

The Anatomy of a High Availability System

Configuring the New Data Center for the Ultimate in Systems Availability

By Bob Bauer
President, Liebert Americas

For the past several decades, Market Intelligence has monitored the shift away from early mainframe orientation toward an exploding PC - and workstation - based market. Yet today, we are experiencing a re-shift back to a centralized environment. The migration back to a more protected environment is driven by a growing demand for systems availability, often nothing less than 7x24.

Demands on servers, e-commerce/internet, and data/telecommunications equipment have made the "no down time" operating parameter the goal. Yet everyone involved in developing, specifying or maintaining systems in today's information-intensive enterprise is aware of a growing problem. The platforms on which these critical systems reside have not necessarily been designed for maximum availability. The challenge for today's data center is to develop an infrastructure that provides the old glass house safeguards, including environmental controls, power protection, cable management and security. The new center also should incorporate the latest advances in site monitoring and deliver a redundant system that allows for on-line maintenance and repair without impacting the processing operation. Also, a modern center should accommodate upgrades and expansions while protecting as much of the initial investment as possible.

In short, a more holistic approach should be followed that takes into account the appropriate protection equipment, site evaluation and proper maintenance. How to identify and configure a data center protection strategy that suits the organization's processing demands and your systems availability requirements is the subject of this article.

The Real Cost of Downtime

How much does it *really* cost your organization when the network crashes due to power glitches? While an exact figure is hard to compute, the real loss is probably much higher than you think. Consider just these few real world examples of the difference between 99.9999% uptime and an environment that demands 7x24 protection.

Some Illustrations of 99.9999% Reliability:

- One hour of unsafe drinking water
- Two unsafe plane landings per airport per day
- Twenty thousand incorrect surgical operations per week
- Twenty-two thousand checks deducted from the wrong account every hour
-

These examples sound extreme because of the magnitude of volume of each of these categories and their importance. Tomorrow's volume of data as the information and e-commerce economy explodes makes it a significant parallel to draw today. Though 99.999% availability may sound acceptable, it is important to understand why proper power protection should be as an integral part of the data center and chief among the critical aspects of overall system availability.

The Evolution of UPS Applications

Today, power conditioning and UPS are available for everything from a single PC to all the sensitive equipment housed in an entire building. There are UPS products designed for specific applications, such as those to protect computer and telecommunications networks — whether it's a single PC or thousands of nodes spread across the globe. In fact, there's a protection product to fit every network power need. This extends from surge protection and power conditioning to three-phase UPS for mainframe database servers.

Three major UPS technologies available today are off-line, line interactive and on-line. The off-line UPS, also called a standby UPS, is a cost-effective choice for small, non-critical stand-alone applications such as isolated PCs and peripherals. In an off-

line or stand-by UPS, the inverter/converter turns on or “switches” in order to make battery power available. Line-interactive technology provides highly effective power conditioning plus UPS back up. This is most applicable in areas where power outages are rare, but where there is frequent power fluctuations. During normal operation, the line-interactive UPS filters the line current to the load and converts a trickle of DC to keep the battery fully charged. When the power fails, the UPS transfer switch shifts from utility to battery.

For the kind of applications housed in most data centers, a more complex power configuration is needed. This, in turn, dictates on-line UPS technology. As opposed to offline or line-interactive UPS, an online system eliminates a wider range of potential power problems such as spikes, surges, including voltage and frequency variations common with standby generator operations. This is because a true on-line device uses an extended battery back-up or standby generator to deliver power to an attached device during an outage. The on-line UPS *continually* recharges the battery, so that if power goes down, backup is immediately available.

The latest on-line UPS systems ensure maximum system availability. To assure that UPS supplied power is always available, even during routine maintenance, large UPS systems are configurable with synchronized busses to allow one UPS to be taken off line without disrupting the critical load. Many UPS systems are equipped with extended battery time and generator capability to ride through long outages. And more and more UPSs, just as the information technology they protect, are being designed for greater power availability. For example, the latest servers are produced with “built in” redundancy capabilities via dual power cords; UPS devices are now being designed to complement this with dual input capabilities.

Your users and your environment will determine the need for any additional features. Different applications have different requirements. With a network server, a graceful shutdown may be all that’s required. A telecom system, however, will need longer autonomy times that require extra batteries, or a standby generator. Telecom systems also require a more comprehensive power strategy.

