A White Paper from the Experts in Business-Critical Continuity™

Addressing the Leading Root Causes of Downtime:

Technology Investments and Best Practices for Assuring Data Center Availability

EMERSON™
Network Power

Executive Summary

Today's data center has evolved into a strategic business asset at the core of business performance and customer satisfaction. However, as the criticality of data center operations continues to increase, so too do the financial and intangible costs of a downtime event.

While availability remains a key concern for CIOs, some underestimate the true business impact of unplanned outages, focusing instead on reducing CAPEX/OPEX while simultaneously increasing their data center's throughput. As a result of this miscalculation, many companies do not have technology investments and best practices in place to adequately prevent and/or address the leading causes of downtime – this according to the "2010 National Survey on Data Center Outages," an independent study conducted by the Ponemon Institute. In addition to providing insight into industry perceptions of data center availability and downtime events, the report uncovers the most frequently cited root causes of preventable outages, ranging from failures of critical data center equipment to human error and accidental "emergency power-off" (EPO) events.

This white paper will examine the most frequently reported root causes of data center downtime and recommend cost-effective solutions, design strategies and best practices for eliminating these vulnerabilities. These recommendations are designed to significantly reduce the frequency of unplanned data center outages while simultaneously addressing key business concerns for CIOs, including improving energy efficiency, reducing total cost of ownership and maintaining end-user satisfaction. This paper also will explore the value of proactive service – including preventive maintenance and data center assessments – and how these proactive investments can be used to identify subtle weaknesses in a business-critical infrastructure, and optimize a facility for performance and efficiency.

Background

For companies that rely on data centers to deliver business-critical services, downtime always has been a key concern. However, as the effects of the economic downturn began to take their toll on once-thriving enterprises, achieving cost reductions through energy efficiency began to overshadow the importance of availability in the eyes of senior management.

According to a 2009 survey of the Data Center Users' Group (DCUG), efficiency was cited as a primary concern by more than 47 percent of data center professionals polled, making energy savings the No. 2 concern overall. Unfortunately, the increased focus on rapid cost-cutting and energy efficiency left many business-critical data centers at an increased risk for equipment failures and unplanned downtime events.

While companies began adopting high-density configurations, virtualization and other strategies intended to boost the capacity of existing IT equipment, many overlooked the need for a robust data center infrastructure to ensure continuity for business-critical applications. As a result, a number of high-profile outages across a variety of industries proved to be more costly than investments that could have prevented them altogether. In addition to an overall disruption of service and (in some cases) loss of customers, these outages translated to hundreds of thousands of dollars in financial losses and future customer business.

According to the Ponemon Institute's "National Survey on Data Center Outages," **95 percent** of companies have experienced an unplanned downtime event within the past two years. Types of outages cited include total data center outages (2.48 events on average), partial data center outages (6.84 events on average) and device-level outages (11.29 events on average). However, even though more than 60 percent of these companies rely on their data center to generate revenue or support e-commerce activity, less than 35 percent believe their data center utilizes critical system design and redundancy best practices to maximize availability.

Furthermore, the majority of data center professionals do not believe they have enough resources to respond quickly to an unplanned outage or failure, with the average downtime event lasting nearly two hours (107 minutes) per outage. When examining the prevalence of downtime in specific industries (Figure 1), data centers serving the healthcare and civic/public sectors accounted for the longest average duration of downtime (with an average of 3 and 2.84 annual downtime events, respectively), followed closely by the industrial and financial sectors. This is particularly alarming considering the criticality of electronic health records, financial services and public information systems dependent on the availability of these systems.

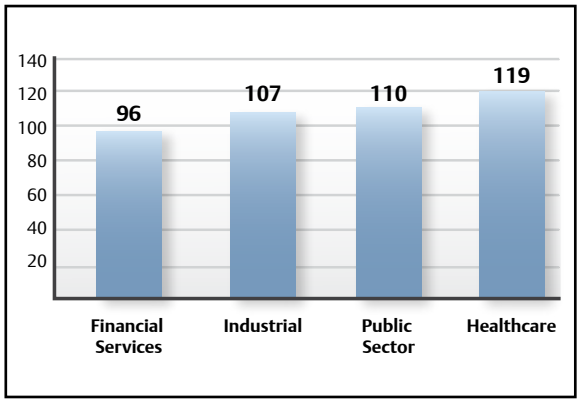

Figure 1. Extrapolated duration of unplanned total data center shutdowns by industry segment.

To gain a better understanding of which vulnerabilities contribute to such a high occurrence of costly downtime events, the survey asked more than 450 data center professionals to cite the root causes of data center outages experienced during the past two years. The vast majority of data center downtime was related to inadequate investments in a high-availability infrastructure. In fact, more than half of data center professionals polled agreed the majority of downtime events could have been prevented.

As shown in Figure 2, the findings uncovered the seven most frequently reported root causes of downtime, ranging from UPS and cooling equipment failures to human error and accidental shutdowns.

