

July 2003

High-Availability Power Systems, Part II: Redundancy Options

SUMMARY

“High-Availability” power systems require Reliability, Functionality, Maintainability, and Fault Tolerance. A few UPS configurations can overcome the “Last Four Hours” problem to achieve Continuous Availability of conditioned power. These configurations also greatly improve the power systems’ fault tolerance, as we shall see.

The “Last Four Hours” – The Industry’s Little Secret

First, the Bad News: All UPS equipment and switchgear, regardless of manufacturer, should be completely de-energized for preventive maintenance at least once per year. This maintenance generally requires between one and four hours of scheduled downtime per year, depending on system configuration.

Preventive maintenance *must* be performed, and the resulting scheduled downtime limits the availability of the power system. Remember: Deferred maintenance is *not* the same as high availability.

Now the Good News: 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.

High-Availability Power Systems

The computing industry talks 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. A system with four “Nines” is 99.99% available, meaning that downtime is less than 53 minutes in a standard 365-day year. Five “Nines” (99.999% available) equates to less than 5.3 minutes of downtime per year. Six “Nines” (99.9999% available) equates to just 32 *seconds* of downtime per year.

These same numbers apply when we speak of *availability of conditioned power*. The goal is to maximize the availability of conditioned power and minimize exposure to unconditioned utility power.

At Liebert, we take this concept one step further and encourage customers to seek *Continuous Availability* of conditioned power. After all, 100% is greater than 99.99999%.

The Road to Continuous Availability

We determine Availability by studying 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. All Liebert 3-phase UPS units are true double-conversion products.
- **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 system must have fault resiliency, to cope with a failure of any power system component without affecting the operation of the critical load equipment. Furthermore, the power distribution system must have fault resiliency, to survive the inevitable load faults and human error.

All Liebert 3-phase power products have field-proven critical bus MTBF in excess of one million hours, and all Liebert 3-phase UPS products use double-conversion technology. These two factors ensure Reliability and Functionality.

With Reliability and Functionality assured, let us look at how different UPS system configurations compare for Maintainability and Fault Tolerance. The table below shows eight system configurations, which will be explained in detail in the balance of this paper.

System Configuration	Concurrent Maintenance?			Fault Tolerance?		Maximum Availability?
	Module	System	Distribution	Module	Distribution	
Single module	No	No	No	Minimal	No	99.95%
Parallel Redundant	Yes	No	No	Yes	No	99.99%
Small Isolated-redundant	Yes	No	No	Yes	No	99.99%
Large Isolated-Redundant	Yes	No	No	Yes	No	99.99%
Distributed Redundant	Yes	Yes	Yes	Yes	Yes	Continuous
Selective Redundant	Yes	Some	Selectively	Some	Selectively	Continuous
Power-Tie™	Yes	Yes	Yes	Yes	Yes	Continuous
Hybrid AC-DC Power System	Yes	Yes	Yes	Yes	Yes	Continuous

Single Module Systems

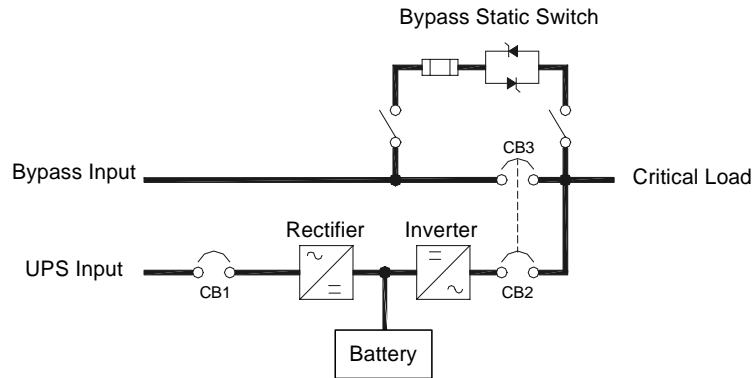


Figure 1: Single Module System

A conventional double-conversion, single-module UPS is shown in Figure 1. The UPS rectifier takes AC power from the utility (or engine generator) and produces DC power. A portion of the DC power is used to float charge the batteries. Most is fed to the inverter to power the load.

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 can schedule system downtime for maintenance or configuration changes. The critical part of single module system selection, then, is finding the UPS module with the highest field-proven critical bus MTBF.

