MAINTAINABILITY: THE MISSING PIECE OF THE AVAILABILITY PUZZLE

J R Sears

Liebert Corporation, USA

RECLAIMING THE LAST FOUR HOURS

The UPS industry has a secret, and it's four hours long. That's the average length of time that the UPS and switchgear must be completely de-energized back to the service entrance each year for maintenance.

There's no way to escape this downtime. Remember: *Deferred maintenance is NOT the same as high availability.*

This does not mean that your critical load will not have access to power during those four hours. It simply means that *most* (but not all) system configurations will expose your critical load to unconditioned generator or mains power during that interval. This unconditioned power is politely referred to as "bypass" power.

Those four hours of unconditioned power cast a dark shadow over your system availability calculations. Therefore, for the remainder of this paper, the term "availability" will refer to *the availability of power that has been conditioned by the UPS and is backed up by a reliable battery system.*

The good news is that certain UPS system configurations permit *concurrent maintenance* – supporting the load equipment on conditioned power while de-energizing one complete power system back to the service entrance. We will describe these system configurations in later sections.

HOW HIGH IS "HIGH-AVAILABILITY?"

The computing and telecommunications industries talk in terms of "Nines" of availability. This refers to the percentage of time in a year that a system is functional and available to do productive work.

In a standard 365-day year, there are 8760 hours. If you must lose four hours of availability for maintenance, your theoretical maximum is 99.95%, or about three "nines" of availability. To achieve four "nines" or 99.99% availability would mean reducing that exposure to 53 minutes per year. Five "nines" (99.999%) availability means less than 5.3 minutes of unconditioned power per year. Six "nines" (99.9999%) of availability equates to just *32 seconds* of unconditioned power per year.

Fortunately, advances in UPS technology make it possible to completely avoid all exposure to unconditioned power. By making careful choices, the facility designer can achieve *Continuous Availability* of conditioned power.

THE ROAD TO CONTINUOUS AVAILABILITY

The high availability systems require four key elements:

Reliability

The individual UPS modules, static transfer switches and other power distribution equipment must be incredibly reliable, as measured by field-documented MTBF (Mean Time Between Failures). In addition, the system elements must be designed and assembled in a way that minimizes complexity and single points of failure.

Functionality

The UPS must be able to protect the critical load from the full range of power disturbances, and only a true double-conversion UPS can do this. Some vendors offer single-conversion (line interactive) three-phase UPS products as a lower-cost alternative. However, these alternative UPSs do not protect against *all* disturbances – including power system short circuits, frequency variations, harmonics and common-mode noise. If your critical facility is truly critical, only a true double-conversion UPS is suitable.

Maintainability

The system design must permit *concurrent maintenance* of all power system components – supporting the load with part of the UPS system while other parts are being serviced. As we shall see, single-bus solutions do not completely support concurrent maintenance.

Fault Tolerance

The UPS system must have fault resiliency, to cope with a failure of any single power system component without affecting the operation of the critical load equipment. Furthermore, the power distribution system must also have fault resiliency, to survive the inevitable load faults and human errors.

This paper was presented at INTELEC® 2001 Conference in Edinburgh, Scotland, October 14-18, 2001

For purposes of this paper, let us assume that the facilities manager and design consultant have chosen a reliable double-conversion UPS from a reputable manufacturer. With Reliability and Functionality assured, let us look at how different UPS system configurations compare for Maintainability and Fault Tolerance.

Single Module Systems

A conventional double-conversion, single-module UPS is shown in Figure 1 of the Appendix.

With this configuration, the critical load is exposed to unconditioned bypass power during the times when the UPS, batteries or downstream distribution equipment need preventive maintenance or reconfiguring. For simple single-module systems, scheduled preventive maintenance generally requires between two and four hours per year. This limits availability to about 99.95%.

In addition, while the single UPS module has some degree of fault tolerance built in, it can't shield the critical load from downstream equipment failures and faults. Nor is there any redundancy in capacity, to protect against a failure of the module itself.

The single module UPS gets good marks for reliability and simplicity. It is ideal for companies that are able to schedule system downtime for maintenance or configuration changes. But it still comes about four hours short of continuous availability.

Parallel Redundant Systems: The Reliable Standard

A parallel redundant system has two or more UPS modules connected in parallel to a common distribution network. The system has enough modules to carry the maximum projected load, plus at least one additional module in parallel to provide redundancy. See Figure 2 in the Appendix.