The implementation of cost-effective solutions, design strategies and/or best practices can enable data center managers to reduce or eliminate the risk of these root causes while simultaneously improving energy efficiency, flexibility, total cost of ownership and end-user satisfaction.

UPS Battery Failure

While batteries may be the most "low-tech" components supporting today's mission-critical data centers, battery failure remains the leading cause of unplanned downtime events. In fact, studies conducted by Emerson Network Power indicate battery-related failures account for more than one-third of all uninterruptible power supply (UPS) system failures over the life of the equipment.

The continuity of critical systems during a power outage typically is dependent on a data center's power equipment, comprised of UPS and their respective battery backups. While the vast majority of outages last less than 10 seconds, a single bad cell can cripple a data center's entire backup system – particularly if adequate UPS redundancy has not been implemented.

All batteries have a limited life expectancy, dictated by the frequency of battery discharge and recharge. However, there are a number of factors that can impact the aging process and shorten a battery's useful life, including high ambient temperatures,

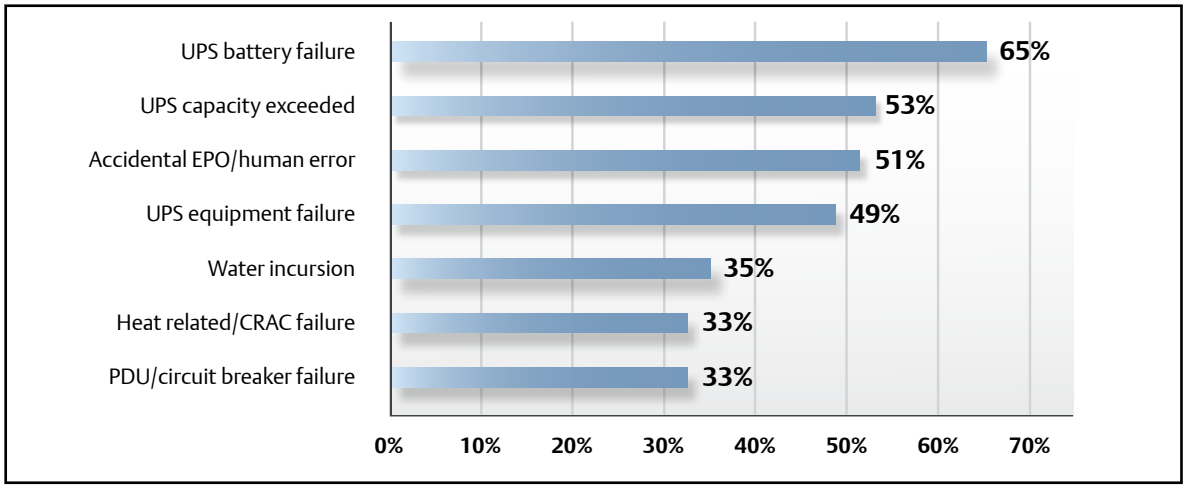

Figure 2. Top root causes of downtime cited by data center professionals. Totals exceed 100 percent because data center professionals cited multiple root-causes over a period of two years.

frequent discharge cycles, overcharging, loose connections and strained battery terminals.

To safeguard backup power systems against unplanned or premature battery failures, data center professionals should take steps to ensure battery maintenance best practices are observed. In addition to standards outlined by the Institute of Electrical and Electronics Engineers (IEEE), manufacturer schedules for maintenance checks should be followed to ensure batteries are properly maintained (correctly installed, fully charged, etc.), serviced and/or replaced before they pose a risk for mission-critical applications.

As outlined in Figure 3, monthly preventive maintenance best practices for battery systems include visual inspections (internal and external), acceptance testing and load testing. Capacity tests also should be performed by a trained service technician at recommended intervals to measure the degradation of a battery over time.

In addition to regular service and preventive maintenance, the installation of an integrated battery monitoring solution enables data center professionals to proactively monitor the performance of individual cells across all UPS systems in the data center 24/7.

Monitoring solutions provide comprehensive insight into battery health – including cell voltage, resistance, current and temperature – without requiring a full discharge and recharge cycle. This allows batteries to be utilized to their maximum effectiveness, safeguarding data center professionals against premature replacements as well as unanticipated battery expirations. In addition to robust on-site monitoring capabilities, integrated solutions typically offer predictive analysis and provide web-based remote access to facilitate rapid service response before battery failure occurs.