Parallel Redundant Systems: The Reliable Standard

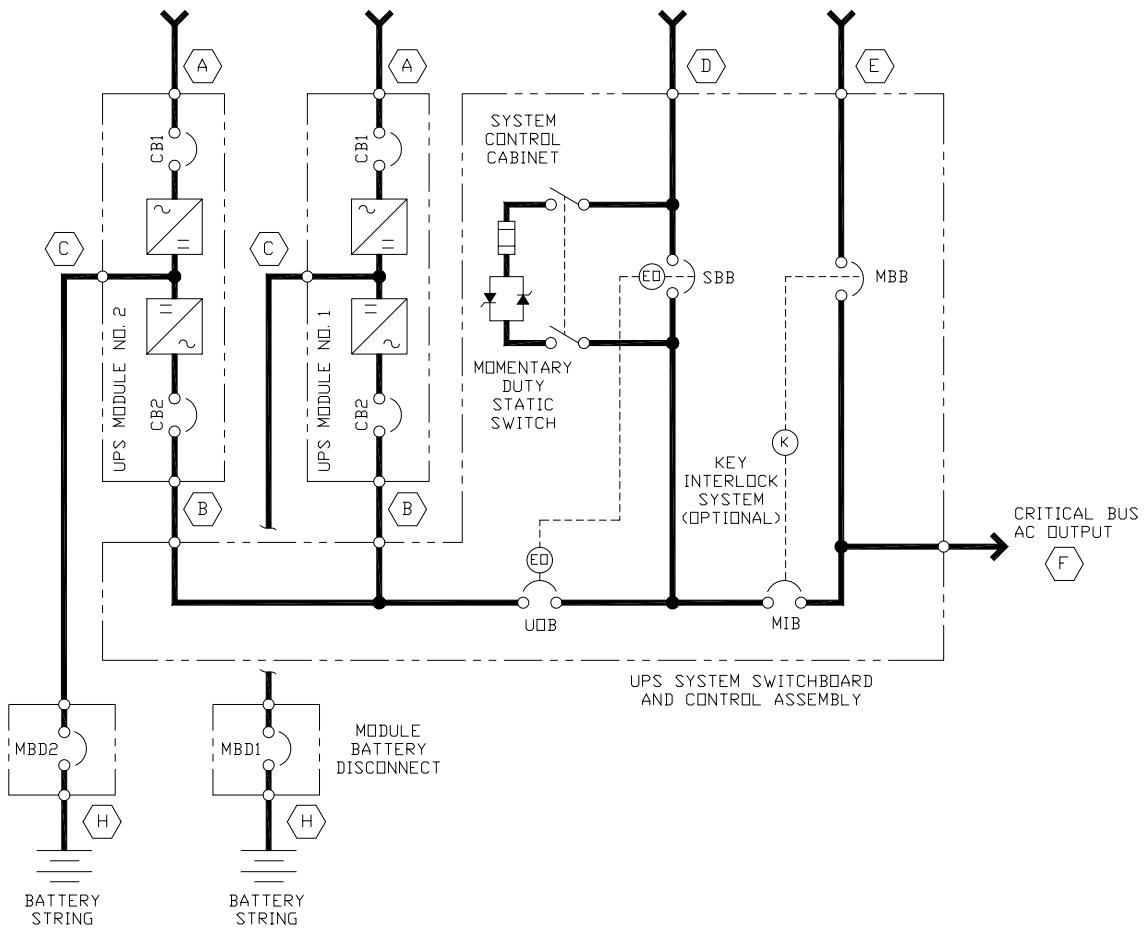


Figure 2: Parallel Redundant System

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. Typically, each module has its own battery plant.

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 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. Furthermore, load shutdown may still be required to service or reconfigure equipment between the UPS and the load. A parallel redundant system can provide up to 99.99% availability.

In addition to extra UPS modules, the parallel redundant system needs to give the operator some measure of system-level functionality. For this reason, Liebert normally provides a System Control Cabinet (SCC) with its Multi-Module Units. Note the system diagram in Figure 2.

The SCC gives system-level information and alarm monitoring. It provides module-to-module isolation, protecting the system from a fault in one module. It also contains the system-level static switch and bypass circuit, avoiding parallel connection of individual module static switches. For systems with custom-built switchgear, Liebert integrates the SCC logic and operator-interface elements into switchgear cabinets provided by the major manufacturers.

The System Control Cabinet is not required for load sharing. Connect any Liebert Series 600 UPS modules of the same model rating in parallel and they will load share within 15% *without any control connections*. This is partly due to inverter design and partly due to close impedance matching during manufacturing. Add the relatively simple control connections and they will load share within a few percent.

The conventional parallel redundant configuration is the industry standard for reliability. It is conceptually simple. The hardware arrangement is straightforward. However, output distribution is still limited to a single load bus per system.

A parallel redundant system works well for companies that can schedule load shutdowns for maintenance or reconfiguring. But “works well” isn’t good enough for applications that demand “never fail” or “never go to bypass,” since parallel redundant systems are still vulnerable to potential failures between the UPS and the load. For applications that must “never fail” or “never go to bypass,” we recommend Distributed-Redundant or Selective-Redundant systems, described in later sections.

Note on “1+1” Systems

A special case of Parallel Redundant products are the “1+1 Redundant” systems such as the Liebert Npower 1+1R. These systems eliminate the System Control Cabinet in favor of a simple Paralleling Cabinet.

