Introduction

Specifying a power supply for use in a particular application is often thought of as a relatively simple job, and usually left until the very end of a project. This often leads to a hasty decision and a compromise on a unit which does not fit the application perfectly, but which is available off-the- shelf and can be in house by Friday. Typically, such a unit winds up being over specified, but the justification is usually thrown out as "this will buy us reliability". Unfortunately, this is not always the case. While running a power supply at reduced load will usually allow it to run cooler, there are many components which do not benefit from this fact, and the way the supply is mounted and cooled can often make a bigger difference. **Ideally, your potential power supply manufacturer should be consulted as early as possible in the design cycle to allow him to offer suggestions**. Very often a 200 watt multi-output power supply must be very tightly regulated because of a single section of circuitry which needs it, adding about U.S.$ 10.00 to the cost of the power supply, while the problem could have been solved by a local voltage reference costing U.S. $ 1.00.

When comparing supplies, you should have a specification written which defines exactly what you need. Never "**ASSUME"** a supply will operate under certain conditions, make sure you specify it. The following items are specifications which are commonly called for in a power supply.

Input Voltage

The input voltage of a power supply should be specified for worst case operating conditions based on the system requirements and the expected operating environment. A range of 90 to 270 is typical. In a very cost sensitive, high volume application, a narrower range such as 170 to 260 may save some money. In a critical industrial control application, a wider range such as 90 to 270 may be appropriate where surges and dips caused by heavy loads can cause problems. If the application environment will be varied, a wide and continuous range input may be appropriate, such as 90 to 270. Input frequency is usually not an issue with switching power supplies, but it is with a supply which uses an off-line transformer. Some switchers, for instance, use a 50 Hz transformer for the housekeeping supply, which can preclude 400 Hz operation.

DC/DC converters are usually specified in a similar manner. Input voltage can either be from a single DC input voltage, allowing the design to be optimized at that voltage, or a wider input range if the converter must run off of an unregulated DC supply. If the converter is to run off of batteries, the input range must include the highest float voltage as well as the lowest full load voltage. This range is often 2:1 or greater. For instance, in a typical 12 volt automotive electrical system, the float voltage while driving may be as high as 15 volts, while the voltage while cranking a cold engine may be 6 volts or less. Also, in a battery system with more than one load on the battery, relatively large input filters will usually be required. Contrary to popular belief, batteries are not particularly well regulated voltage sources, and their equivalent series resistance goes up dramatically as the battery is discharged.

Surge current is sometimes specified on larger power supplies, usually on line operated supplies of 300 watts or higher. The intent here is to limit surge currents to prevent false tripping of circuit breakers or interference with other supplies on the same line. On power up, supplies with large input filter capacitors can draw peak currents of over 10 times the steady state current. However, it should be remembered that many other loads have similar characteristics. For instance, light bulbs typically draw a surge current of 10 times average current when turned on. If the environment requires it, specify it. If not, don't. A typical specification for limited surge current is 2 to 4 times the steady state current. This usually means that either a thermistor is used in series with the input, or a current limiting resistor is used with a bypass switch, or the input filter inductance can be made large enough to limit surge currents.

Hold-up time is another important input parameter. It is usually specified in cases where the load is a digital system which will be interrupted by a momentary power glitch. Studies have shown that over 95% of all power failures last no more than 1 to 5 AC cycles. This means that by specifying a holdup time of 16 to 80 ms, your system can ride through the vast majority of power failures. Most power supplies have hold-up times specified at 20 mS. For critical applications, specifying a long hold-up time of 100 mS can usually eliminate the need to consider a UPS for the system. The cost impact for such long hold-up times is usually about 10% or less, although it can be impractical on larger supplies. This approach usually has the added benefit of dramatically improving overall MTBF, since the only extra components are a few extra capacitors as opposed to an entire UPS system.