While proactive battery monitoring and maintenance are critical to maximizing UPS availability, data center professionals also

Recommended Task	FLOODED IEEE 450			VRLA IEEE 1188		
	Monthly	Quarterly	Annually	Monthly	Quarterly	Annually
Battery system voltage	●			●		
Charger current and voltage	●			●		
Ambient Temperature	●			●		
Visual inspection	●			●		
Electrolyte levels	●					
Pilot cell voltage and specific gravity	●					
Specific gravity all cells		10%	●			
All cell voltages		●			●	
All cell temperatures		10%			●	
Detail internal visual inspection			●			
AC Ripple Current and voltage						●
Capacity test			5 Years			●

Figure 3. Monthly preventive maintenance best practices for battery systems include visual inspections (internal and external), acceptance testing and load testing.

should keep charged spares on-site to cover any cells that may have expired between service visits. Depending on the criticality of the data center, data center professionals should have enough batteries on hand to replace between 5 and 10 percent of batteries in all cabinets. Because these batteries are fully charged and ready for rapid deployment, service visits to repair battery strings can be significantly reduced –cutting costs without compromising availability.

Exceeding UPS Capacity

High-density configurations have become more common in recent years as data center managers seek to achieve increased throughputs from their existing infrastructures at the highest efficiencies possible. In fact, the average power draw for an individual rack can exceed 10 kW in some high-density data centers. With such high densities common during peak hours, the capacity of a single UPS system can be quickly exhausted, leaving critical IT equipment unprotected and the UPS system at risk for overload and failure.

According to the Ponemon Institute, more than half of all downtime events were the result of exceeded UPS capacity, with outages ranging from individual rack-rows to total data center shutdowns. Because the UPS is the first line of defense in a data center's power infrastructure, critical IT equipment must be backed by adequate UPS protection to ensure that all systems will be fully supported.

If operating under a single-module UPS configuration, it is important to have a thorough understanding of the typical loads experienced throughout the data center. By measuring output multiple times per day via an integrated monitoring and management solution, data center professionals can gauge the typical power draw of their IT equipment over time to confirm whether their infrastructure is backed by enough UPS capacity. Some solutions, such as the Aperture Integrated Resource Manager, also give data center professionals the ability to scale UPS systems automatically at the rack-level and shift loads in real-time based on available capacities.

Establishing a redundant UPS architecture also enables data center professionals to increase capacity of their backup power system, with the added benefit of eliminating single points of failure. Parallel (N+1) redundancy remains the most cost-effective option for high availability data centers and is well-suited for high density environments where future growth is certain.

In a parallel redundant system (Figure 4), multiple UPS modules are sized so enough modules are available to power connected equipment (N), plus one additional module for redundancy (+1). In these configurations, all UPS modules remain online and share the load equally. This enables data center professionals to design their redundant UPS systems with additional capacity to accommodate increased, non-typical loads resulting from equipment failures, modules being taken offline for service and rapid power fluctuations common in virtualized data centers. However, in order for this configuration to provide adequate protection, it is critical to ensure that the total IT load **does not** exceed the total capacity of N UPSs.

Because N+ 1 configurations require a dedicated static switch, the facility manager cannot simply "add on" modules in the future to increase capacity. Therefore, it is a best practice to maximize the size of UPS systems used in an N+1 configuration based on projected capacity needs rather than

investing in multiple small UPS systems. In addition to being more cost effective from a long-term CAPEX perspective, limiting the number of modules in a parallel redundant system minimizes service costs over the life of the equipment and is more energy efficient. An alternative option is a 1+N configuration, in which specialized UPS modules with integrated static switches are paralleled via a paralleling cabinet, enabling data center professionals to "add-on" UPS modules as capacity needs change.

UPS selection also should be a consideration when seeking to minimize the occurrence of UPS capacity overload. Many vendors, including Emerson Network Power, have introduced intelligent UPS systems designed for rapid, seamless deployment in critical IT environments.

In addition to high operating efficiencies (up to 97 percent efficiency through the integration of an "always-on" inverter), some intelligent UPS solutions have significant overload capacity built-in and are able to handle bursts of more than 25 percent above the UPS's total capacity.

Many intelligent UPS systems, including the Liebert NXL, also are capable of achieving superior performance and availability through redundant components, fault tolerances for input currents and integrated battery monitoring capabilities.