The obvious benefits are lower cost and reduced complexity. In smaller systems (below 150 kVA), system-level controls are less important than they are in larger systems.

The trade-off is reduced system-level functionality. The various 1+1 products generally do not have system-level monitoring, controls or a system-level static switch. These configurations rely on the individual module static switches to manage load transfers and handle temporary overload conditions.

For calculating availability, these systems can be treated approximately the same as regular parallel redundant products. They still require between one and four hours per year of downtime, so availability is limited to approximately 99.99 percent.

Isolated Redundant Systems: Calculated Tradeoffs

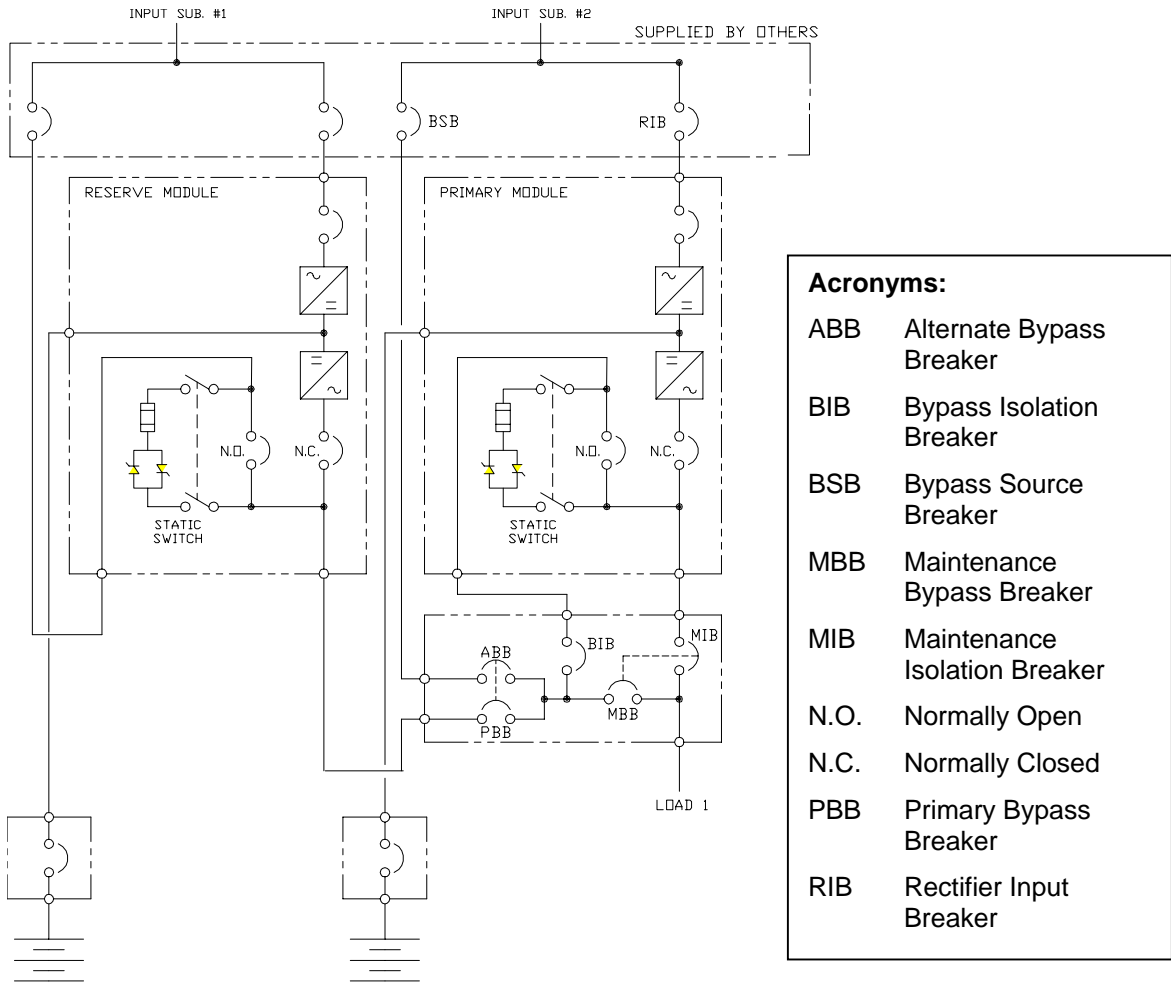


Figure 3: Simple Isolated Redundant System

The isolated redundant configuration has one or more single module systems designated as primary, plus an additional module as a reserve. Each primary module powers its own load bus. The reserve module is configured to serve as the first bypass power source for all the primary modules. See Figures 3 and 4.