Under normal operations, all the modules are on-line and sharing the load equally. If one module fails or needs to be taken off-line for maintenance, the other modules have enough capacity to carry the full system load. It is still necessary for the system cabinet and switchgear to go to bypass (or completely shut down) for between one and four hours per year to service bus bars, circuit breakers and other system-level components. A parallel redundant system can provide up to 99.99% availability in an ideal world where system downtime can be limited to one hour per year. Real life is more like three-to-four hours per year.

A parallel redundant system works well for companies that can schedule load shutdowns for maintenance or reconfiguring.

Distributed-Redundant Systems: Continuous Availability

This configuration features two or more *independent* UPS systems, each capable of carrying the entire critical load. Each system provides power to its own independent distribution network. There are no power connections between the two UPS systems, which can be either single modules or parallel redundant systems. See Figure 3.

The goal of Distributed Redundancy is to bring power system redundancy to every piece of load equipment, as close as possible to the input terminals. This redundancy is the key to ensuring both maintainability and fault tolerance throughout the facility.

If your downstream loads have dual power cords, then it's easy to design power distribution equipment that will give them complete and direct access to the output of both independent UPS systems. The facility operator can simply shut down one complete UPS system and the dual-corded loads will continue to function on the other UPS. Single-corded loads present challenges, as we shall see.

A few years ago, the computer industry appeared to be moving strongly toward dual-corded power supplies. In recent years, the financial model of Internet Data Centers has pushed operators toward high-density, 1U servers, which are inherently single-corded loads.

The balance of this paper will describe ways to add maintainability and fault-tolerance to a room that must support both single-corded and dual-corded equipment.

As a further challenge, we shall also examine creative ways to extend this maintainability to buildings with multiple tenants and sites with a mixture of AC-powered and DC-powered equipment.

Selective-Redundant Systems: Ideal for Multi-Tenant, Mixed-Use Facilities

Some tenants in a building might need the 24 x 7 protection and maintainability of a dual-bus, Distributed-Redundant system. However, other tenants have less-critical missions and could be well served by a conventional single module or parallel redundant system. One approach to satisfy both types of operations is a Selective-Redundant System, shown in Figure 4.

Picture a multi-story office building housing your business plus tenants that rent space from you. A conventional single-bus UPS system with standby generator can provide basic outage protection, enabling you to provide premium power for the entire facility.

Within your facility, certain areas are critical. You may have a data center on the third floor, an army of engineers and programmers (and their servers) on the seventh and eighth floors, a tenant with a corporate website and database on the tenth floor, and various network equipment closets on every floor. How do you satisfy these requirements?

The solution is Selective Redundancy – creating Distributed-Redundant power systems by selectively "hardening" critical portions of your facility. Add enough additional UPS capacity to create a dual-bus system for each area. See Figure 6 for an example system.

A typical system might have a 2x800 kVA parallel-redundant UPS on the ground floor to power the entire facility. Then you could add a 2x500 kVA system for the data center, a separate 30 kVA system for the seventh-floor servers, and a separate 100 kVA system for the tenth floor tenant.

Such a configuration creates three independent Distributed-Redundant Systems within the facility. The UPS systems must be kept in sync, even when one or more are operating on batteries or the standby generator. This will enable individual loads in each critical work center to be switched transparently between the facility-wide UPS and the dedicated UPS.

Hybrid AC-DC Power Systems: Optimized for Telecomm and Hosting Facilities

Traditional telecommunications power systems consist of multiple parallel-redundant rectifiers that convert commercial AC power to –48 VDC power that charges lead-acid storage batteries and supplies power to critical load equipment. Long battery support times are the norm: from 1 hour to more than 24 hours, with the typical range being 3 to 8 hours.

By contrast, most data center equipment runs on AC power, provided from a UPS. The typical UPS has approximately 15 minutes of battery storage plus a standby generator for long-term power outages.

The convergence of voice and data transmission has forced AC and DC equipment to co-exist in the same facilities. Hybrid AC-DC Power Systems can make it a peaceful coexistence.

The Hybrid system begins with the Distributed-Redundant Power System explained earlier in this document. Two independent UPS systems support two independent power distribution systems. Each distribution system powers a set of DC rectifiers, creating 48 VDC for the telecommunications equipment. The DC load equipment has dual power cords, so it can receive power from both sets of rectifiers. See Figure 5.