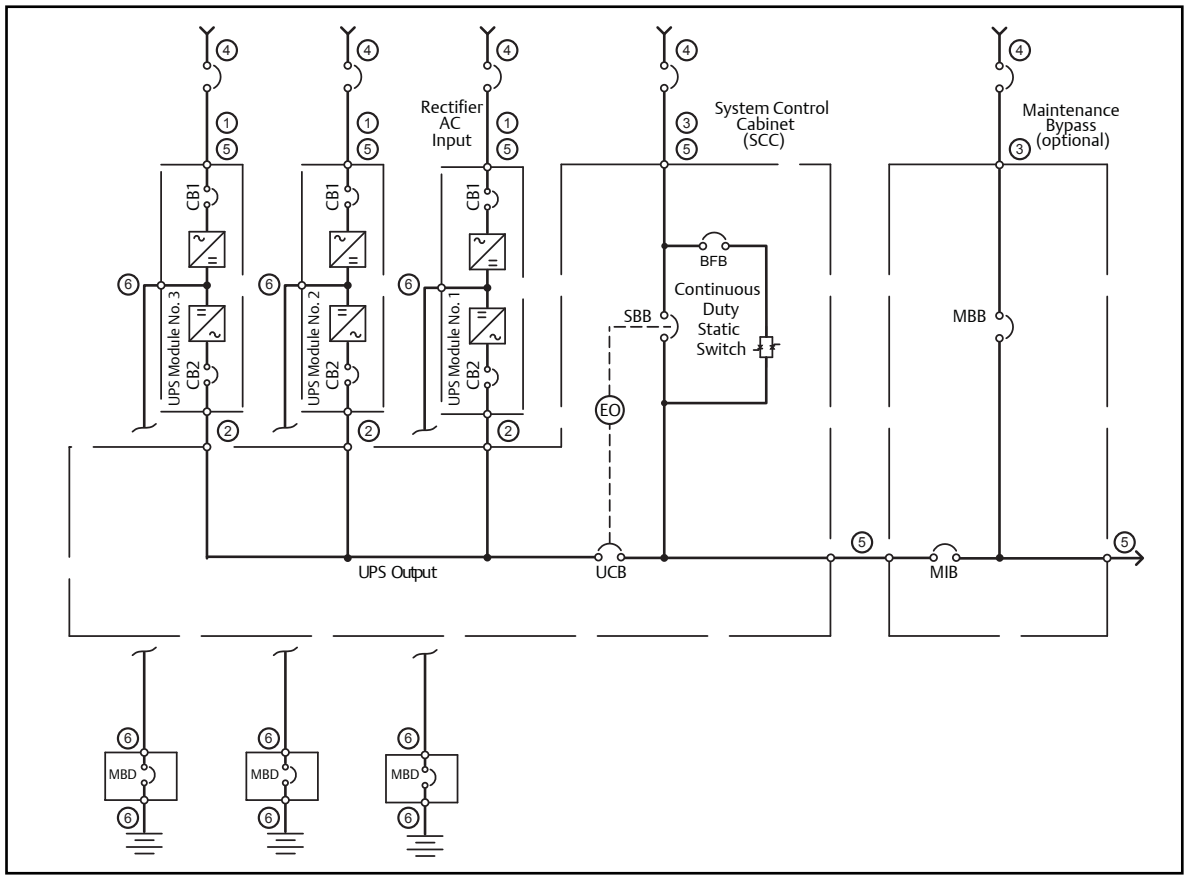

Figure 4. A typical parallel-redundant system configuration.

UPS Equipment Failure

In addition to exceeding capacity, general UPS equipment failure is another leading cause of downtime cited by survey data center professionals, with 49 percent reporting an equipment failure within the past two years. With this in mind, it is important to consider a number of key factors that can affect the reliability and overall lifespan of UPS equipment.

System topology is a significant consideration when evaluating UPS reliability. Generally speaking, two types of UPS are found in today's data centers: line interactive and double conversion. While double-conversion UPS systems have emerged as an industry standard in recent years, many legacy and/or small data centers may still be using line-interactive systems in their critical infrastructures. Therefore, when evaluating each topology, data center professionals should consider the criticality of their data center operations.

In addition to distributing backup power to the load, the UPS battery in line-interactive systems conditions the power before it flows to IT equipment. Because of this dual role, line-interactive UPS batteries can drain rapidly, putting the system at an increased risk for failure. While these systems provide adequate protection for some data center applications, they are not recommended for business-critical applications or facilities that experience a high occurrence of utility failures.

Double conversion UPS systems, on the other hand, condition power through a "double-conversion" process, in which AC power from the PDU is converted into DC power, which is converted back to AC power when it is delivered to the rack or row. This enables the battery to be dedicated to the load and eliminates the need for power transfer if the primary utility fails. While double-conversion UPS systems typically require a larger initial capital investment, they have been proven to be more than twice as reliable as their line-interactive counterparts – making them ideal for truly mission-critical environments.

Data center professionals also should consider the overall durability of their UPS system. Some UPSs, such as the Liebert NXL, are designed with integrated fault tolerances