This redundant configuration permits individual modules to be isolated for maintenance purposes without interrupting or disturbing the critical load. The operator can transfer the load from any primary module to the reserve module in order to perform preventive maintenance or to investigate an alarm condition. The entire system still needs to be de-energized for between one and four hours per year for system-level preventive maintenance. Like the parallel redundant configuration, the isolated redundant configuration is limited to about 99.99% availability.

Since the reserve module is the first bypass power source, all the primary modules synchronize their outputs to the output of the reserve module. If utility power input to a primary module should fail, the module will draw power from its batteries and continue to support its critical load without interruption. If utility power is not restored before the primary module exhausts its battery, the load will be switched to the reserve module and the primary module will shut down. The other primary modules will automatically switch their bypass inputs to the secondary bypass source, if one is available.

If a primary module should fail, the load will automatically be transferred to the reserve UPS module. Since the reserve module is no longer available as a bypass source to the other primary modules, these

modules switch their bypass inputs to the utility source. When the first primary module returns on-line, all primary modules switch back to having the reserve module as their bypass source.

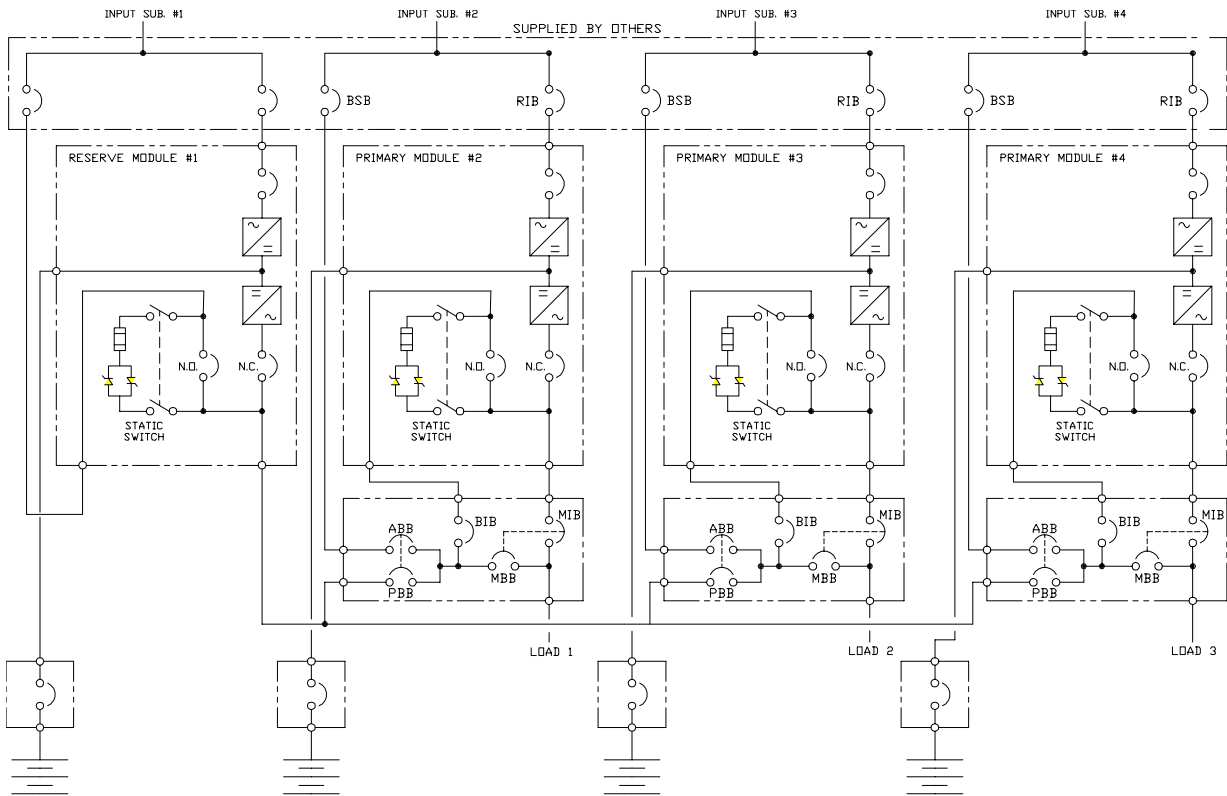


Figure 4. Four-Module Isolated Redundant

Three Benefits, Plus...

There are three key attractions of the isolated redundant configuration. First, the redundant UPS module permits taking another module off-line without putting the critical load to bypass. Second, the first bypass source is a UPS instead of the utility. Third, a customer can add redundancy to an existing UPS site by installing a new UPS as the first bypass source in an isolated redundant configuration.

All three of these factors mean higher availability of conditioned power than the single module configuration, and approximately the same as parallel redundant.

Five Cautions

The customer should be aware, however, of the many drawbacks of isolated redundant systems.

