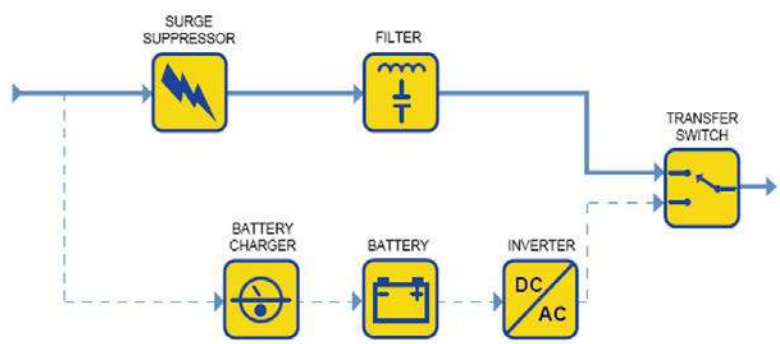

Courtesy of APC

Line Interactive UPS

The line interactive UPS uses a simple design that is very common in small to mid-businesses. The same AC/DC inverter is used for both charging the batteries and converting battery power to output, it is simply reversed if it's charging vs discharging. This means that the inverter is already energized, which means that the unit can switch over quickly. This design also helps for line conditioning of the output voltage to provide a more consistent power stream. This design is very efficient and reliable, while also providing reduced transfer times. The batteries being used regularly does mean that they will wear down more quickly than the standby design. This type of UPS is used in most UPS units of 5,000 VA or less.

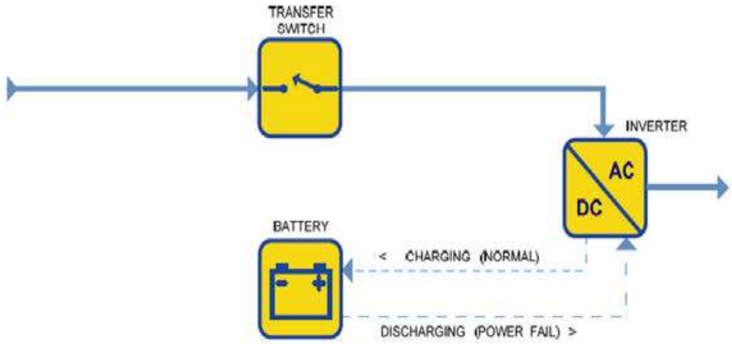

Courtesy of [APC]

Double Conversion On-Line

A double conversion UPS is similar to a standby UPS in reverse. Power is always sent to the batteries, which are always charged, and then it is always inverted out of the batteries to the output. This type of UPS is always online by default, and therefore doesn't have a transfer time for normal AC failures, only when there is a component failure in the UPS itself does a transfer have to happen to bypass mode. Because of this power conditioning is very high. The efficiency of this double conversion is low though, as power is lost on both sides of the conversion. This type of UPS is used in most UPS above 10,000 VA.

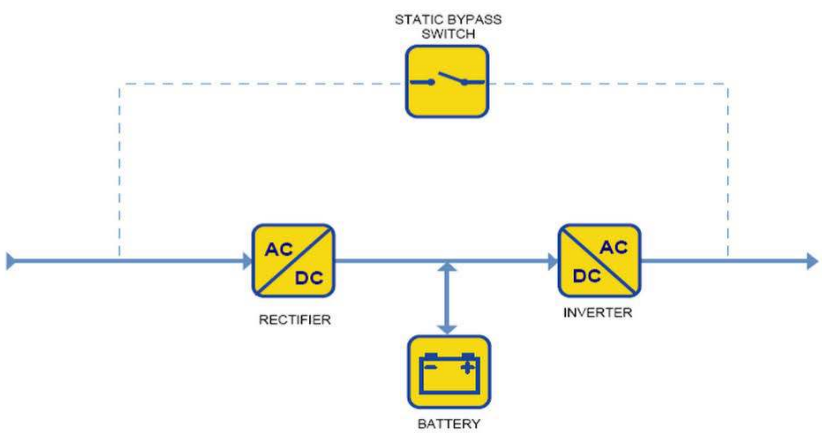

Courtesy of [APC]

Other Types of UPS

There are other types of UPS that are combinations of these styles or some form of hybrid of them. They are less common though, and follow the same basic premises of the above UPS, so further exploration of their features and functionality won't be in this document.

Sizing a UPS

Sizing a UPS can be a very complicated procedure. There are a lot of variables that go into it, both from a wattage and volt-amps side, as well as a runtime size.

Power Supply Efficiency

The rating on computer power supplies is not the wattage going INTO the power supply, but the wattage going OUT of the power supply. Have you noticed that on power supplies they often list their efficiency on the 80+ standard? These are the platinum, bronze or silver metrics that show how efficiently the power supply converts from AC to DC power. Luckily, most all of them are gold now (or platinum) which means that between 85 and 90 percent of the AC coming in goes out as DC. Normally the highest efficiency number is in the center, approximately 40 to 60 percent utilized. So now we don't just take the "1,100 W power supply," we also have to multiply it by how efficient the power supply is at that power draw.

Let's take our 1,100 W power supply unit (PSU) and 500 W PSU (output DC power):

	10% load	20% load	50% load	100% load
Power Supply Efficiency	86%	90%	90%	87%
Output Watts 1100W (DC)	110	220	550	1100
Output Watts 500W (DC)	50	100	250	500
Input Watts 1100W (AC)	128	244	611	1264
Input Watts 500W (AC)	58	111	278	575

In this table we can see the power supply efficiencies at various different loads. For example, at 20 percent load, a 1,100 W power supply gives 220 W DC, is 90 percent efficient, and therefore consumes 244 W of AC power to supply that. This is an important distinction to keep in mind, that the DC watts isn't what's actually being used by the system, rather the end result of the system consuming an (always larger) amount of AC wattage. The efficiencies of the power supplies don't change that much across the spectrum (three to four percent) but over a few hundred power supplies, the difference of having a properly sized power supply running in its sweet spot, rather than a larger one running under-utilized could potentially add to some real savings.