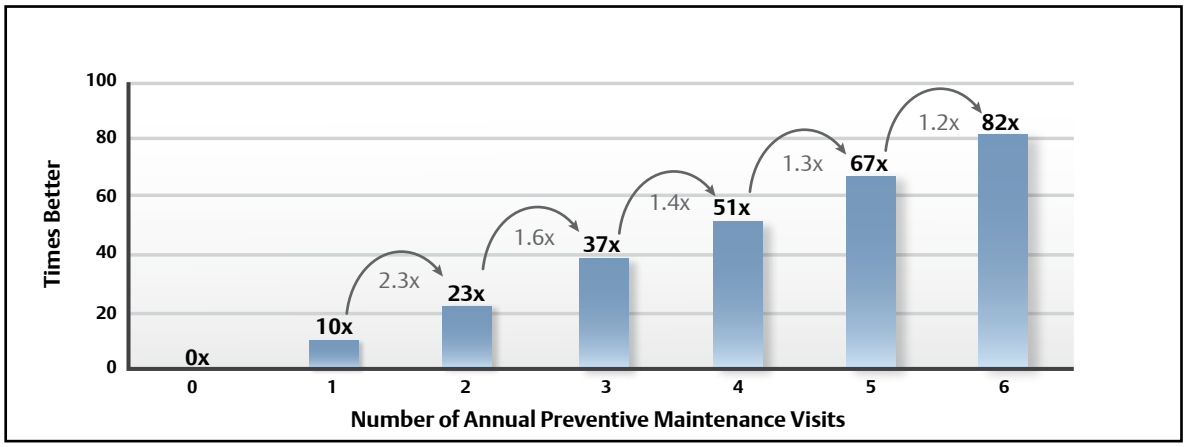

Figure 5. An increase in the number of annual preventive maintenance visits can be directly correlated with increases in MTBF.

for input current and a variety of redundant internal components, including fans, power supplies and communications cards. These features enhance reliability and enable the UPS to maintain availability between service visits even in the event of an internal component failure.

However, regardless of the type of UPS used, preventive maintenance is necessary to maximize the life of critical equipment and minimize the occurrence of unplanned outages. All electronics contain limited-life components that need to be inspected frequently, serviced and replaced periodically to prevent catastrophic system failures. If not serviced properly, the risk of unplanned UPS failure increases dramatically.

While most UPS systems can achieve a lifecycle of 10 years or more under normal operating conditions, it is not uncommon for a well-maintained system to remain in operation for 20 years or more. In a recent study of 5,000 three-phase UPS units with more than 185 million combined operating hours, the frequency of preventive maintenance visits correlated with an increase in mean time between failures (MTBF). As shown in Figure 5, our visits annually increased MTBF more than 50-fold compared to no visits (or more than 25-fold compared to semi-annual maintenance).

It is important to note, however, that while preventive maintenance increases the reliability of a data center's UPS systems, redundancy still is required so that a system can be taken offline and serviced. Without adequate UPS system redundancy, unplanned downtime risk can increase dramatically during a routine service visit.

PDU and Circuit Breaker Failures

As evidenced by the Ponemon Institute's survey findings, UPS-related issues are among the most common root causes of downtime for a data center's power infrastructure. However, it also is important to be mindful of downstream factors that can impact the availability of an entire data center – namely, circuit breaker and power distribution unit (PDU) failures.

A third of surveyed data center professionals identified PDU and/or circuit breaker failures as a root cause of downtime. While these failures do not occur as frequently as UPS-related outages, they nonetheless should be a significant area of concern for data center professionals because a single failure can bring down an entire facility. However, through effective capacity monitoring and management the likelihood of circuit breaker and PDU overload can be reduced significantly.

To maximize visibility into the stress placed on the data center's power infrastructure, data center professionals should consider investing in a PDU that has integrated branch circuit monitoring capabilities. Branch circuit monitoring solutions – such as the Liebert Distribution Monitor (LDM) – utilize branch circuit sensor modules and individual current transformers (CT) to monitor current input/output for the main panel board as well as individual branch circuit breakers, reporting the current and alarm conditions for each breaker.

Many branch circuit monitoring systems also are designed to work in-concert with the data center's building management systems. By pairing circuit monitoring data

with performance information for critical facility equipment, data center professionals gain a comprehensive view of the data center's power infrastructure from the utility to the rack. Many data center management solutions – such as the Aperture Integrated Resource Manager (AIRM) – enable loads to be shifted at the rack-level in real-time. This enhanced functionality allows data center professionals to make precise capacity management decisions based on holistic data across interdependent systems, reducing the likelihood of equipment overload failure downstream.

Beyond enhancing capacity monitoring and management, the addition of static transfer switches to the power system configuration can safeguard data center professionals against downtime.

Most servers in today's market are designed with dual power supplies, capable of either powering the server independently or sharing the load equally. Because of this internal redundancy, these types of servers are less prone to downtime due to faults passed along from the PDU. However, it is easy to overlook that legacy servers in today's data centers are designed with a single power supply and an integrated transfer switch between cords. Because these systems depend on the capacity of a single integrated power supply, there is an increased risk of system shutdown if primary power is disrupted.

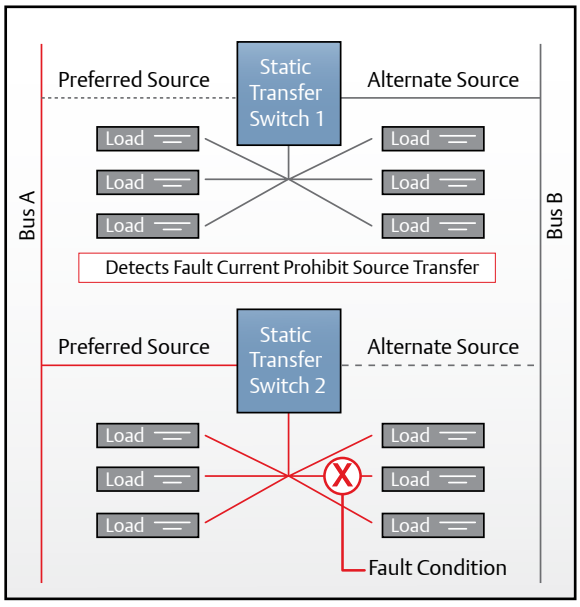

Figure 6. If an STS is being used in a system with a load fault, the fault will remain isolated and will not be passed along to the second bus.