First, this approach requires more complex switchgear than parallel redundant systems. A two-module system (one primary, one reserve) requires at least one additional circuit breaker to permit choosing between the utility and the other UPS as the bypass source. A four-module system (three primary, one reserve, with three independent distribution networks, as shown in Figure 4) would require three separate maintenance bypass cabinets with the additional breaker in each one. This is several times the complexity of a parallel redundant system with a common load bus.

Second, system MTBF will tend to be *lower* for isolated redundant than either single module or parallel redundant systems. The redundant module helps availability, but adds *complexity in series* rather than parallel. Over the life of the system, the probability of a critical breaker failure becomes uncomfortably high.

Third, isolated redundant systems with two or more primary modules need a special circuit for selecting whether the reserve module or the utility will be the bypass source.

Fourth, the additional complexity is not cheap. The configuration can be relatively simple and low-cost for a two-module system, but the complexity mounts quickly in larger systems.

Fifth, each load bus is restricted to the capacity of a single module.

In summary, the isolated redundant configuration can improve system maintainability and fault tolerance. However, some isolated redundant configurations can actually worsen system reliability. It is most effective in two-module systems featuring one primary and one reserve module. It is less effective in systems with two or more primary modules, which are better handled in a parallel redundant, Distributed Redundant or Selective Redundant configuration.

Distributed-Redundant Systems: Continuous Availability

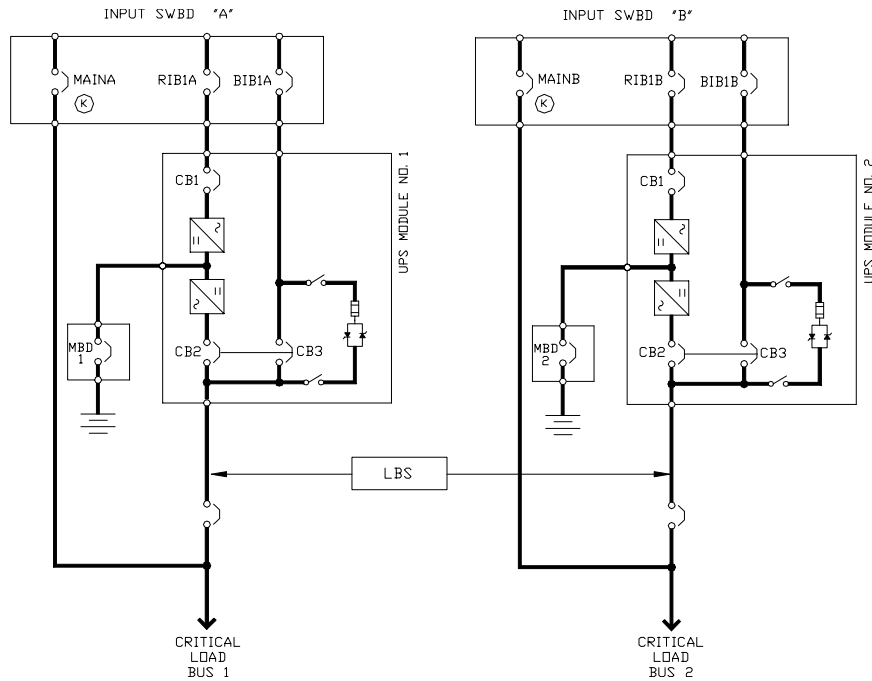


Figure 5. Distributed-Redundant System

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 distribution networks. See Figure 5 above and the separate paper on Distribution Options.

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.

Note: 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. The enabling technology, then, is the Liebert Load Bus Sync™ (LBS) option. The LBS option keeps the output of the two UPS systems within a few degrees (adjustable) of each other, even when the bypass is not available or the units are operating on battery. This means that downstream equipment can be switched transparently between the two sources. The Liebert Static Transfer Switch (STS), for example, can be used downstream to make manual or automatic transfers between the two independent, synchronized UPS systems.

Four Key Advantages

Distributed-Redundant systems provide a significant improvement in *maintainability and fault tolerance* compared to single module, parallel redundant or isolated redundant systems. With dual distribution networks, the user can switch all the downstream loads to one bus to permit taking the other UPS, breakers and distribution equipment off-line for maintenance or load reconfiguration. This means the owner can have a very active maintenance program without *ever* exposing the critical load to bypass power. Distributed-Redundant systems are capable of Continuous Availability.

Second, Distributed Redundancy can be simpler and less expensive than parallel-redundant or isolated redundant. The most-common approach is to begin with a pair of single module UPSs and the LBS option panel. This requires no system-level control cabinet or switchgear. It does not even require maintenance bypass cabinets for the individual systems, since each UPS system can function as the

“wraparound” bypass for the other. Distributed-Redundant systems are also scalable: for larger applications, the customer can specify a conventional parallel-redundant UPS for each side.