Efficiency is usually specified for one of 3 reasons: Either the supply will use enough power that the cost of the power compared to the cost of the supply is worth considering, or the input power is limited, as when a DC/DC converter is operating from a battery, or if the power dissipated by the supply will be difficult to dissipate. A typical power supply has had a fair amount of optimization done in it's design, and so there is usually not a lot that can be done to improve efficiency. However, in cases where the efficiency is a very important consideration, it can usually be raised. The penalty to be paid is primarily in cost, and to a lesser extent in size.

Outputs

The outputs of a power supply must be specified for the particular type of load. A 12 volt output used to power analog circuitry will have different requirements than a 12 volt output used to power a stepper motor. Also the type of output must be specified. Most power supplies are constant voltage, current limited types, which means that the output maintains a constant output voltage until the load current reaches the limit point, after which the voltage begins to drop in order to maintain a constant output current. However there are applications where the supply must deliver a constant current regardless of load resistance or output voltage. Some typical applications are laser supplies, charging systems for batteries or capacitor banks, and motor controls where torque needs to be controlled. In fact, for some applications, a constant power characteristic is required, for powering filament loads or similar applications.

Output voltage should be specified with a tolerance. If it may need to be adjusted with the final system, this field adjustability must be specified. If the currents are high and/or great precision is required, remote sense capability should be called out. This allows the supply to sense the actual voltage at the load, and adjust the supply such that the desired voltage is maintained at the load, not at the supply. There can be a significant voltage drop in the wires connecting the load to the supply.

The output current should be specified both with a nominal number, which should be a best guess as to what the average load will be, as well as a minimum and maximum number. If the maximum current differs substantially from the nominal current, then a duty cycle and time period should be provided. It should be noted that peak power requirements lasting for milliseconds or more are essentially steady state to the semiconductors in the supply. Only the thermal design can take advantage of a supply having a high peak to average current ratio. The minimum current drawn by the load usually becomes a consideration in multi-output supplies. This is because the usual topology for a multi-output supply involves a single control loop on the primary output, with secondary windings providing the extra outputs. The secondary windings are then left in what is referred to as a "semiregulated" state, or, they can include secondary regulators to provide fully regulated outputs. In order to maintain reasonable output voltage on all the secondary outputs, the primary output must have a minimum load on it. For this reason, most multi- output supplies specify a minimum load at least on the primary output. The secondary outputs also require a minimum load if they have no further regulation on them. Secondary outputs can be fully regulated either by putting linear regulators on the outputs, or by using magnetic amplifiers (magamps). Magamps are more efficient than linear regulators, but they require an extra inductor and associated control circuitry. When specifying the regulation on secondary outputs, try to specify a minimum load if you know that your application will always provide it, and don't overspecify the amount of regulation you need. This will allow your supplier to provide the most cost-effective unit possible.

Regulation should be specified as a percentage change under specific load changes. It can either be specified as "total" regulation, meaning that the output will be in the specified window under all conditions of line and load conditions, or it can be specified as individual components. Manufacturers typically specify the individual components in order to enable the user to estimate performance under his own conditions. Unfortunately, this often leads to confusion. The definitions of the various regulation components are as follows:

Line Regulation: The change in output voltage resulting from a change in input voltage. Usually specified as a percentage under specified input changes (1% from 90 to 270 VAC, etc.).

Load Regulation: The change in output voltage resulting from a change in load current. Usually specified as a percentage under specified load changes (1% from 25% to 75% load, etc.).

Thermal Regulation (stability)**:** The change in output voltage over a specified temperature range or time period. Usually specified as a percentage per degree, or a percentage over time.