The installation of static transfer switches (STS) upstream from the IT load assures that single-cord loads will be powered in the event of bus failure, maintaining the availability of critical IT equipment. Static transfer switches also provide increased protection during service times, which ensures constant power is delivered to the servers.

In addition to maintaining system availability, STS also enable faults to be compartmentalized and localized to a single bus, preventing widespread outages and increasing the facility manager's ability to quickly identify and remedy root causes. As shown in Figure 6, if an STS is being used in a system with a load fault, the fault will remain isolated and will not be passed along to the second bus. While all single-corded loads connected to the bus will fail in this scenario, distributing the loads among multiple STS minimizes the overall effect of the fault.

Cooling Issues: Heat-Related CRAC Failures and Water Incursion

While many of the leading root causes of downtime are directly related to the data center's power infrastructure, cooling system failures can be equally detrimental to availability. Cooling-related failures were cited as a root cause of at least one outage by more than a third of data center operators polled. Water incursions and heat-related computer room air conditioner (CRAC) failures were cited as the leading causes of cooling-related downtime (35 and 33 percent, respectively).

As high-density configurations become more common, so too does increased heat density.

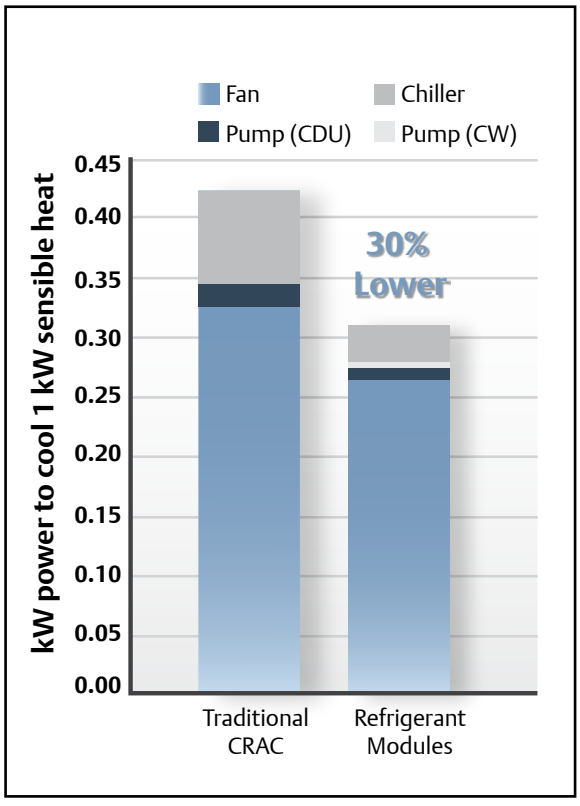

Figure 7. Row-based cooling can lead to energy savings of up to 30 percent in high-density data centers.

As a result, availability of critical systems has suffered. One commonly reported issue is heat-related CRAC failure. Many legacy CRAC systems were not designed to accommodate the high heat densities common in today's high-density data centers. As a result, heat-related failure has become a common concern in data centers seeking to get more from their existing infrastructure.

One way to minimize the risk of heat-related CRAC failure is to optimize air flow within the data center by adopting a cold-aisle containment strategy. In a cold-aisle configuration, rack-rows are arranged facing each other, creating distinct cold-aisles (where cool air from the CRAC enters the rack) and hot-aisles (where hot air is exhausted from the rack). Because the cold-aisles are isolated and sealed off, hot air expelled from the rack is not able to re-enter the cooling environments. This increases the effectiveness of the CRAC system and ensures that cooling capacity is utilized as efficiently as possible.

Many data center professionals choose to increase cooling system effectiveness by utilizing a row-based cooling solution. In fact, the use of row-based cooling solutions can reduce the annual cooling related power consumption by nearly 30 percent (Figure 7).

However, while using row-based cooling modules is a sound strategy for increasing cooling effectiveness and efficiency, some row-based cooling solutions create vulnerabilities that can undermine data center availability – the most significant being water incursion.

Most row-based cooling solutions in high-density environments use fluid to quickly remove high-heat from the rack. These solutions typically fall into one of two

categories: water-based and refrigerant based. While chilled water-based solutions are extremely cost effective, they bear an intrinsic risk for leakage. If not detected early, water incursion within the rack can lead to the corrosion and failure of sensitive electronic equipment, resulting in downtime as well as costly equipment repairs and/or replacements.