Yet any manufacturer who claims a single power protection product will assure 99.999% availability for the mission critical networks and computer systems that run a data center simply isn’t telling the whole story. What’s omitted: there is far more to assuring computer uptime than just adding one or two strategically placed back-up power supplies into the room.

The Show Must Go On: Configuring for 7x24

Total protection against all threats is the only way to assure near 100% uptime of mission-critical operations. The threats that your systems face are many and extend far beyond power to encompass environmental factors, equipment and software reliability, security issues and disaster recovery. While the ability to ride through power outages is important, it’s only one slice of the critical protection pie. Truth is that “dirty” power with sags, surges, harmonics and other aberrations are far more frequent and damaging to equipment than outages.

The sensitivity of today’s computer hardware and software also needs to be considered. These highly sophisticated electronics and digital controls demand equal protection from hazards found in the surrounding environment – including temperature and humidity fluctuations as well as airborne dirt, dust, and other contaminants.

Whatever the cause of the outage – improper environmental control, a power outage, dirty power – true continuous availability can only come from an analysis of all the threats and a protection strategy to address them.

Analyzing the Threat

You should begin by looking at potential points of failure within your existing system. This is not as straightforward as it sounds. Of course, the data center was designed in order to protect against many common sources of downtime, yet there are other, more “hidden” threats that may be overlooked.

Even a simple LAN configuration can transmit electrical noise among all nodes on a network, so it’s critical to protect each node in order to protect the data stream. Overload neutral conditions are becoming common as more PCs are added to office environments. Harmonics generated by PC power supplies do not cancel on the neutral line. These harmonics lead to a RMS current in the neutral line approximately 1.7 times that of “hot” conductors, leading to potential overloads.

Switching devices, routers and modems as well as other connectors need to be taken into consideration, too. Any power assisted cabling devices that are not protected can go down during a blackout. The results for the users are the same: they won’t be able to use the network, even if their file server and nodes have power support. The same is true for Wide Area Network gateways. A PC-to-host session can be interrupted if the gateway goes down and the host database may be corrupted by the interruption.

Again, you also must consider the environmental “threats” that affect the sensitive equipment housed within the center. The total cooling load within this critical space, for example, can impact equipment performance. This is a combination of the equipment heat load in the room; the number of people in the room; their migration into and out of the room; infiltration and the effects of windows (the solar heat load) or wall construction; plus the load resulting from bringing in outside air for pressurization. Controlling relative humidity is equally important. Too much moisture in the air can corrode switching circuitry while excessively low humidity can cause static electricity.

Before installing sensitive electronics inside a data center, you need to conduct a site audit that includes other, less obvious factors that contribute to interrupted service. These may include ground impedance factors such as the type of soil, the humidity, and the methods of grounding used. Proper grounding is, quite literally, the foundation of a good power plan. Take for example, the ground wire used for all electronic systems as a signal reference. If this is in any way compromised, the potential is there for much more than signal confusion but an actual systems malfunction.

Measuring by Redundancy

After you’ve determined the threats, you need to consider the way in which your system of choice is configured. System configuration – even within a controlled environment such as a data center – remains the critical difference in availability. MTBF or Mean Time Between Failure has become an industry-wide measurement of product reliability. Redundancy objectives serve as the best realistic predictor of your necessary level of *availability*. These objectives are formed around the question: how much protection is enough? At this level, the focus shifts from a protection “product” to a protection “configuration” that not only provides for power availability in the event of failure, but also switchover capability for routine service. For ultra-critical loads, the power system needs to be about 10 times more reliable than the load – *and redundant* – to avoid compromising the initial investment as well as the overall business plan. This almost certainly points to the most substantial form of power protection, a distributed redundancy configuration.

Just what is distributed system redundancy as it relates to power protection? In its basic form, distributed redundancy involves creating two (redundant) UPS system busses and redundant power distributed systems. This eliminates as many single points of failure as practical, all the way up to the load equipment’s input terminals. In order to provide “fault tolerance,” some method of allowing the load equipment to receive power from both UPS power busses must be provided. To protect against fast power system failures, such as circuit breaker trips or a power system fault, you need a commensurately fast switching method.

Static transfer switches (STSs) have been applied to accomplish very fast break-before-make transfers between two AC power sources. It is important that the two AC power sources be designed as independent as practical to eliminate any common failures. Switching between the two power sources needs to be break-before-make for the same reason. A number of distributed redundancy power distribution configurations can be devised. Keep in mind, however, that redundancy needs to be as close to the load as possible to achieve its goal – namely, keeping power available at the load equipment level.