Cross Regulation: This refers to the fact that in some multi-output supplies, the voltage of one output can be affected by load changes on another output. Usually specified as a percentage change for a specific change on another output (5% for a load change of 25% to 75% on output #1)

Ripple and Noise

Ripple is a term used to describe periodic variations in the output voltage which are usually fed through the output filter. The ripple frequency is the operating frequency of the primary transformer, which is typically 50Hz for a conventional linear supply or at the switching frequency of a switching power supply. Actually, most supplies have a combination of 50Hz ripple as well as high frequency ripple. A limit of 1% is typically placed on the outputs of switching supplies, and this is usually a good compromise between ripple and transient response. Lower ripple usually means higher output inductance and higher capacitance, which slows transient response. Noise is often not specified, and one common reason is that it is very difficult to measure consistently. The best and most consistent way to measure noise is with an oscilloscope and a piece of 50 ohm coaxial cable terminated at both ends with a 50 ohm resistor. A high frequency bypass capacitor is sometimes used at the supply output since most loads will have one on their input. To be thorough, any noise specification should include a bandwidth as well since the observed noise will be limited by the capability of the measuring instrument.

Interface Signals

Often there will be a need to interface the system to the power supply at the logic level. Sometimes this is to enable the system to do an orderly shutdown upon the detection of an imminent power failure. This signal is really a power good signal. The best way to specify it is to specify it as active low, such that when the input capacitors start to discharge, an open collector driver goes high and stays high. Other times, it is desirable to be able to turn the supply on and off with a logic signal. The question here is where the common ground point is. A common way to do this is with an optocoupler which drives the pulse width modulator. This allows most of the circuitry to float and yet the entire supply is controlled from a simple TTL compatible signal. Another typical interface signal is a synchronization input and/or output. This can be used to synchronize 2 supplies in a system to prevent a beat frequency from developing. Other times it is used to synchronize the power supply to the system clock in a large digital system. With the supply synchronized, the possibility of the coincidence of power supply switching noise with data transitions will be minimized.

Mechanical - Size & Weight

At first glance, specifying the size and/or weight of a power supply might seem relatively simple. However, there should be several considerations here. First, as much of an allowance as possible should be given on size. Ideally, a maximum size should be given to allow flexibility in design or second sourcing. The only part of the mechanical specification which usually needs to be firm is the mounting screw pattern. Allowing variations in size and shape will make it easier to find more than one manufacturer with a supply which will fit. Even when designing for a custom supply, sticking to an industry standard footprint will ease the search for a second source. Looking for a second source on that "L" shaped supply with the cutout 65 mm up the short side can be difficult at best, even though it is just a single 5 volt supply. The best approach is to investigate similar supplies to get a feel for the expected size & weight, as well as the typical footprints, and then write your specification such that these supplies would meet it.

Thermal Considerations

Probably the single most important factor in specifying a reliable power supply is the thermal design. The situation to be avoided is the one where a package concept is defined by marketing, the electrical design is done by the electrical engineers, the mechanical engineers have no choice but to design brackets to hold the whole mess together, and when the prototype goes together in the custom molded plastic chassis, the inside air temperature is measured at 95° C. The thermal design of the power supply cannot be isolated from the rest of the system, nor can data sheet references to allowable temperature rise be used alone. The power supply will generate heat, usually a fixed amount plus a variable amount based on output loading. Most supplies specify efficiency at full load, but it is poor practice to operate anything at full rated load. Consequently, the actual efficiency in your system may be lower than you think. In the absence of actual data, a good guess is to decrease full-load efficiency by 5% or so when operating at 60-80% full load. The heat generated by the supply will be the difference between the input power and the output power, or:

Input Power = Output Power / Efficiency
Power Dissipated = Input Power - Output Power