Let’s add one more factor into this, though: the power factor.

Power Factor

One of the most important parts of UPS sizing is the power factor. The power factor is essentially a relationship between the watts and volt-amps, or how the device gets the watts in question. Power factor is normally expressed as a number (0.65) or a percentage (65%). The good news is that most new server room equipment (in the last two decades) use a “power factor corrected” power supply, which means the power factor is basically one. So, for pretty much all new server equipment you can assume that wattage and volt-amps are one in the same … when they are operating at 100 percent load, which is what the measurements are made from. However, when those same PSUs are under lighter loads, they get a lower power factor. For example, while the power factor may be .99 at 100 percent output, it may only be .96 at 50 percent utilization, or .71 at 10 percent utilization. You’ll notice earlier, the power supply is most efficient converting from DC to AC in the middle of its utilization, and as mentioned above, the power factor is best at full utilization. Luckily, to go from wattage drawn by the server to VA we simply cross multiply and divide with the power factor. So, looking at 278 W of AC power with a 96 percent power factor, we get 289.6 VA used.

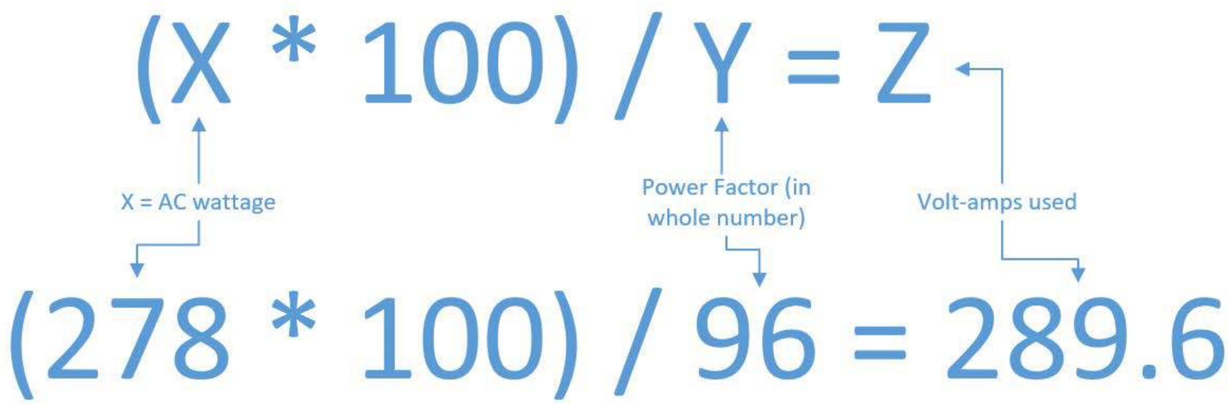

Let’s look at everything again, this time adding in power factor.

	10% load	20% load	50% load	100% load
Power Factor	71%	86%	96%	99%
Input Watts 1,100 W (AC)	128	244	611	1,264
Input Watts 500 W (AC)	58	111	278	575
VA consumed 1,100 W	180	284	637	1,277
VA consumed 500 W	82	129	289	581

Now we can see the VA numbers at each of these power supply wattages. That lets us clearly see the gap between volt-amps and AC wattage on some of these and how inefficient the overall power draw of the power supplies are at certain utilizations. Further, we can also see where power supplies are the least efficient in terms of DC output and volt amps consumed.

	10% load	20% load	50% load	100% load
Output Watts 1,100 W (DC)	110	220	550	1,100
Output Watts 500 W (DC)	50	100	250	500
VA consumed 500 W	82	129	289	581
VA consumed 1,100 W	180	284	637	1,277
VA efficiency	61%	77%	86%	86%

So we can see that less than 50 percent load gives us an efficient (87-90 percent) AC – DC conversion, and also the best power factor rating (96-99 percent), which gives us a much more efficient overall utilization at those loads on both the 500 and the 1,000 W power supplies. The big difference is how much wattage is given out at those respective loads.

In graph form:

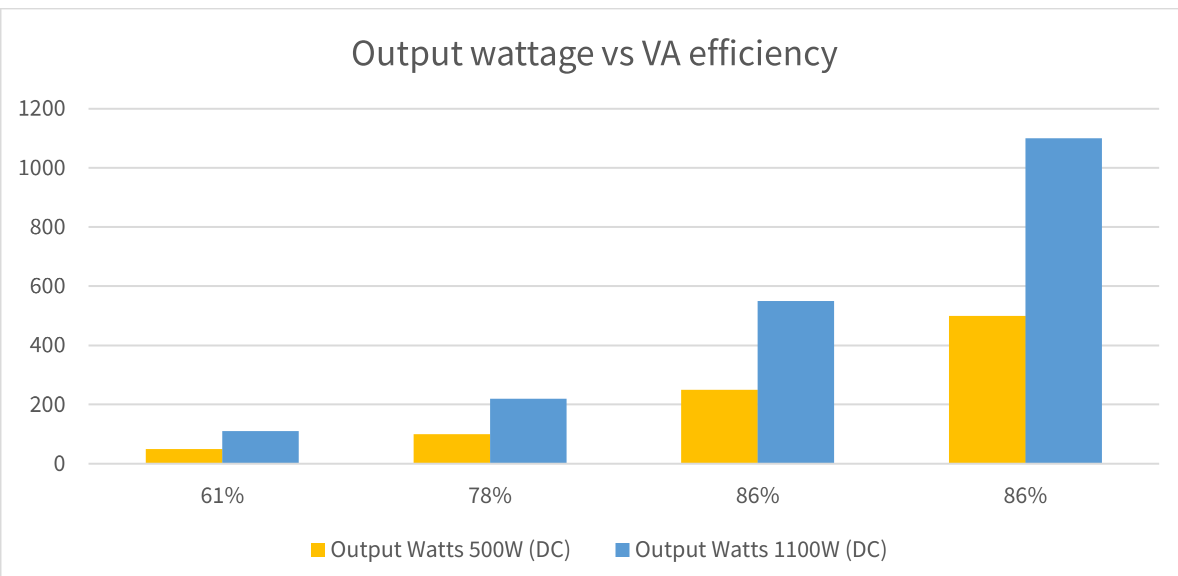

Given an equally efficient power supply, we can see that the efficiency of the same output wattage changes drastically between the two power supplies. For example, at 110 W used, the 500 W power supply is 17 percent overall more efficient on VA draw than is the 1100 W. This is important, because it means that if you've oversized your servers by spec'ing them all with 1,100 W PSU when they only need 500 W, then your power supply efficiency may be low. When coupled with the power factor, that may mean that the power supply draws way more VA than is necessary, putting more relative load on the circuits, and the UPS, and is also producing more heat than a more efficient power supply. Remember the earlier discussion about 208 V being more efficient than 120 V? Well, there's more to that, because 208 V also makes these same power supplies operate two to three percent more efficiently than 120 V normally. A PSU designed exclusively for a 208 to 240 V power supply can be even more efficient than that.

Redundant Power Supplies

Well that was tiring, but there's more to think about! This is all assuming that the server only has one power supply. If my server has two power supplies, what does that do to my power factor and efficiency? Well, if your power supplies are both in active-active mode, whatever the load is gets divided in two. Even if you have a 50 percent power supply load, which should give you a power factor of .96 and an efficiency of 90 percent, if you have two power supplies and distribute the load between them, each is now down to only 25 percent load, which means that the power factor is down to .86 rather than .96 but the efficiency stays around 90 percent. This is the main reason that many server manufacturers no longer run their power supplies in active-active, but rather in active-standby. That way power supply one operates at higher efficiency and power factor up until the point that it dies, then the other power supply immediately takes over, also at the higher efficiency. It keeps the PSU in its most efficient spot, where it also functions best.

Desktops

Okay, now we've gone pretty deep into power factors and efficiencies on server room equipment, but what about consumer equipment? We regularly need to put a UPS on desktops too. Consumer equipment often still doesn't run a power factor corrected power supply, unless it is Energy Star 4.0 certified. Those power supplies have a power factor as low as 65 percent. That means if you look at a 400 W desktop computer, the VA rating of it could be 615. Because of this, UPS manufacturers design their low-end UPSes accommodating for this load power factor disparity. For example, if you look at an APC Back-UPS 500, the 500 is the VA rating, but if you look at the wattage rating of the system, it's only 300 W. Say you had a new Energy Star 4.0 certified desktop and plugged it into this UPS that had a 400 W power supply. The efficient desktop probably has a power factor of near one, so its wattage equals its VA. However, while this UPS has enough VA, it doesn't have enough wattage to run that desktop.

The Core of VA Sizing

You can never exceed either the wattage or the VA on a UPS, because it will cause it to overload. You may be currently thinking, "Holy crap, this has gotten complicated!" But the best way to look at it is a minor oversizing. If after all my math and sizing of my server I come to a conclusion that I need 7,000 W, how do I figure out what size VA UPS I need? Well, the safest path is to simply multiply it by a reasonable worse case VA efficiency. For example, if all of the equipment in your datacenter is similar and you can pull the power factor of it (say for example, like the ones on the previous page) we know it's least efficient load is a VA efficiency of 61 percent. However, we know our equipment is probably running higher than 10 percent all the time, so we can step up to 20 percent or more percentage utilized, and a 78 percent or more VA efficiency. If we cross multiply and divide, we find that at 78 percent VA efficiency, we would need 8,974 VA to run the load. Therefore, a 10,000 VA would be the lowest possible UPS we would want. If all of our equipment is running at a higher efficiency, this simply means we have more run time, more head room for expansion. Of course, in real life, we wouldn't size the UPS so close to the load. We'll pad that number for future growth and some less efficient equipment to go up to something maybe like a 12,000 VA.

Runtime

Runtime is really where the rubber meets the road on UPSes. It doesn't do me any good if the UPS can hold my entire load but only for one minute. Most UPSes can add additional batteries on as a way of combatting the runtime problem. One would like to think that the more batteries you put on the UPS, the longer the runtime and it was as simple as that, but as we've found out, with all things there is complexity. Notably it comes to (once again) efficiency.