Third, the simplicity of the Distributed-Redundant approach makes it *inherently more reliable* than isolated redundant systems. The redundancy is in parallel, rather than in series. With two distribution networks and downstream static switches, the critical load equipment can have the functionality of “dual-power-cord” devices. There is less risk for the manager who schedules preventive maintenance. Best of all, the electrical independence of the distribution networks ensures that a load fault on one will not propagate to the other. The intact portion of the critical load can still be carried by the intact UPS and distribution system.

Fourth, the LBS option can be used to add a Liebert Series 600T or Npower UPS and a second distribution system to an existing Series 600, Series 300 or Npower installation. A special variant of the LBS can be used to sync a Liebert UPS to any competitor’s UPS. Another variant of the LBS -- the Extended Load Bus Sync or ELBS -- can be used to sync *three* or more Liebert UPS systems with complete functionality.

Why is the LBS option important?

The LBS keeps the output of two independent UPSs in sync, even when they are operating from two different power sources. 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.

The LBS option consists of an interface card in each UPS and a small wall-mounted panel with a simple selector switch. A minimum number of reference signals are brought out from each UPS to the LBS cabinet. There are no other connections between the logic or controls of either UPS, to ensure maximum system independence and isolation. The LBS remains dormant, content to monitor the output of both systems until their outputs drift apart by more than a predetermined amount. If the two systems begin to drift out of sync by more than a few degrees (adjustable), the LBS will activate and sync one unit to the other, according to the position of the selector switch. Once sync has been reestablished for five seconds, the LBS drops out until needed again.

The LBS concept is appropriate for virtually any three-phase power application, 10 through 6,000 kVA.

The Need for Double-Conversion UPS Products

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” by the IEC) can only sync to their own input power sources.