This power is the power dissipated by the supply. However, if the supply is mounted in the same enclosure as the load, the enclosure must usually dissipate almost ALL the supply input power, assuming that the load is not performing much useful work in the mechanical sense. The difficult problem here is how to calculate temperature rise, how to decide if a fan is needed, where to place vents for good air flow, etc. While months could be spent trying to calculate the amount of heat conducted, radiated or convected away from a particular surface, the simplest and surest way to attack the problem is with a simple physical model. Assuming you have a prototype enclosure to work with, it is quite easy to use resistors as heat sources of various types in the expected locations. If your system is almost complete, except for the supply, then the system should be assembled and run off of an external lab supply. A dummy supply (this can be as simple as a single aluminum plate) can then be installed with chassis-mounted resistors (Dale RH series or equivalent) attached to simulate the actual supply. If the supply is to be conduction cooled, the dummy chassis temperature can be expected to closely approximate the actual supply chassis temperature. If the supply uses convection cooling extensively, then the dummy chassis temperature will not necessarily be a good indication of the actual supply chassis temperature, but the temperature of the inside air will still be fairly close to the actual. The purpose of this sort of model is to evaluate the cooling capacity of a system. If the housing is a large, intricate molded plastic assembly, a model such as this is the only practical way to investigate temperatures. Armed with this test data, you will be better able to specify the highest ambient the supply will have to operate in, what the baseplate temperature will be (assuming conduction cooling) and what airflow will be required. You will also be able to evaluate the effect of power supply efficiency on the overall system by simply varying the power in the resistors. For more accurate testing, a dummy chassis can be set up using a combination of conduction and convection cooling, as appropriate. This is accomplished by using one or more resistors to simulate the conduction cooled power supply components, and air cooled resistors to simulate the convection cooled components. The power supply manufacturer should be able to give you the appropriate wattages for various supplies.

Cooling Method

Armed with a feeling for how well the supply can be cooled in the proposed package, the next step is to specify the cooling method. If there will be a metal baseplate which will be cooled reasonably well, then conduction cooling may be the most efficient way to get the heat out of your box. However, if there is not a good mounting surface to act as the cooling base plate, then convection cooling may be the best approach. Most power supplies use a combination of conduction and convection cooling, and so this should be considered. If the supply is going into a plastic enclosure, and the desired supply requires some conduction cooling, then a sub-chassis or just an aluminum mounting plate to increase surface area can be considered. Once again, the previously mentioned test methods can be used to model the system.

Temperature

Operating Temperature is really a part of the overall thermal design, although it is frequently specified separately. It is also a fairly deceiving specification, because it is really tied closely to airflow and to the thermal characteristics of the supply itself. A supply which is rated to operate to 70°C can overheat in a 25°C office environment if the supply is not mounted and cooled properly. Also, "environmental chamber" tests can be misleading, because a chamber is designed to maintain a uniform thermal environment, usually by high rates of air flow. This can give a false sense of security in a power supply test, since this air flow has the effect of dramatically decreasing the thermal resistance of the power supply heatsink system. It is not uncommon for semiconductor heatsinks to run 20° C cooler in a high rate of airflow at a given temperature.

Humidity

Power supplies destined for an office environment usually need not have a humidity spec called out, but many consumer applications and most industrial applications specify a humidity range. The most innocuous specification is to require operation in up to 95% humidity, non-condensing. This means that temperatures will not be changing fast enough to cause condensation on the surfaces of the supply. Barring condensation, any supply should work under high humidity conditions, except, perhaps, a particularly high voltage supply. If the supply will be located in an area where condensation will be possible, condensation must be included in the operating specification. This may force the pc boards to be conformally coated, or possibly the supply would be totally sealed or potted. While this becomes a major feature of the supply, it is a necessary one.

Salt Spray

Although this is generally considered a military specification, there are many commercial applications where resistance to salt spray is indicated. There are many marine applications where the end product will be exposed to this sort of environment, and as electronics penetrates the marine marketplace further, more applications can be expected. Military specifications such as MIL-STD-810 give detailed procedures for salt spray testing, but for a commercial product, a more benign test may be appropriate. The best approach here is to discuss the requirement with the manufacturer to determine the best way of assuring everyone that the supply will perform adequately in the expected environment.

Altitude

There is usually no problem in meeting an altitude specification, since it simply means a bit more attention must be paid to high voltage insulation. Convection cooling is much less effective, but air temperatures are cold enough at high altitudes that the effectiveness of the cooling air is not dramatically affected. In addition, airborne applications typically have plenty of air volume, which also helps. Commonly, 3200 m is called out if the unit is not expected to fly, or if it would always be in a pressurized environment (Humans generally require help in breathing beyond 3200 m). If the unit must operate at higher altitudes, this should be called out. Space applications are more demanding, since air cooling is usually not possible. In addition, temperatures can usually be expected to vary widely due to solar heating.