A certain amount of efficiency is always lost when a UPS converts from DC to AC. It's the same reason a 12 V DC inverter in a car gets warm when you use it a lot. Some electricity is always going to be lost. Unfortunately, this is not a linear amount, there's a certain amount of overhead regardless of the load on the UPS, which means the actual efficiency varies greatly depending on load. For example, here's an efficiency graph of an APC 8,000 VA UPS:

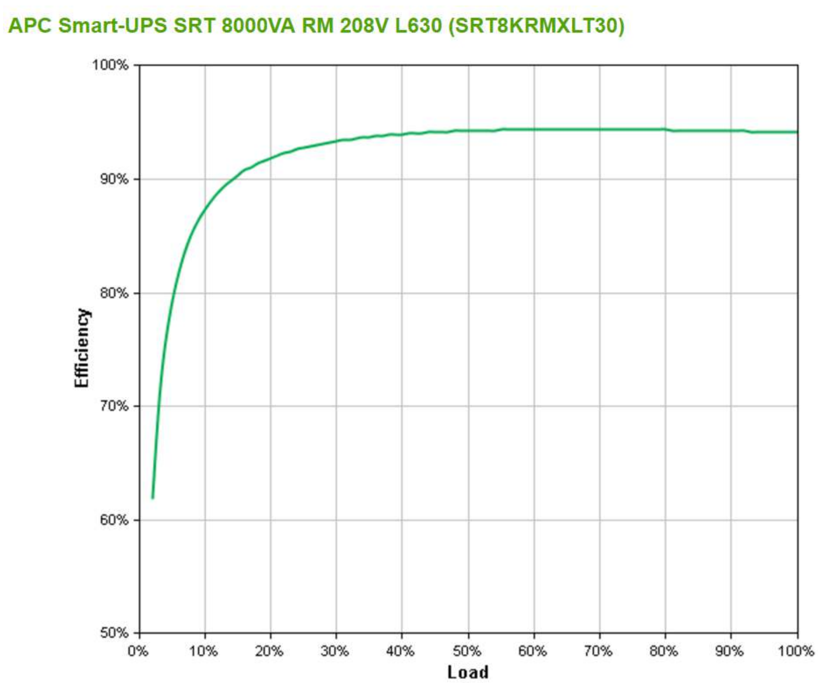

Courtesy of APC

We can see that as long as the UPS is at 25 percent or more utilization, it’s 92 to 94 percent efficient. Lower than that, though, and the overhead of the DC/AC conversion puts a real tax on the efficiency. Similarly, other UPS can have efficiency lower than 90 percent even at a good load. It’s important to look this up when comparing UPS to validate that the two are similar. If we look at a runtime chart of this UPS, we can see what the runtime is with varying numbers of battery packs on it. ‘A’ has none, ‘I’ has 10 extra.

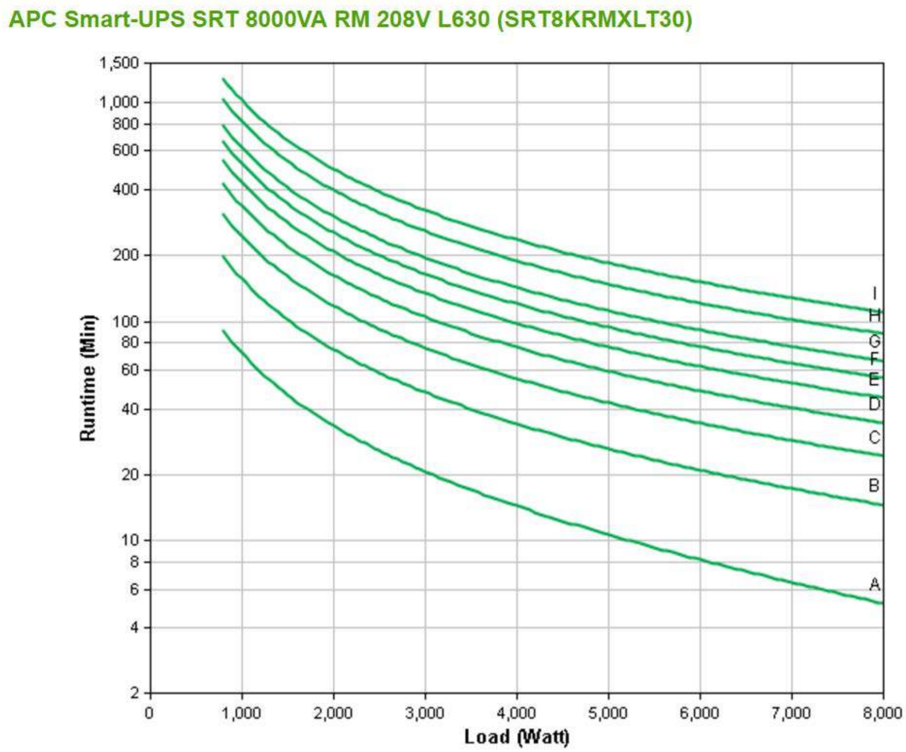

Courtesy of [APC]

It's hard to see in the graduated chart, but each battery actually adds more run time than the previous, but keeps the same curve. Which looks like this in graph form at maximum (8,000 W) load:

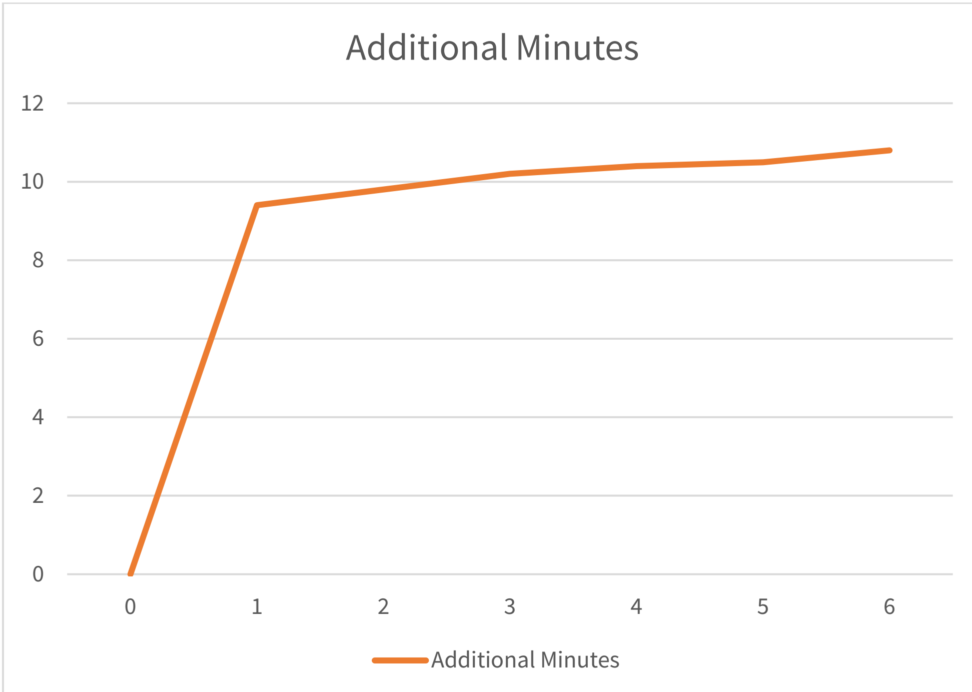

Adding one battery adds 9.4 minutes of run time, but add another to that and it's an additional 9.8. It's not a linear growth, with each additional battery actually adding more runtime than the previous. The difference from one to six batteries is 1.4 minutes, which doesn't sound like much, but it's a difference of 15 percent, which could be 20 or 30 minutes depending on how many batteries and how large of a load you have.

Unfortunately, this means we can tell that it's not simply a matter of multiplying by how many batteries -- there is more to the equation. While there are several different math equations for predicting this number on the internet, none of them seemed to work in practice. Their results didn't line up either with APC's quoted runtimes, or with real world experience with this UPS with similar loads. The recommendation for calculating runtime must be found on the manufacturer's website and based off of their calculations, using the load you previously calculated.

So how long is long enough? Once again, it depends. Does the office have a lot of short outages but rarely long ones? Then a short runtime is probably acceptable. Do you have a generator? A lot of companies use UPSes simply as a hold-over until the generator can kick in, and not with the intention of running the environment off the UPS for any real length of time. Or do you need it to last long enough to gracefully shut down all your VMs and get the environment into a stable state before power loss? All of these factors change how long the UPS really needs to run. If you want something for HOURS of runtime, you're much better off having a lower run time UPS and then a generator backing it. However, that is sometimes too cost prohibitive, and larger UPSes have to be bought to keep the environment going for the whole duration.

Circuits and Receptacles

As with the PDUs, the circuits on the UPSes are very important. The type and layout of circuits strongly influences how the PDUs will connect into the UPS. While most small (desktop) grade UPSes have a single circuit on them, most larger UPSes have multiple circuits. To start, what's a circuit? With regards to a UPS, a circuit is a series of electrical plugs that are controlled by a single breaker. Each breaker can supply a certain number of amps, normally between 15 and 30. When it exceeds that, the breaker will flip, to prevent damage to the UPS, or a potential fire. A circuit is the fault domain within the UPS. If that breaker gets flipped, everything that is downstream of it loses power.

This is one of the reasons that it's very important to pay careful attention to the circuit layout on the back of a UPS when PDUs are plugged into the UPS and then equipment is plugged into those PDUs. Circuits will often have three to six plugs on them on the UPS. If two PDUs are plugged into the same circuit, and then the downstream server has both of its power supplies plugged into those PDUs, it's no better than if they were in the same PDU. If the circuit fails, both go down, as seen below.

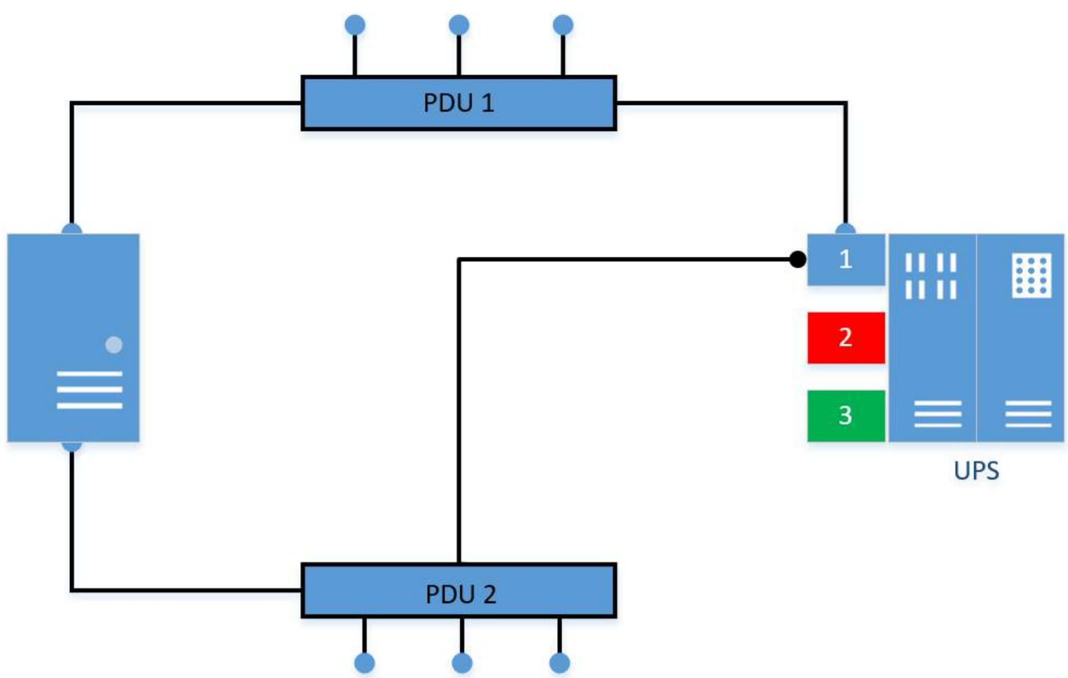

Each circuit also supplies its power in one voltage, either 110 or 208 V, as well as a maximum amperage that it can supply before the breaker is flipped. Depending on the UPS, this could be as low as 15 amps on 110 V or as high as 30 amps on 208 V.