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

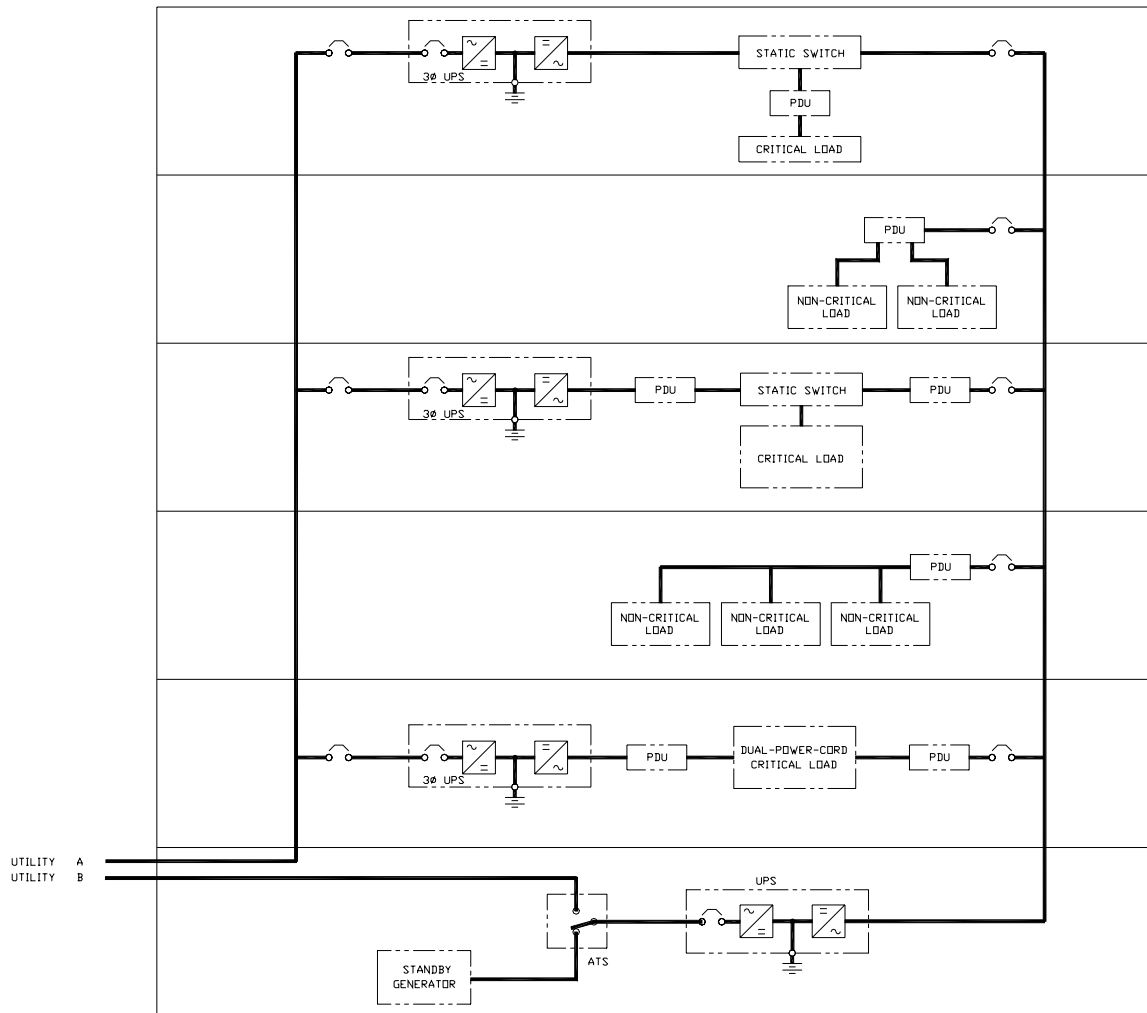


Figure 6. Selective-Redundant System

Some portions of your business are absolutely critical and need the 24 x 7 protection and maintainability of a dual-bus, Distributed-Redundant system. However, other parts of your business are less critical 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.

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 3x1000 kVA parallel-redundant UPS in the basement to power the entire facility. Then you could add a 2x750 kVA system for the data center, a separate 30 kVA system for the

seventh-floor servers, a separate 100 kVA system for the tenth floor tenant and Liebert Little Glass House or Nfinity single-phase UPS in the network closets.

Such a configuration creates three independent Distributed-Redundant Systems within the facility. The Liebert Extended Load Bus Sync (ELBS) option will keep all the 3-phase UPS systems in sync, even when one or more are operating on batteries. Individual loads in each critical work center can be switched transparently between the facility-wide UPS and the dedicated UPS.

Note: As a practical consideration, no UPS system should be more than 300 feet from the ELBS control box, so the maximum span between UPS systems is about 600 feet.

Power-Tie™ Systems: A Special Case of Distributed Redundancy

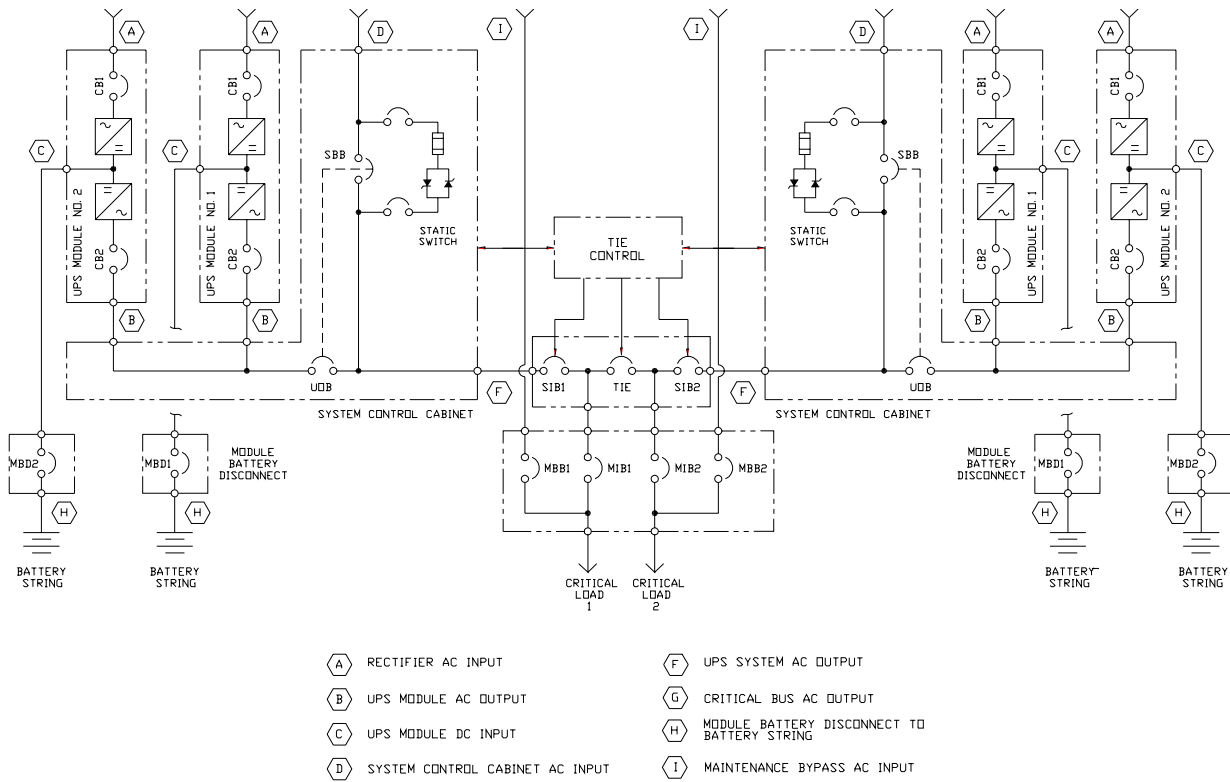


Figure 7. Power-Tie™ System

The Liebert Power-Tie concept features two independent UPS systems, each capable of carrying the entire critical load. A tie breaker and a simple control panel enable the two UPS systems to operate independently, or with either UPS system powering *both* distribution networks. See Figure 7.

