

The Truth Behind Redundant Data Center Power

by Paul VanMeter and Ricky Vasquez

Introduction

Reliable, redundant power is one of the most critical values that a data center can offer its customers. Yet what most data centers offer is some form of backup that falls short of truly redundant power availability.

For example, multiple Seattle, Wash., data centers went offline for several hours in July of 2009 as the result of a fire in an electrical vault, leaving the city's Fisher Center data center complex without power and generator support. As a result, a leading credit card payment portal was offline for hours, leaving thousands of merchants unable to process payments through their websites, along with interrupted service for other major websites, web hosting providers and dozens of others that depended on the data center availability to support their businesses.

According to the Ponemon Institute's Sept. 2010 National Survey on Data Center Outages, 95% of 452 data center managers had experienced unplanned outages in the previous two years, with respondents averaging 2.48 complete shutdowns in the previous two years, with an average duration of 107 minutes. Nearly 88% of those surveyed said they had experienced loss of primary utility power in the previous two years.

In addition to outages from a primary utility, other points of electrical failure can be at the transformer, the automatic transfer switch, the uninterruptible power supply and independent power distribution units.

With businesses relying on their data centers for almost all of their daily operations, regardless of the point(s) of failure, outages are costly. The exact cost depends on the business. Pingdom, a company specializing in uptime monitoring, estimated the cost of an outage that hit a major electronic payment provider in 2009 at between \$7 million and \$32 million.

Outages will happen, from either "Acts of God," like the storm system that knocked out power in cities from Chicago to Dallas in early February of 2011, a backhoe cutting a power line, simple equipment failure, an electrical fire or from a host of other factors. The only guarantee is that a failure of primary power is likely to occur. But customers don't care if there were storms, rolling blackouts or some other issue causing an outage. They expect unflinching delivery of data center services.

2N vs. N+1

Knowing this, data centers rely on redundant power with generators and backup systems to continue operations. But all redundant power systems are not created equal.

Some data centers offer “N+1” electrical system redundancy. While this provides some element of assurance of continuing operations in the event of a power interruption, N+1 systems have their own weakness that make them inferior to fully redundant, or 2N (sometimes referred to as A/B) systems.

N+1 systems feature redundant equipment, so if any single piece fails, there is a backup. But the weakness of these systems comes from designs that include common circuitry or feeds at one or more points, rather than two completely separate feeds. Equipment, including circuitry, is man-made. It will fail at some point.

So despite all of the additional money invested in backup generators, UPS systems, switches, transformers and other systems, this common circuitry can fail, resulting in interruption of service. Therefore, an architecture offering two completely separate systems, with no single points of failure, is far superior to a single system with redundancies.

Operating During an Outage

While some data centers with low loads can rely on a UPS (and a backup in case the primary UPS system fails) to handle power for a short outage, those systems alone are not sufficient to handle the growing emergency power needs of today's systems.

So in addition to the UPS, most data centers will rely on a backup generator in the event that



the electricity from the local utility falters even partially, as in the case of a brownout, or fails completely. Generators will provide emergency power, but typically only for a limited period of time. Key considerations are the generator's capacity both in terms of power and of fuel. If the facility is cut off from fuel supply for an extended period, how long can the generator operate at top capacity? For example, a 2MW generator, powerful enough to power 1,000 typical homes, with sufficient fuel capacity can run a data center for several hours. But unless a data center has large fuel reserves of several thousand gallons, the fuel supply could run out before utility power returns, particularly in the event of a flood or severe storms, when power can be out for several days.

Even a generator with a sufficient fuel capacity may not be enough to provide truly redundant power. The data center should have a second generator as a backup in case the primary generator fails. Both should be out of harm's way. Even if the data center is in an area that rarely

floods, the generators should sit well above flood levels. The generators should also be routinely tested and inspected.



Similarly, data centers must ensure that generators properly synchronize with UPS systems so there is no interruption of service when moving to generator power. Lack of proper synchronization and failure in portions of the electrical infrastructure preventing transfer of the electrical load between different power sources were cited by a leading data center as two of the reasons for an outage in June of 2009.

Smaller Footprint Equals Lower Costs

Beyond a data center's power availability and reliability, another essential consideration today is the distribution of power throughout the facility. Many data centers were designed to provide only 75 to 100 watts per square foot. But today's high-density racks and servers need much more than that, often two times as much. One choice is to rent out additional space in a data center, meaning paying for space that's unneeded just to gain the additional power. The other option, as many have found, is to move to a facility that was built to offer 200 watts or more per square foot and allows for cabinets that operate at 8KW or more as well as a variety of circuits, including 110v, 208v and three-phase 208v, which provides the customer with more

versatility regarding equipment. The larger the floor space, the more complex the cooling system. By putting more equipment in a smaller footprint, a customer will pay less for his space, a benefit that will recur on a monthly basis. Facilities with lower power availability may restrict power set up due to the density.

Better Maintenance Availability

All the UPS systems should keep equipment powered seamlessly, preventing resets while generators kick in. Beyond the protection against outages, the 2N design provides the data center with much easier maintenance scheduling. Any unit can be pulled off line and repaired or replaced without fear of losing power to subsequent systems. According to a commentary by Kenneth Brill, executive director of the Uptime Institute, in *Forbes*, the top reason for catastrophic electrical failure is the lack of electrical maintenance.

"Electrical connections need to be checked annually," Brill wrote. "Many sites cannot do this because IT's need for uptime and the facilities department's needs for downtime are incompatible." A 2N architecture avoids this conflict so that facilities can conduct preventative maintenance on a regular basis.

Conclusion

Reliable power is one of the most critical factors in the ongoing, reliable operations of a data center. Data centers that can offer truly redundant, 2N power and high density connections offer their customers a critical value proposition.

If considering a data center, it is important to carefully inspect the design to ensure that

power is readily available and that fully redundant (2N) systems will deliver performance even in the event of an extended power outage either due to the failure of an internal component, like a UPS system, or an extended outage from the local utility.

Even though a data center may promise redundant or backup systems it is important to verify exactly what that promise entails. Physically inspect the entire electrical system design—from the utility feed to the final delivery of power to equipment—to ensure that separate, independent power feeds are indeed being delivered. This will ensure that if any one system fails, there is no loss of power to equipment.

About the Authors

Paul VanMeter is CTO and Co-founder of Colo4. VanMeter has been recognized as a leader in the colocation service industry, spearheading the successful growth of Colo4 since 2000. With more than 20 years of experience in the colocation and telecommunications industries, Paul brings together facility and connectivity resources to support companies that depend on technology and server-based systems.

Ricky Vasquez is Director of Facilities at Colo4. Vasquez is responsible for daily oversight and management of Colo4's data center to ensure the ongoing flawless operation of the facility. With over a decade of industry experience, Ricky leads Colo4's Facilities team, which is responsible for maintenance and operation of the facility, including electrical systems, mechanical systems, security and electrical contractors, janitorial services, and all other aspects of physical operations.

Colo4
3004 Irving Blvd.
Dallas, Texas 75247
214.630.3100
www.colo4.com