Yes, technically, it could be an even higher voltage and amperage, but at that point it's a high-end, datacenter-grade UPS, not a normal server room unit, which is outside the scope of this document. It's also possible that the protected side of a UPS is hard wired straight into a wall panel (circuit breaker box), but again, those are very large, datacenter grade UPSes, above the scope of this document. If you have questions relating to this type of configuration, you can reach out to us to discuss it.

Due to these circuit layouts, there are several considerations when a new UPS is being acquired:

1) How many receptacles are on the back of the UPS?
   Are there enough to plug in all of the downstream PDUs (or direct servers)?
2) What type of receptacles are on the back of the UPS?
   Are they the same plug as your PDUs have? As we covered in the first part of this document, there are a LOT of different options, and depending on the UPS, you could have any of them. In NEMA: 5-15R, 5-20R, L5-30R, L6-20R, L6-30R, or various other IEC plugs, normally either a C13 or C19. Make sure that the UPS has enough (and the right type) of plugs to accommodate what you need to plug into it.
3) What circuit are the plugs on?
   People often pay attention to the available plugs on a UPS but forget to pay attention to WHERE those plugs are. For example, the SRT3000RMXLA from APC, the UPS has three circuits and has 5-20R on it, and you may assume that if you buy three PDU with 5-20P you're set. In reality, however, that UPS has two 20 amp circuits, each with four 5-20R, but the third receptacle is a L5-30R. So now across the three PDU, two would have to share the same circuit, and it would be a lower amperage circuit. With the largest amperage circuit (the L5-30R) being left unused completely.
4) Can a single circuit support its necessary PDUs?
   Always keep in mind what the downstream draw of a circuit is. Remember that circuits are limited on how much power they can consistently supply. That means if eight downstream devices were plugged into a 15 amp circuit each pulling two amps, the circuit were to be overloaded, which would cause the breaker to trip, and the whole circuit would go offline.
5) What is the UPS's total amperage available?
   Just because the UPS has the three circuits above (20, 20 and 30) doesn't necessarily mean that it can supply 70 amps. Even the NEC deregulated rating of 56 amps is still way over what it can actually supply. Looking at the technical specifications for the UPS, we see that it supplies 3,000 VA at 120 V, which mean the actual output amperage can't exceed 25 amps because we know VA/volts gives us amperage (3,000 VA/120 V). So if that L5-30R is heavily used, the 5-20R circuits are unusable. Similarly, if both 5-20R have a decent load on them, the L5-30R is unusable. So why are there so many circuits? Well, the L5-30R and L5-20R are to have different options for connectivity, it would then assume that EITHER the L5-30R or the L5-20R are used, not all of them at the same time.

Care and Feeding

A lot of problems with UPSes come from a lack of proper sizing. Overloaded circuits or lack of enough batteries to keep the equipment online for an appropriate duration. However, the vast majority of UPS problems come from simply not taking care of the UPS itself. UPS units require regular (if infrequent) care and feeding to make sure that they survive properly. Most UPSes run a self-test automatically, and can even alert you when it runs and how well it went. They can also do battery relearn cycles and other things to keep themselves working well. However, there's one thing they can't handle themselves, over time batteries just go bad.

Batteries

Batteries operate through a chemical reaction. That reaction can only happen so many times before the components that make it up have to be replaced. It's the same reason that your cell phone or laptop's battery gets worse and worse the longer you have it, or that the battery on RAID cards must be replaced. The proactive things the UPS can do itself to keep the batteries healthy (as listed above) only go so far.

Most manufacturers often provide a three-year warranty on the batteries in their UPSes, and that three year number isn't a fluke. Batteries normally only last for three to four years when in constant use before they stop effectively holding a charge. In reality, what that looks like is the batteries will say that they're fully charged, but as soon as they take load, they'll very quickly go from something like 92 percent to five. The UPS often won't even know the batteries aren't healthy and will THINK it has 45 minutes left, and then all of a sudden it will just shut off. This is why this is so dangerous. If the batteries aren't proactively replaced every three years or so, the issue WILL happen, and you'll be completely unprepared.

Management, Monitoring and Alerting

It's great to have a good UPS that can run all of your needs, but if the UPS doesn't get monitored and/or chatty when it has a problem, it isn't going to truly live up to its potential within your organization. Modern UPS units and even PDUs have a load of networked information available, including status, alerts, temperature, humidity, etc.

Remote Monitoring and Alerting

One of the simplest tasks that needs to happen with a UPS is the ability to remotely see their current status and be alerted to changing conditions. As with everything else, if this isn't set up, then it can't work. Different UPS models come with/without networking modules in them, and sometimes on low-end UPSes the cost of that network module may be nearly the same as the UPS itself.

It's very important to include this module, as it's what allows the unit to be able to alert the world to its current status as well as initiate remote (or USB) shutdowns if needed. Additionally, higher-end PDUs are also now network accessible. These PDUs can be plugged into the network in the same manner that the UPS is so that load, internal errors and other things can be centrally monitored.