An immediate advantage of this configuration is elimination of battery storage for the 48 VDC system. Instead, the DC equipment can rely on the batteries and backup generator supporting the AC UPS equipment. Engineering and installation costs should also be considerably less. System complexity and parts counts are reduced, thereby improving general reliability. This configuration is particularly applicable where the system operation is co-dependent on having both the AC- and DC-powered load equipment operational.

Most importantly, the Hybrid AC-DC System provides the utmost in maintainability and fault tolerance, for the highest possible overall system availability at a lower cost than conventional AC and DC power systems.

DISTRIBUTION OPTIONS FOR OPTIMUM MAINTAINABILITY

Let's assume you have chosen a Distributed-Redundant UPS system with dual-load-bus power distribution. Your next task is designing a power distribution system to give maximum fault tolerance and maintainability to your critical loads. Figure 6 shows the different ways that load equipment can be connected.

Option One: All Dual-Power-Corded Loads

Dual-power-cord machines can be fed by two conventional Power Distribution Units (PDUs). If the machines can operate equally well on either cord, you will not need a Static Transfer Switch to transfer the loads between the UPS systems. All power system components are serviceable while the load is on-line. If the loads are rack-mounted servers, all your electricians need to do is provide dual outlet strips inside each cabinet. This is the simplest and least-expensive configuration.

Option Two: Manual Circuit Breaker Transfers

This option has a dual-input Power Distribution Unit with manually operated breakers feeding single-input loads. Breakers are used to transfer the PDU and load between UPS 1 and UPS 2 when one UPS or the other requires service. This is a simple, low-cost approach but it only works for certain cases.

In particular, the UPS systems must be kept in synchronization to enable safe, closed-transition transfers to occur (More on this subject later.). Furthermore, the system designer must provide some amount of interlocking between the UPS system logic and the distribution circuit breakers.

This arrangement adds a measure of maintainability but does not provide any additional fault tolerance. This limits overall system desirability compared to other possible configurations.

Option Three: Static Switch on Primary Side

A Static Transfer Switch (STS) ahead of the PDU provides uninterrupted transfers between UPS systems. The fast-switching capability of the static switch provides protection against “fast” power system failures as well, such as source failures, breaker trips and even operator error. This arrangement adds both maintainability and fault tolerance.

For these systems, the designer must consider several factors. First, the UPS systems must be kept in constant synchronization at all times to enable safe automatic transfers between sources. Second, putting the STS upstream of the PDU transformers minimizes the size of the STS required, since it can operate at the higher voltage and lower current of the UPS output – typically 480 VAC in the United States. Third, this arrangement works best with larger UPS systems where the load of any one PDU is a relatively small portion of the total connected load. Fourth, this arrangement sometimes will experience problems if a UPS suddenly fails or if there is a fault or breaker failure upstream of the STS. The STS will faithfully transfer the load, but the input voltage drop might cause the PDU transformer to lose magnetization. So the good UPS must not only contend with the added load, it must support the magnetizing inrush of the PDU transformer. If there are multiple PDU transformers all re-magnetizing at the same instant, there can be problems.

Option Four: Static Switch on Secondary Side

This arrangement also adds both maintainability and fault tolerance while overcoming some of the technical limitations described above. This configuration has the STS receiving power from two PDU transformers, while supporting various single-corded loads through circuit breaker panels.

As before, there are many considerations. First, the two UPS systems must be kept in synchronization to ensure safe, open-transition transfers. Second, both PDU transformers remain magnetized and ready to support loads at all times; this prevents a magnetizing inrush in the event of a load transfer. Third, this arrangement has optimal maintainability, since either transformer can be de-energized for maintenance while the other supports the connected loads.

Option Five: Point-Of-Use Transfer Switches

Hosting facilities frequently provide dual power strips to each cabinet, only to discover that a large portion of the load equipment is single-corded. This effectively prevents the facility operator from ever shutting down the branch circuit that feeds a particular equipment cabinet.

To give maintainability to these circuits, the operators can install point-of-use transfer switches. These devices, typically rated 20 to 60 amps, can be installed under the raised floor or in the equipment rack. Like static transfer switches, the smaller switches receive power from both UPS systems and can perform fast transfers between power sources.

In addition to adding maintainability and fault tolerance, the point-of-use transfer switches give the facility manager a great deal of flexibility. The small switches can be deployed quickly and reconfigured easily as the mission changes.