With this in mind, the use of a refrigerant-based row-based cooling solution is a best practice for minimizing the risk of cooling-related equipment failures. Unlike water-based systems, refrigerant-based cooling does not rely on an electrically conductive cooling element, minimizing the risk of catastrophic system failures in the event of a cooling fluid leak. At room temperature most refrigerants become a vapor, further reducing the risk of damaging fluid leaks.

If refrigerant-based cooling solutions are cost prohibitive, integrating a comprehensive leak detection system – such as the Liebert Liqui-tect – into the cooling infrastructure is essential to mitigating the risk of system failures due to water incursion. By employing leak detection modules installed at critical points throughout the data center, leak detection systems signal an alarm when moisture reaches potentially hazardous levels. This enables data center professionals to pinpoint leakages within seconds and take corrective actions before sensitive electrical equipment is damaged or destroyed.

Human Error and Accidental EPO

Often the most cost-effective root causes to address, human error and accidental "emergency power-off" (EPO) remain leading causes of data center outages. In fact, more than half of all data center professionals to the Ponemon survey reported at least one outage as a direct result of accidental shutdown or user errors within the past 24 months.

Because these events represent significant yet wholly preventable threats to the availability of mission-critical data centers, data center professionals should observe and enforce the following rules and policies to minimize the potential for catastrophic errors and accidents:

1.Shielding Emergency OFF Buttons
Emergency OFF buttons are generally located near doorways in the data center. Often, these buttons are not covered or labeled and can be mistakenly shut off during an emergency, which shuts down power to the entire data center. This can be eradicated by labeling and covering emergency OFF buttons to prevent someone from accidentally pushing the button.

2.Strictly Enforcing Food/Drinks Policies
Liquids pose the greatest risk for shorting out critical computer components. The best way to communicate your data center's food/drink policy is to post a sign outside the door that states what the policy is and how strictly the policy is enforced.

3.Avoiding Contaminants
Not keeping the indoor air quality clean can cause unwanted dust particles and debris

to enter servers and other IT infrastructure. Much of the problem can be alleviated by having all personnel who access the data center wear antistatic booties or placing a mat outside the data center. This includes packing and unpacking equipment outside the data center. Moving equipment inside the data center increases the chances that fibers from boxes and skids will end up in server racks and other IT equipment.

4.Documented Maintenance Procedures
Following a documented, task-oriented procedure can mitigate or eliminate the risk associated with performing maintenance. This step-by-step procedure should apply to all vendors, and data center professionals also should ensure that back-up plans are available in case of unforeseen events.

5.Accurate Component Labeling
The inaccurate labeling of power protection devices, such as circuit breakers, can have a direct impact on data center load availability. To correctly and safely operate a power system, all switching devices and the facility one-line diagram must be labeled correctly to ensure correct sequence of operation. Procedures also should be in place to regularly double-check device labeling.

6.Consistent Operating of the System
As data center professionals become comfortable with the operations of their data centers, it is common to neglect established procedures, forget or skip steps, or perform "short cuts" and inadvertently shut down the wrong equipment. This further reinforces the point that it is critical to keep all operational procedures up to date and follow the instructions to operate the system.

7.Ongoing Personnel Training
Data center professionals should ensure that all individuals with access to the data center, including IT, emergency and security personnel, have basic knowledge of equipment so that it's not shut down by mistake.

8.Secure Access Policies
Organizations without strict sign-in policies run the risk of security breaches. Having a sign-in policy that requires an escort for visitors, such as vendors, will enable data center managers to know who is entering and exiting the facility at all times.

The Value of Data Center Assessments

As we have highlighted throughout this white paper, a number of technologies and best practices can be adopted to address the leading root-causes of downtime outlined by the Ponemon Institute. However, many data center professionals and senior executives have yet to recognize the value in identifying potential weaknesses in their data center before they result in an unplanned downtime event.

According to the survey, only 13 percent of data center professionals have conducted a data center assessment or audit following an unplanned downtime event (Figure 8). Unfortunately, these results represent a significant missed opportunity for data center professionals to identify and address vulnerabilities unique to their data center.

In addition to power and cooling systems, a number of factors can impact the availability and performance of critical systems. These factors – some as in-depth as arc flash vulnerability, others as subtle as obstructed floor tiles – can go undetected by data center professionals, resulting in decreased performance and an increased risk for downtime. When considering the ever-increasing cost of downtime, conducting a comprehensive data center assessment is the most cost effective way to ensure that all vulnerabilities are identified and remedied before they impact data center operations.

A comprehensive assessment of the facility as well as all thermal and electrical systems also can offer detailed insight into how an existing data center can be optimized for efficiency without compromising the availability of critical systems.

For example, by conducting a computational fluid dynamics (CFD) simulation, data center professionals can gain a thorough understanding of how cooling system configuration affects cooling efficiency and facility performance. Based on the findings of this simulation, existing equipment can be reconfigured for optimal performance, increasing availability and efficiency with little to no capital expenses.