The modern UPS also is fully web enabled, so simply browsing to its IP address lets you see the current status of it and how well it's functioning. These UPSes can be configured to send emails/SNMP alerts for various things as well: on battery, self-tests, internal temperature problems, battery faults, circuit load concerns etc. Additionally, there are additional UPS software packages that can be installed that provide the ability to see your entire UPS/PDU infrastructure in a single pane of glass.

The modern UPS/PDU doesn't simply stop at monitoring itself -- it can be a great tool to monitor the related environmentals. Common options are things such as temperature and humidity probes. These allow for the UPS/PDU to also provide alerting for things that normally directly impact the ability of the UPS to function, an overheating room or a water leak. It's a very simple way to "kill two birds with one stone" and put multiple monitoring options under one roof. Some UPSes also offer the ability to check the status of any additional power sources (for example, another plug in the same room). This can be useful if, for example, you want to see if the whole building lost power, or just the circuit that the UPS was plugged into.

It's also important to make sure that every piece of the environment that is responsible for getting these alerts/monitors to your admins has battery backed power. Keeping your servers and switches with backup power doesn't help if the internet handoff device from the carrier goes offline. No one will get the alerts and the battery will run out and let everything crash before you ever know there's a problem … unless you configure server shutdown.

Server Shutdown

Monitoring your environment to detect problems is great. Being alerted is even better. However, if there isn't a system configured (and tested) to actually shut down servers, then the odds of a hard outage at some point become much more likely. Many people think that with, for example, a one hour run time on their UPS that they have plenty of time to log in and manually shut down the servers should something go wrong. But what if it happens when you're out to dinner? It's important to have server shutdown pre-configured.

Most modern UPSes come with free software to shut down your environment. It's important to make sure that it's aware of the type of environment you have, and covers all the types of servers in your environment. If you have physical servers versus VMware or Hyper-V, they all have different types of shutdown methods. Many types of this software will allow you to set priority shutdown orders, so that VMs are taken down in the proper order to help with dependencies. Sometimes, you just have to make a script for shutting down parts of your environment. Making sure that everything goes down gracefully is your first priority. The second priority is setting WHEN they go down …

Avoid the temptation to set the shutdown to kick off as soon as the UPS goes on battery power. Most power issues tend to be transient and only last a few seconds to a couple minutes. If the power comes back online too quickly, the servers will already be in the process of shutting down and the outage from automatic shutdown will be much worse than the actual power event. The other end of the spectrum can also be bad; if the shutdown doesn't start until the UPS reads "batteries low" when the shutdowns start the increased load of everything shutting down may lead the UPS to a premature depletion and take the environment down before everything is cleanly down.

Of course, the proper plan is to test how long it takes to shut your environment down gracefully, and then plan the shutdown time on UPS to accommodate. This is also useful knowledge when you are learning how much runtime you need on your UPS. It needs to be able to sustain a few minutes of outage, and then sustain the environment while it gracefully shuts down. Having a generator in the building is no excuse for not having this configured. If the generator and UPS are properly configured, it should come online and supply power before the UPS would kick the shutdown off. If for some reason the UPS gets to a low battery state, it probably means the generator isn't coming online for some reason (bad fuel, 99 percent of the time) and therefore servers need to go down gracefully anyway.

This is an important part of sizing a UPS, they need not only to hold the power while the generator kicks in, but also be able to last through a shutdown in case the generator DOESN'T kick in.

Cooling

Cooling is frequently a pain in IT's side. Server rooms and datacenters are extremely dense from a BTU-per-square-foot perspective. A LOT of heat comes out and it comes out in a very small area. Getting proper cooling laid out is, in itself a very large topic. For the purposes of this document however, keeping that cooling online is the main premise. Again, if you'd like to discuss the design of your server room or datacenter cooling system, reach out to us.

A very frequent problem with power design/implementation is not having a way to keep the air conditioning running when the power goes out. Keeping the servers running off of a UPS for an hour after the main building goes dark is a great thing, but if the server room gets to 120 degrees during that time, then the servers are going to go down anyway for a different and potentially more stressful reason. This is something else that has to be played into the automatic shutdown procedure, how long can things keep running in the server room without cooling before they have to be shut down? Little transient power outages may not be a big problem as the AC will be back online before the temperature rises too much, but depending on the load and size of the server room, even 10 minutes without cooling can be enough to cause a server room to overheat and the servers to shut down.

The seemingly natural solution of putting the air conditioner on the UPS is not often not a great idea. Excluding the fact that air conditioners pull a lot of power (between three and four thousand watts wouldn't be uncommon), the way that the air conditioner uses power (inrush load on an inductive motor) it can take significantly more power to get it started up than it does to keep it running. This means that the air conditioner can create some very dirty electric loads on the UPS itself, which leads to having to massively oversize the UPS to accommodate for an air conditioner. An example of this is watch the lights in your house when the AC kicks on. When the AC first starts up the lights dim for a bit and then they go back to full brightness. This happens even on your house which is going directly to the power grid and it's all due to the immediate current draw it takes to get the AC running.

Because of this heavy and sudden power draw of the AC, it's best not to put it directly on a UPS. For most people this isn't the worst situation. Most UPSes aren't planned to run for more than 15 to 30 minutes, and somewhere during that span the servers are going to be shutting down. Normally, unless it's a very small server room with a very large amount of equipment, the ambient air can absorb the additional heat load enough to keep things functional until they're shut down. Where this really becomes a problem is on extended run times.

It's impractical to expect most server rooms/datacenters to run for more than 30 minutes without any form of cooling, but it's also impractical to put air conditioning on the UPS. What are we to do?!