The Critical Nature of Environmental Factors and Controls

As mentioned, the security of the “glass house” computer room also extends to environmental protection – a factor vital to the proper operation of sensitive electronics. Computers have changed, but one thing remains constant: excessive heat or humidity can damage or impair the operation of critical computer systems and peripherals. Whether it’s a large mainframe center, a network and web site servers, or an ancillary room housing minicomputers for R & D – sensitive computer and telecommunications equipment simply works better when properly cooled. Assuring maximum availability of computing systems cannot be fully achieved unless environmental protection considerations are factored into the equation.

Of course, as information needs grow with stunning speed and unpredictability, environmental factors – along with power quality – become part of a much more complex issue in the modern data center. The latest environmental solutions permit the same level of centralized protection in vulnerable areas like an open office or a factory floor. These solutions can range from a lockable enclosure that acts as a virtual airtight and soundproof “vault” for critical equipment to modular systems that provide cable management, air conditioning and power protection over flooring, through ceiling panels, even by mounting on the wall.

Ancillary Equipment: Devices that Make a Difference

The ability to have real-time knowledge and control over the quality of the air and power itself can only come from connectivity and communications from making the protection equipment in the center an active part of your overall facility management process. This starts with building into each piece of equipment a basic communications capability. These capabilities range from simple remote monitoring of a single environmental or UPS unit to an integrated communications system that can oversee power, environmental and security points – all monitored from a single location.

The latest power communications software and hardware provide multiple communications options for each level of power protection. Levels range from a simple shutdown interface to comprehensive, in-band SNMP-based software that goes beyond power protection and control to oversee air conditioning and up to ten user-programmable inputs. Keep in mind the latest UPS communications are also available in a “redundant” form – designed to provide multiple solution paths in the event of a problem. The latest in-band/out-of-band redundant communications strategy is a prime example. In this strategy, the UPS provides out-of-band communications separate from the network wire, assuring emergency contact with administrators or the equipment manufacturer even in the event of network failure.

In addition to standard control consoles and remote monitoring units, you now have the ability to perform wide-area oversight. Designed for large, complex computing and telecommunications systems, these provide real-time monitoring and alarms from any piece of analog or digital equipment. The system can monitor power, temperature and humidity, smoke or water detection, and security ... everything from temperature at a remote telecommunications shelter to battery status on the starter motor of your standby generator.

Beyond the Box: the Importance of Service and Support

In order to provide a high level of customer service, many mission-critical enterprise systems must be able to communicate information not just within a single facility, but across the online business enterprise anywhere in the world. So the computer system in an office in Toledo, Ohio, or Paris is just as vital as the computer on the 80th floor of the World Trade Center. Yet the application requirements may vary greatly.

Specifying and maintaining these differing requirements with a consistently high level of availability requires access to a number of support services. If you're like most data center managers, you are already burdened with too many tasks and too little staff, so you may want to consider going outside for service and support. Yet any of these after-sale functions, regardless of whether they are staff or contract maintenance, should include system sizing and configuration as well as an objective assessment of the data center in which the critical equipment operates.

A data center manager also must be a savvy planner to prevent the loss of service, performance and investments in technology. A proactive maintenance program is critical to long-term effectiveness of any program designed for system availability. This includes everything from fulfilling product warranties to the total management of the entire protection system including UPS, power and precision air conditioning, batteries, switchgear and generators.

Contracting support services can help ensure 7x24 oversight as well as provide critical after-sale training, maintenance, repair and analysis. The service provider must be able to respond quickly, so single-source accountability is key. Look for a supplier with access to a large product line and the local expertise to tailor a solution whether the air conditioning system requires low ambient operation, or a high-rise type condenser or 50 Hz power.

Whoever said, “ignorance is bliss” wasn't responsible for the operation of a data center. Small problems can crop up from nowhere and suddenly become big problems if no one is aware of what's happening. For protection to be truly without fail, you must move beyond the reliability of any single power protection device to developing a high availability system. With a solution in place that addresses power, air and servicing, you can be assured of 7x24. For today's mission-critical applications, these are benefits that reach all the way to the bottom line.

###

[LIEBERT WEB NOTICE AND CONDITIONS](#)

Copyright © 1995 - 2000 Liebert Corporation.

For more information, contact webmaster@liebert.com.