Shock and Vibration

This is a very important specification, since almost any supply needs to have shock and vibration specified. Even if the supply will be desk-bound, a vibration spec is a good idea because it can be tailored around shipping vibration. Most manufacturers take this into consideration automatically, since power supplies which can not survive shipment to the customer are a most obvious problem. If the power supply will be operated in a stationary location, it is a good idea to let the specification take the form of a shipping specification with the supply in it's shipping container. Most power supply manufacturers can certify that the shipping container used will meet minimum shipping requirements. Some customers like to add a shock specification as an indication that the supply should be "rugged". This can sometimes be taken to the extreme, especially if testing is required. Often, the intent is to allow the unit to survive "bench handling", which is a term often used to describe an occasional fall to the floor. If this is the intent, try to say so clearly. A simple drop test may be easy to verify informally, while a detailed shock specification may require outside laboratory testing. If the supply is to be used in a mobile application or similar high vibration environment, a more direct vibration spec should be imposed. MIL-STD-810 offers many versions for different types of environments.

Mean Time Between Failures (MTBF)

Specifying an MTBF must be done with some detail for this specification to be meaningful. MTBF can be determined either by calculation or by demonstration. Demonstrated MTBF is an ideal goal, but it is usually impractical. Many units must be run for many hours, and the variations due to manufacturing anomalies can make consistent results difficult at best. Usually, a calculated MTBF is used as a verification that the design is relatively sound, but the temperature which this calculation is based on must be specified. MIL-HDBK- 217E calls out the procedure for summing up the failure rates of all the individual components, giving a total failure rate. This number should be used for comparative purposes only, because the actual failure mechanisms in power electronics are never adequately accounted for in a relatively simple procedure.

Burn-In

This is a crucial part of the power supply manufacturing process. It should be a part of the power supply specification mainly to guard against a manufacturer "cutting corners". When reliability is a primary consideration, specifying a longer and more aggressive burn-in is usually a cost effective way to find weak components. A manufacturer may use a burn-in time as short as 1 hour, or up to a week or more. A "cycled" burn-in is usually more effective than a constant load burn-in. Turning the supply on and off forces the unit to cycle thermally, which often reveals bad solder connections and bad semiconductor wire bonds. It also surges the circuitry, stressing different components under transient conditions. Since most failures occur relatively early on in the cycle, 24 hours is usually a good compromise. Periodic review of burn-in data will show up many interesting details, and can be used to adjust burn-in times. For instance, if 99% of all burn-in failures occur in the first hour, a 24 hour cycle might not be justifiable on a cost basis. Working closely with the power supply manufacturer will yield large dividends in terms of overall reliability and cost effectiveness.

Configuration Control

When large quantities of units are purchased on an OEM basis, the user may need to exercise control over changes. If a manufacturer decides to move the mounting holes of the supply to "improve" the unit, the customer who just ordered 10,000 cabinets to mate with the "old" hole pattern may want to be informed in advance. When buying an "off-the-shelf" unit, a power supply manufacturer may be reluctant to grant configuration control to the customer unless the volume is substantial. Usually, some arrangement can be worked out to protect the user from major changes to a product or from a unit being "dropped" with little or no notice.

Process Control

This is similar to the configuration control described above. In some cases, a user has particularly strong feelings about certain processes he either wants to see or wants to preclude. A common example would be a restriction on using certain hazardous materials (capacitors with PCB's, beryllium insulators, etc.) or on certain processes (no halogenated solvents on unsealed electrolytic capacitors). Once again, if there is a good reason for it, specify it. But by all means, discuss it with the power supply manufacturer. The supply manufacturer should ordinarily be considered the authority on such matters.