THE NEED FOR UPS SYCHRONIZATION

Unless all the downstream loads are true dual-cord loads, the output of the UPS systems must be kept in sync, to avoid the possibility of an out-of-phase transfer and the resulting load disruption.

Any UPS will typically sync to its bypass source. As long as both UPSs are tied to the same bypass source, they will automatically stay in sync in normal operation. However, if the UPSs are operating on batteries, on different backup generators or on asynchronous bypass sources, their outputs will tend to drift out of sync.

Fully functional Distributed-Redundant systems require the UPS modules to sync to each other under all circumstances, even when the modules are operating on separate batteries or non-synchronized backup generators. *Only true double-conversion UPS modules can do this*. Single-conversion UPS products (defined as “offline” or “line-interactive”) can only sync to their own input power sources.

CONCLUSION

Each of the system configurations described in this paper has its own advantages and disadvantages. Single module and parallel redundant systems are proven reliable but do not permit concurrent maintenance. The Distributed-Redundant, Selective-Redundant and Hybrid AC-DC systems all permit concurrent maintenance and can provide Continuous Availability of conditioned power, given the proper power distribution system.

APPENDIX: SYSTEM DIAGRAMS

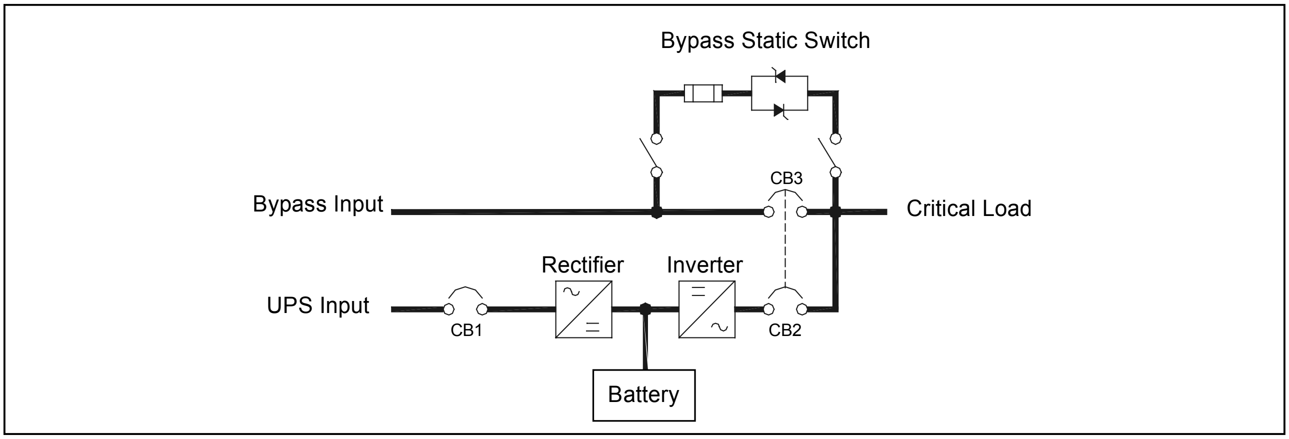

Figure 1: Single Module UPS System

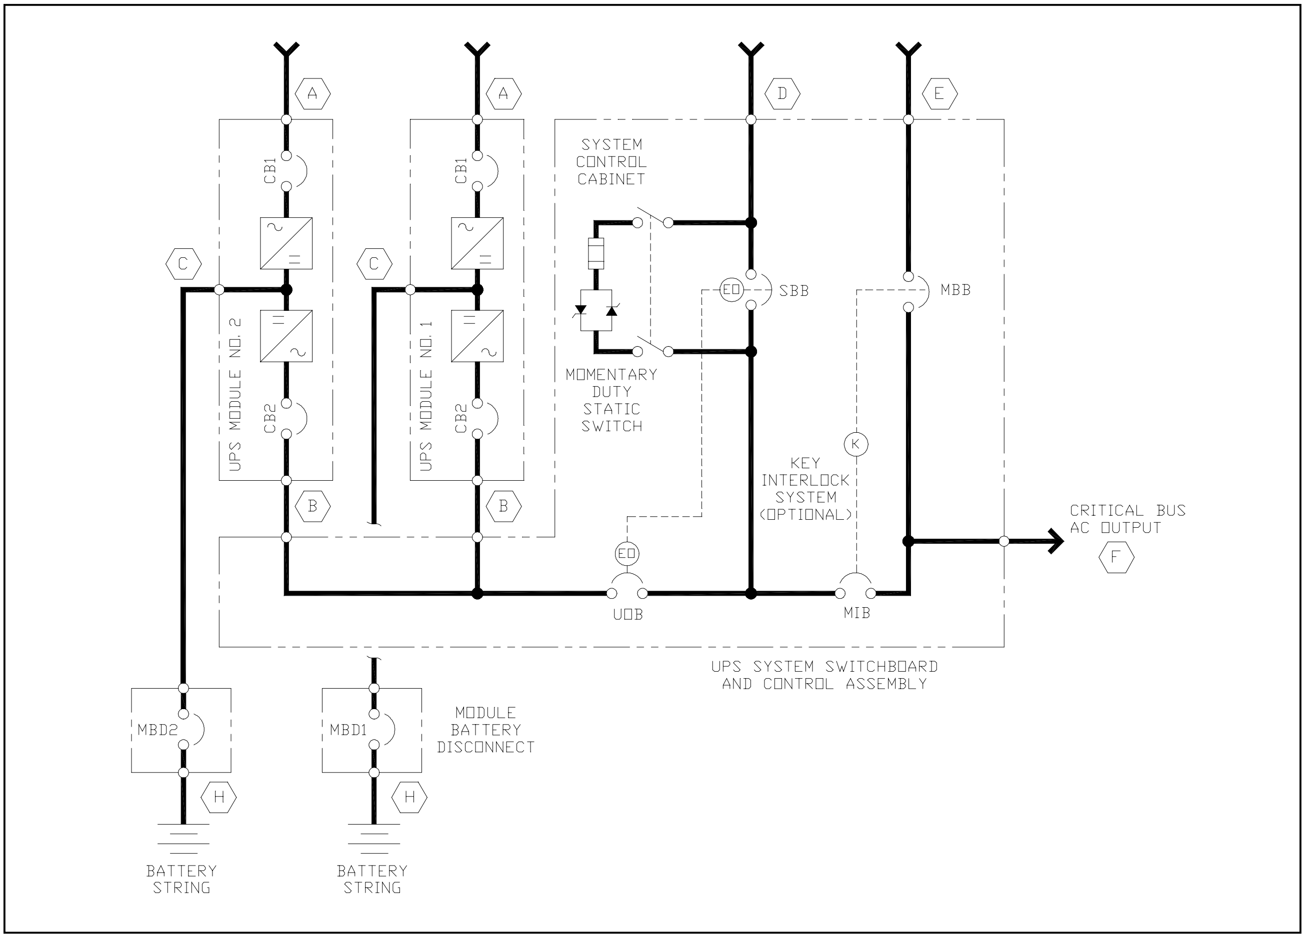