The key advantage of this configuration is maintainability. The operator can transfer all the loads to one UPS system or the other at the push of a button, while the Power-Tie controls make the process foolproof. The controls use fiber optics (for maximum isolation) to transport the module sync signals between systems. Less-critical signals are transmitted through optical isolators and isolation amplifiers. In addition, all status signals have impedance protection between the signal wire connections and the system power supplies and ground.

Taken together, these steps ensure that a fault in any one UPS module or in the inter-system control wiring will not cause the critical bus or the other UPS system to fail. Even if several control wires of the Liebert Power-Tie system were severed or shorted together, the UPS systems would continue to function and support their respective critical loads.

The Power-Tie improves system maintainability by facilitating system-to-system load transfers without using the UPS bypass sources. The tradeoff is an increase in cost and complexity. In all other respects, the Power-Tie systems have all the features and benefits of other Distributed-Redundant UPS systems.

Hybrid AC-DC Power Systems: Optimized for Telecomm Equipment

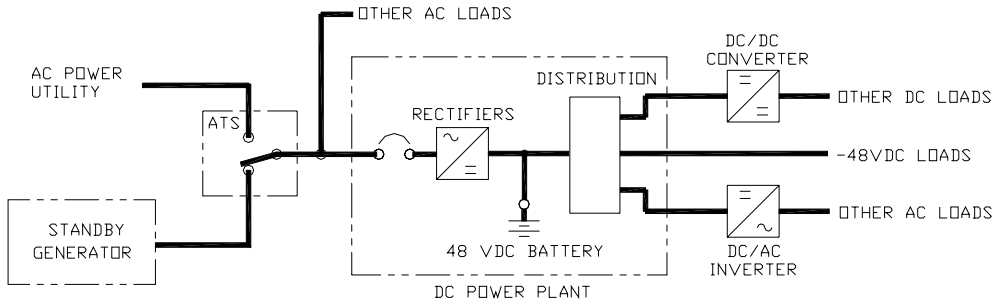


Figure 8. Traditional DC Power System

Most facility managers are aware of the growing amount of equipment that requires 48 volts DC input power. This includes telephone equipment and the equipment that connects to the Internet backbone.

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. See Figure 8.

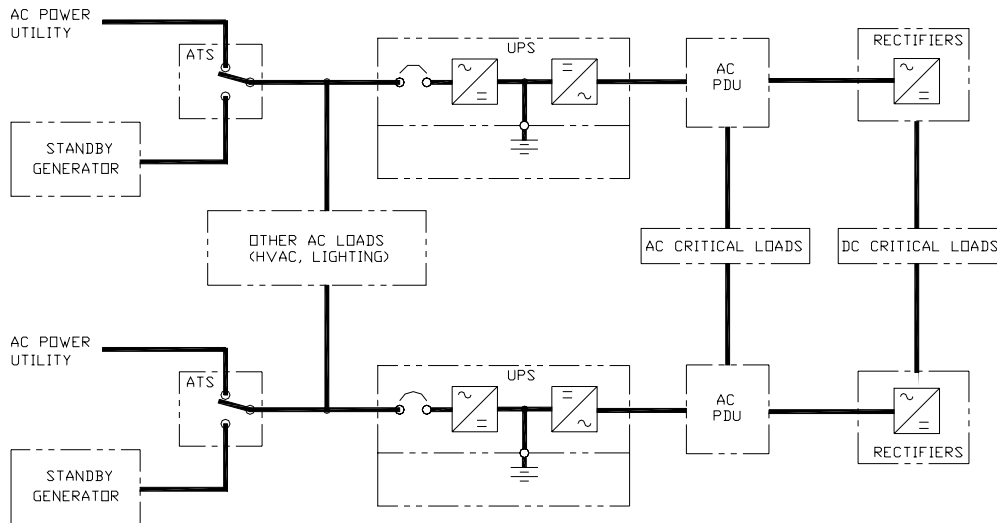


Figure 10. Hybrid AC-DC Power System

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. Liebert Hybrid AC-DC Power Systems make it a peaceful coexistence.

Liebert 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 10.

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. Your critical load deserves nothing less.

Conclusion

Each of the system configurations described in this paper has its own advantages and disadvantages. The single module system will continue to be the simplest, least-expensive approach for applications that can comfortably schedule downtime for preventive maintenance and system alterations. The parallel redundant system improves system availability and reliability and simplifies maintenance of individual UPS modules. The isolated-redundant configuration enables the facilities manager to add redundancy to an existing single module installation. The Distributed-Redundant, Selective-Redundant, Power-Tie and Hybrid AC-DC systems provide Continuous Availability of conditioned power.

The choice is yours.