Approvals

Having agency approvals on a power supply can make the process of obtaining approval for the entire system much easier. Since most requirements are centered around the hazards of the line voltage, and most power supplies have relatively low voltage outputs, the supply is usually center stage in the agency investigation. The most basic approval is the appropriate UL approval. Unless a supply is a completely packaged standalone unit, it would not be possible to obtain a UL approval on the unit itself. Instead, it is usually done under the UL recognized component program. Under this program, the component is recognized for use in a particular type of equipment or application. The end product must obtain UL approval, but this is much easier if the individual components (such as the power supply) are already covered as recognized components. The next approval which is usually sought is CSA, the Canadian agency which closely parallels UL in the United States. The requirements are usually quite similar to UL, and CSA will usually want to see a UL report first., there are many labs which will certify a product as meeting to VDE standards. The same holds true for IEC standards, as well as FCC standards. The major point to keep in mind is that the cost of obtaining all these certifications is relatively high, and so it may not be practical to impose them on a manufacturer if the dollar volume of the product is relatively low. If the potential volumes are low, and the agency approvals required are extensive, then an off-the-shelf product which carries the approvals already may be the only viable alternative.

Military Specifications

If a power supply is associated with military hardware, then very often it will be required to meet various military specifications. One of the most common system specifications which involves the power supply is MIL-STD-461. This specification limits the amount of "noise" a power supply may generate and defines what types and levels of interference it must tolerate. Meeting the limits of this specification can be difficult, adding size and weight to the supply. If security is an issue, very often "TEMPEST" requirements are imposed. This places restrictions on the system to prevent any intelligence from inadvertently leaking out. In general, it requires the power supply to have very high attenuation from output to input. Another specialized requirement is for a unit to survive a "nuclear event". This requires specialized surge protection as well as special component selection and circuit design techniques.

Another standard is MIL-STD-454. It relates general requirements for electronic equipment, and is very often a "flow down" from the prime contractor. It does not require exceptional effort to meet. "Weapon-Spec" and "Hi-Rel" soldering are terms used to describe high reliability soldering. Originally referred to as WS-6536, MIL-S-45743 and later as MIL-STD-2000, it specifies in great detail how a reliable solder joint should be made, and how it should be controlled. In order for a manufacturer to qualify, someone from the organization must be qualified as a certified instructor at one of several locations operated by the military. That person must then set up an effective in-house training program for production personnel. Aside from the cost of setting up such a program, there are recurring operational costs to maintain such a stringent system, which obviously affects the unit cost.

Very often, a contract will have quality requirements called out. Most typical is MIL-I-45208. It contains detailed requirements for equipment calibration, handling of discrepant material, and vendor surveillance, to name a few. When very high quality standards apply, then MIL-Q-9858A is called out. This relatively brief specification refers to many others, and usually requires considerable effort to comply with. As always, the key is not to overspecify.

Environmental test methods are elaborated on in MIL-STD-810 as well as others.

Space Applications

When a power supply must be designed to operate in space, there are usually a complete set of the most stringent requirements to be met. Aside from the obvious environmental problems such as shock, vibration, temperature and radiation, there are quality requirements, component selection restrictions, approval of non-standard components, special circuit design techniques, EMI requirements, and extensive documentation requirements, to name just a few. As an additional requirement, size and weight are usually at a premium as well.

The quality requirements usually involve traceability of all components, and extensive vendor surveillance. Very often the program will require MIL-Q-9858A compliance, or, as a minimum, compliance to the relevant parts.

Component selection must usually be attempted from a program-approved parts list, with exceptions handled in a specified way. Usually this means that justification must be written for the use of the non-standard part, and a specification control drawing must be generated. Certain parts, such as opto- couplers, must usually be avoided because of their inherent susceptibility to radiation. Although FET's can be used, specific circuit design techniques must be used to assure that a shift in gate threshold voltage over time will not cause problems. Several manufacturers now offer "radiation hardened" FET's for these applications.