Figure 2: Parallel-Redundant UPS System

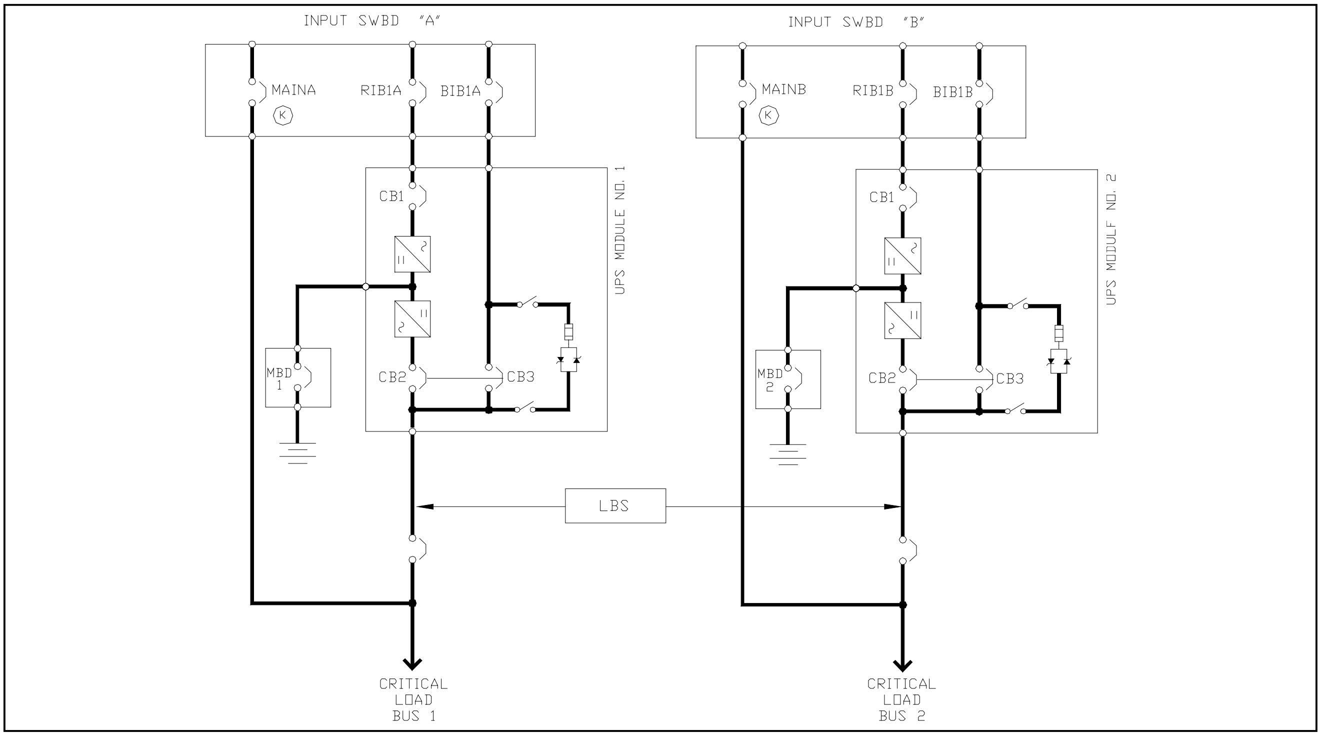

Figure 3: Distributed-Redundant Power System

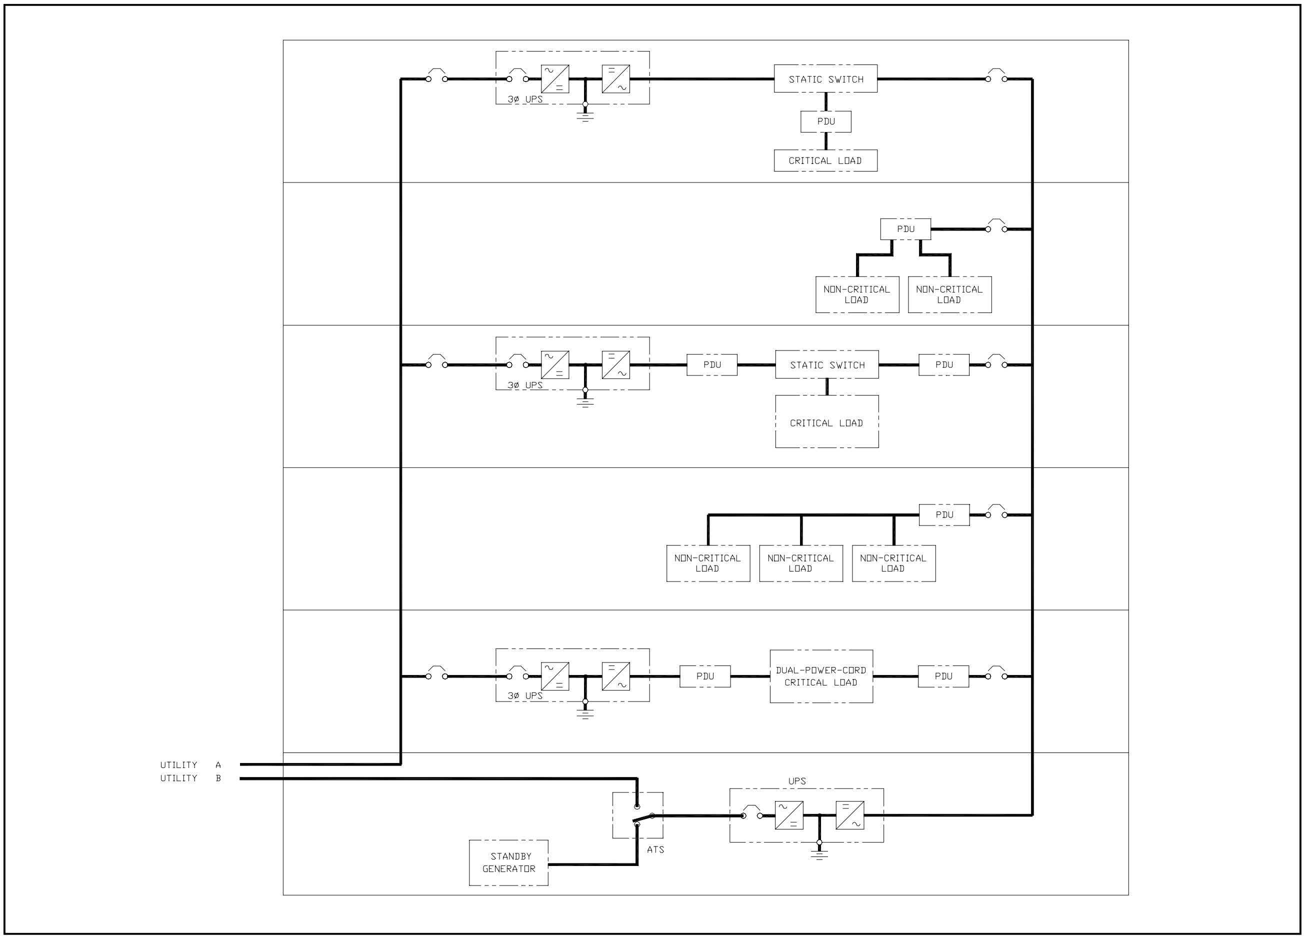

Figure 4: Selective-Redundant Power System

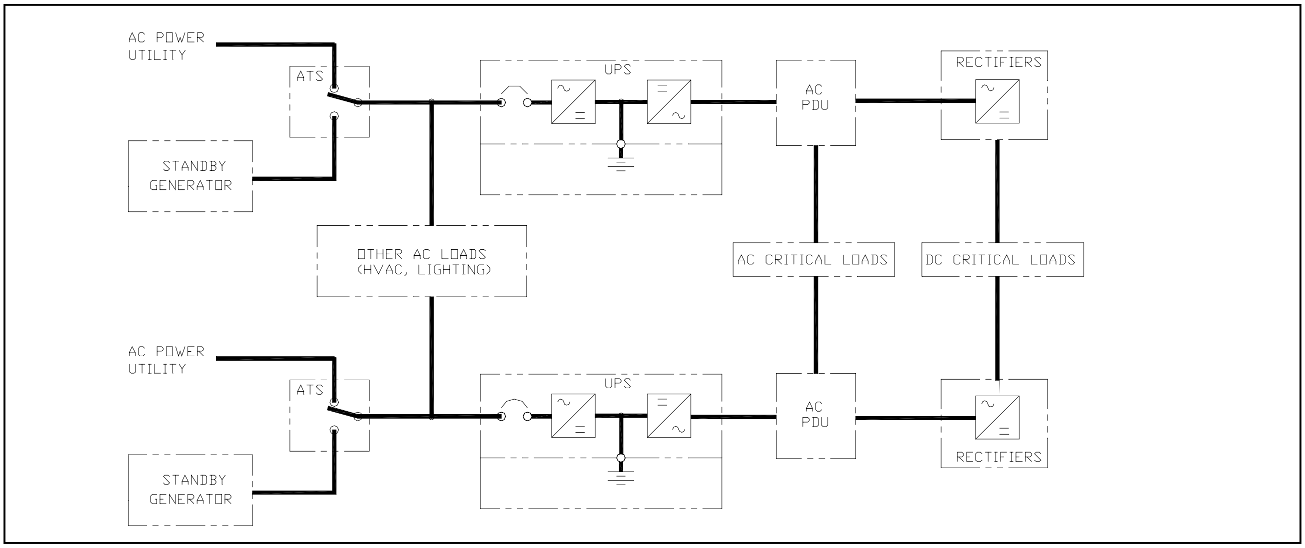

Figure 5: Hybrid AC-DC Power System

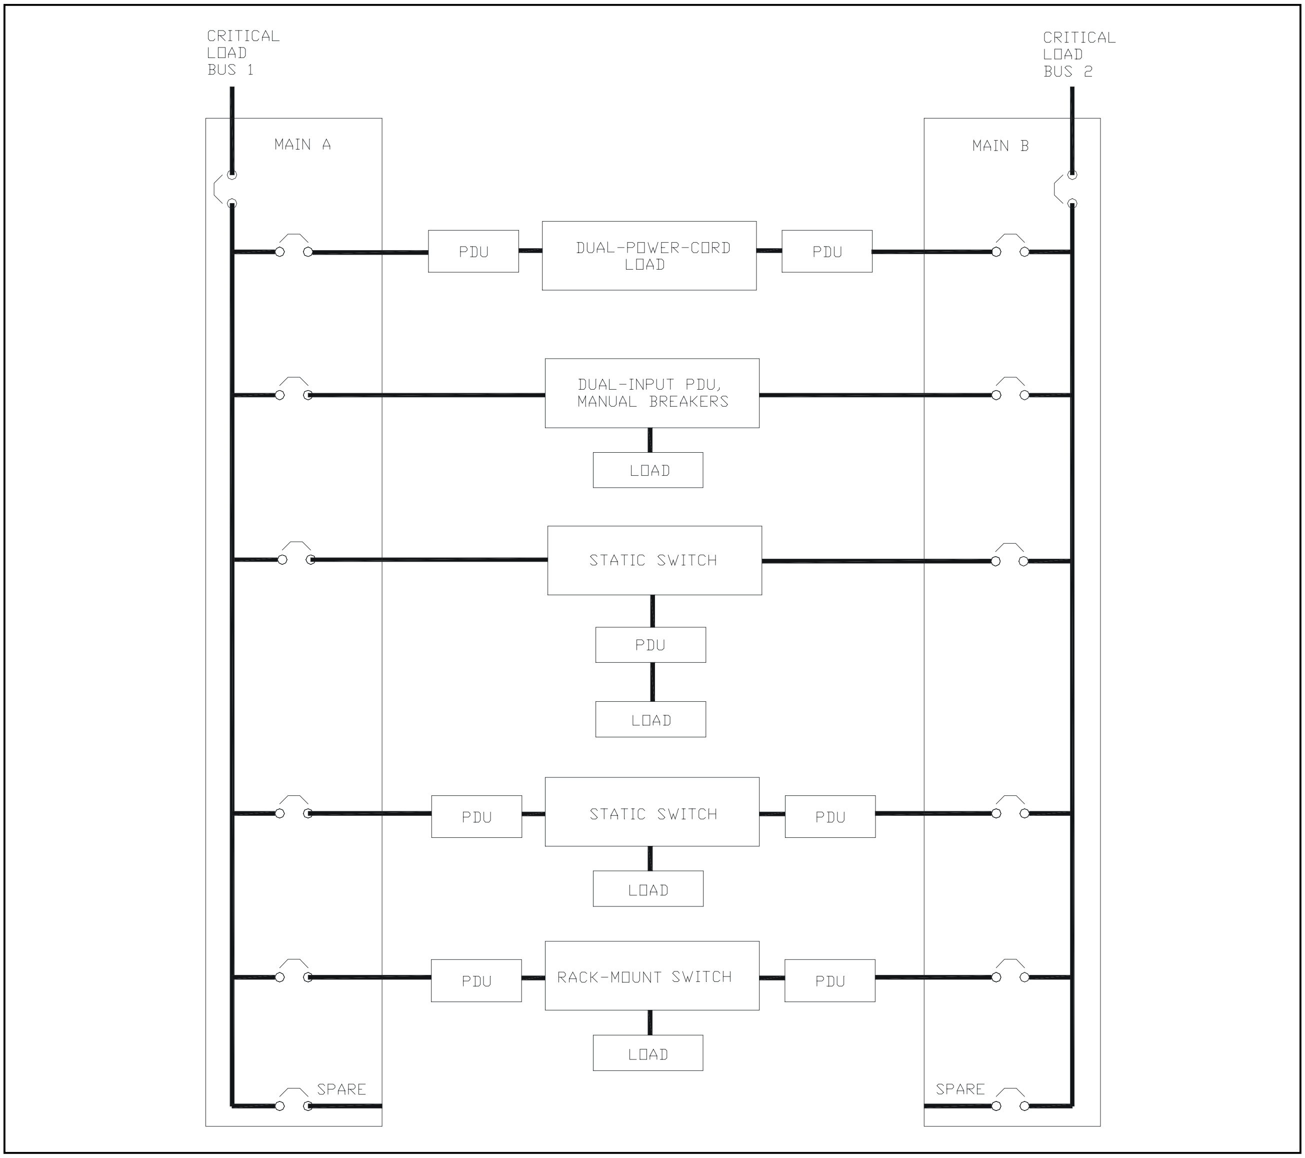

Figure 6: AC Distribution Options