Generators

Generators are outside the scope of this document (and most IT organizations), but to try to go full circle on things that need to be accommodated, let's talk about them briefly. For starters, why would we have a generator vs a BIG UPS? It really comes down to the intended use. UPS units aren't horribly efficient, they are large, they use electricity themselves, they have to be cooled, and their batteries need to be replaced way more often than most people do. So why do we use UPSes?

Modern UPSes give very clean power and they are very responsive to failover. Most modern UPSes take effect within a few milliseconds of an event. Modern generators can provide relatively clean power, although not as clean as a UPS's. Additionally, the generator must engage, cycle on and then start providing power, which can take a few minutes to fully come online and be ready to go. In IT, the world has already ended in that timespan, and we're going to have to spend (potentially) hours powering everything back up again. However once that generator is online, it can run (essentially) forever. Keep feeding it fuel and it will keep pumping out electricity. Speaking of fuel …

Fuel Types

The type of fuel that the generator uses is also actually a very important consideration. Most everyone defaults to diesel, but that may or may not be the best option for a few reasons. Diesel fuel is NOT clean fuel. When it burns it lets off a lot of nasty emissions, these are actually regulated by state/federal laws depending on where you are located. Further, diesel has a self-contained tank of fuel that needs refueling. Depending on what kind of disaster has happened, are you going to be able to get new fuel to it to keep it running? That fuel that it's storing can also get old, or too cold and gel, and essentially become useless unless special additives are put in it. Ever seen an old diesel truck try to start when it's really cold outside? Your generator may suffer the same struggles.

Propane is another option. It has a long shelf life, burns cleaner, is normally in lower demand in disasters, and engines start well in cold weather. Generally speaking, propane may need slightly less maintenance than diesel engines. Of course, it also has a more complicated fuel system since it's a pressurized liquid, can be more expensive to install/operate, and may have a shorter all around life (if actually running).

Natural gas is another option. Natural gas also burns cleaner than diesel, starts well in cold weather, and is normally directly plumbed into a fuel supply. That means that there isn't a worry of retrieving more like with diesel or propane. On the other side, an earthquake may break that direct plumbing and cause a fuel shortage. It also is only readily available in certain areas on the country, and can be more expensive to install/operate and may have a shorter life expectancy (if actually running).

So which one do you use? Well, it depends on emission regulations, part of the country you're in, and the types of disasters you are likely to have happen. If you have it available, are in a colder climate and you aren't prone to earthquakes, natural gas may be a good option. If you want the cheapest option and don't mind the regular care and feeding of the fuel, diesel may be your winner. You have to weigh the pros and cons. Regardless of which system you have though, there is one thing that is always true …

Testing

Conservatively, probably half of all problems and failures with generators come from a lack of regularly scheduled testing. Generators, much like disaster recovery plans/sites are something that you hope you never (ever) have to use. With that, it's easy to ignore and/or forget about the maintenance you should be doing to them. Just like a DR plan however, when you need a generator to work, you REALLY need it to work. Generators are complicated pieces of equipment.

Essentially you have a truck motor with all the constituent parts (oil/fuel/air filters, radiators, exhaust systems, fuel supplies, pumps etc.) coupled to a really large alternator, which then goes to a voltage converter/regulator. These parts need to be regularly tested, generally monthly, but check with your manufacturer for their recommendations. Also, pay attention to their recommended maintenance schedules for generic service. Further the test needs to run for a substantial amount of time to validate that it is working as expected and have the load actually applied to the generator to make sure it can sustain it as expected. These tests are very difficult, without a special type of switch to move the load, a transfer switch.

Transfer Switches

Having the generator able to come online is great, but it doesn't give you much if someone has to re-wire the whole datacenter to plug things into the generator. This is where an automatic transfer switch comes in. This is essentially just a big switch that switches the power feed from grid power to the generator. The UPS will hold the load during the spin up and switch to the generator, at which point the generator starts charging the UPS again for the switch back. These switches, while they sound very simple, are also pretty complicated in practice. They are necessary, though, and properly getting them installed and configure for failover is very important.

Sizing

Sizing a generator isn't as simple as sizing a UPS. "Wait, what? Sizing a UPS was simple?! There were seven pages on it!" A generator needs to be able to power the UPS (including its inherent efficiency shortcomings), as well as the ancillary things that need to stay online. We mentioned earlier that the air conditioning probably wouldn't be hanging off of the UPS but rather would be on the generator itself. That heavy load previously mentioned needs to be accounted for on the generator's load, not to mention the large surge that the AC imposes. Some subset of lights are often also on the generator, as well as possibly some emergency systems (fire alarms, security systems, etc.) that may not be on the server room's UPS, but on small UPS throughout the building.

Conclusion

Power is the ONE thing that IT simply can't function without. Nothing that IT does can work without electricity of some type, and yet it's often the one thing that IT administrators and managers completely and utterly neglect. Like most things in IT, power has a complete lifecycle, including proper sizing, preparation and architecture all the way through care and feeding and then proactive maintenance. Trying to keep track of it as someone not used to electricity can be a real chore, but a vital one that literally keeps the lights on. Hopefully this journey through all things power has helped to clarify and explain some of the most common power problems and concepts for future use.

About Mirazon

Mirazon is an American IT consulting firm with offices in Louisville, Kentucky and Kansas City, Kansas. Originally founded as a Microsoft services provider, Mirazon has expanded into value-added reselling and managed services to round out its IT consulting portfolio. The engineers at Mirazon regularly design, maintain, troubleshoot and update networks, storage, Wi-Fi, cloud environments and backup and business continuity systems.

If you have further questions about anything mentioned above, you may contact Mirazon by

- Calling: 502-240-0404
- Emailing: info@mirazon.com
- Or visiting mirazon.com