www.spkaa.com
Ph: 888-310-4540

SPK and Associates
900 E Hamilton Ave, Ste.100
Campbell, CA 95008

Datacenter Power Plan for a new site

Before anything, determine the service level agreement you need to meet. Most datacenters fall into three categories:

1. HA – typically a Collocation facility. This is where you need fully redundant N+1 power, network, and cooling for applications that are always up.
2. Corporate critical daytime but available for scheduled maintenance evenings or weekends. Some potential outages may occur for unexpected hardware failures but critical infrastructure is redundant.
3. Engineering datacenter. Desired to be up but no one is willing to pay to create a higher availability with redundant architecture.

In this case, we'll focus on #2. It fits the majority of customers and we can discuss how it differs from #1 and #3 without going into a lot of detail.

Generation:

This starts at the generation. Are you on a power grid that's reliable? Are you subject to blackouts? Speak with your local utility for historic outages. Minimally, a Uninterruptable power supply (UPS) is needed for critical equipment. If you have outages that last over 20 minutes, a generator is recommended. If a company is over 200 people, Generators are easily justified based on lost productivity cost and hardware failure.

UPS or Generator:

What size UPS or generator? A typical power budget for one rack is 6000 watts. This is roughly three 115v circuits at 18 amps/each. Electricians limit it to 16 amps each circuit but the breaker will not typically blow until 19-20amps. For a datacenter with 10 racks, this is 60kilowatts. UPS are typically rated in Volt-Amps. Kilowatts and Volt-Amps are somewhat interchangeable for the lay person however there is a difference. If you want to read about the difference, here more information:

http://thethriftygeek.com/technical-articles/watts-vs-volt-amperes/

With UPS and generators, efficiency is gained with size. It is better to buy one larger UPS/Generator or generator than many smaller units. Maintenance costs also are cheaper. However, there is an exception. Collocations typically have N+1 generators. This means they have at least one spare in the case of for all circuits. For mission critical corporate datacenters, two medium UPS/Generators may be used rather than one larger UPS to provide added redundancy. For engineering datacenters, a small UPS for critical machines may be used.

Power Distribution:

Power distribution to the rack is important. Each circuit should have a digital amperage indicator. If you have to ask if' it's worth the money, you've not seen an entire rack go down after accidentally overloading a circuit. Machines that you never planned to power off stopped. No shutdown so OS and applications are in an unknown state and possibly broken.

www.spkaa.com
Ph: 888-310-4540

SPK and Associates
900 E Hamilton Ave, Ste.100
Campbell, CA 95008

Here's a PDU company that I've used reliably:

http://www.tripplite.com/en/products/power-distribution-units.cfm?gclid=CJ_h5e6oqKwCFUoaQgodLEJpCA

Which PDU? It depends. Most like the zero U PDUs so they don't waste rack space for power. I strongly recommend the L5-20P or L5-30p twist lock plugs. The household NEMA 5-15p/NEMA 5-20p is not recommended. The main reason – accidental removal. They're blades with no locking mechanism. L5-30 twist locks will optimize circuit breakers within the power panel to maximize utilization of the panel.

Rack power management:

Within the rack, vendors document power draw. Remember that steady state power and startup power is different. Some PDUs allow staggered start of equipment to allow for staggering the initial surge of power required when spinning up disk drives. If you don't have this, keep to 15-16amps/circuit or lower. An extremely rough rule of thumb for servers is 1.5 amps per rack U. Here's Dells method for documenting power within a rack. One other feature that's nice is "delayed start":

http://www.dell.com/content/topics/topic.aspx/global/products/pedge/topics/en/config_calculator?c=us&l=en

Server:

On the server, dual power is desired for all non-redundant corporate servers. One of the most common failure points in a machine is the power. A $50-$75 2nd power supply will save a lot of time and aggravation without a lot of cost. If the server's are clustered or otherwise redundant, then this isn't necessary. At that point, lowest price with N+1 servers is the goal.