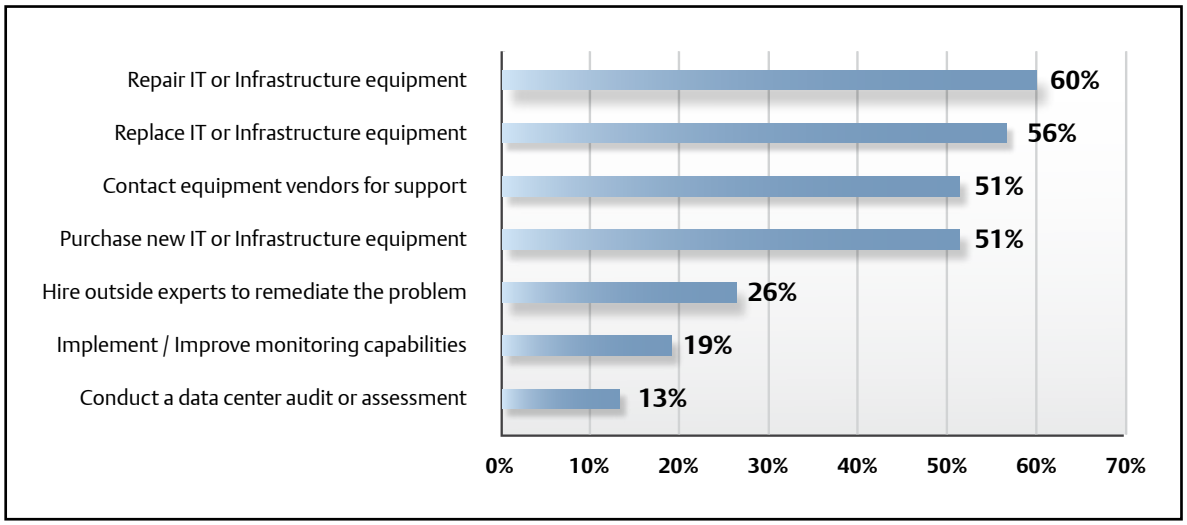

Figure 8. Organizations' response to fixing or correcting root causes as reported by the Ponemon Institute.

Conclusion

As evidenced by the findings of the Ponemon Institute, unplanned data center outages remain a difficult and costly challenge for organizations that have focused their attention on cutting operating costs while maintaining or increasing throughput. Unfortunately, this has led many organizations to overlook actions that must be taken to ensure data center efficiency does not compromise the availability of critical systems.

While root causes reported by survey data center professionals span the entire data center, most can be mitigated cost-effectively by adhering to best practices outlined in this white paper. However, while the root causes addressed represent the most common causes of downtime cited by survey data center professionals, conducting a comprehensive data center assessment is the only way to adequately ensure that all potential vulnerabilities are identified and addressed prior to causing a data center outage.

References

Emerson Network Power e-book, *Avoiding Trap Doors Associated with Purchasing a UPS System*, http://www.EmersonNetworkPower.com, 2010.

Emerson Network Power white paper, *Balancing Scalability and Reliability in the Critical Power System: When does 'N+1' Become 'Too Many +1'?*, http://www.EmersonNetworkPower.com, 2004.

Emerson Network Power white paper, *Choosing System Architecture and Cooling Fluid for High Heat Density Cooling Solutions*, http://www.EmersonNetworkPower.com, 2008.

Emerson Network Power white paper, *Data Center Assessment Helps Keep Critical Equipment Operational*, http://www.EmersonNetworkPower.com, 2007.

Emerson Network Power white paper, *The Effect of Regular, Skilled Preventive Maintenance on Critical Power System Reliability*, http://www.EmersonNetworkPower.com, 2007.

Emerson Network Power white paper, *High-Availability Power Systems, Part II: Redundancy Options*, http://www.EmersonNetworkPower.com, 2003.

Emerson Network Power white paper, *Implementing Proactive Battery Management Strategies to Protect Your Critical Power System*, http://www.EmersonNetworkPower.com, 2008.

Emerson Network Power white paper, *Longevity of Key Components in Uninterruptible Power Systems*, http://www.EmersonNetworkPower.com, 2008.

Ponemon Institute study, *National Survey on Data Center Outages*, http://www.Ponemon.org, 2010.

Emerson Network Power white paper, *Protecting Critical Systems during Utility Outages: The Role of UPS Topology*, http://www.EmersonNetworkPower.com, 2004.

Emerson Network Power technical note, *Using Static Transfer Switches to Enhance Data Center Availability and Maintainability*, http://www.EmersonNetworkPower.com, 2010.

Emerson Network Power
1050 Dearborn Drive
P.O. Box 29186
Columbus, Ohio 43229
800.877.9222 (U.S. & Canada Only)
614.888.0246 (Outside U.S.)
Fax: 614.841.6022
EmersonNetworkPower.com
Liebert.com

SL-24656-R10-10 Printed in USA