
Who tripped the circuit?

The clear-cut case for branch circuit monitoring in data centers

Abstract

The power demands of data centers are exponentially larger and more variable than ever before, creating a real management challenge. How can you utilize the power distribution system to fullest potential, without compromising reliability?

Organizations have invested millions in uninterruptible power systems (UPSs), battery backup, generators and system-wide monitoring to protect their electrical systems. Yet there might still be a proverbial weakest link at the local level: the branch circuit that serves critical equipment. Even the largest and most sophisticated enterprises are suffering unplanned downtime due to a most rudimentary failure: tripped circuit breakers.

The risk is increasing, thanks to blade servers, high-density architectures, variable power consumption patterns and other factors. The good news is that there is a simple, low-cost and surprisingly straightforward solution. New or existing electrical infrastructures can easily be equipped for 7x24 monitoring at the branch circuit level, to ensure reliable, continuous power for essential business systems.

Optimize power utilization and availability on branch circuits.

Improve the management and reliability of the power distribution system.

February 8, 2007

Contents

ABSTRACT1

WHY ARE BRANCH CIRCUITS GETTING OVERLOADED?3

DATA CENTERS ARE MORE POWER-HUNGRY THAN EVER.3

IT EQUIPMENT HAS MORE VARIABLE POWER CONSUMPTION THAN BEFORE.4

CIRCUITS GET LOADED FOR DAY-TO-DAY AVERAGES, RATHER THAN POTENTIAL MAXIMUMS.5

MANY IT COMPONENTS ARE SERVED BY TWO POWER SUPPLIES.5

MOVES, ADDS AND CHANGES ARE CONSTANT.5

IS BRANCH-LEVEL POWER DISTRIBUTION A WEAK LINK?6

WHY TRADITIONAL APPROACHES HAVE FALLEN SHORT6

EATON BRANCH CIRCUIT MONITORING SYSTEM7

DETAILED INSIGHTS INTO CIRCUIT-LEVEL POWER CONDITIONS.7

EARLY WARNING OF PROBLEM CONDITIONS.8

LOCAL AND REMOTE ACCESS TO POWER MONITORING DATA.8

MONITORING OF ENVIRONMENTAL CONDITIONS.8

TRENDING AND LOAD PROFILING FOR MORE ACCURATE MANAGEMENT AND PLANNING8

EASY INSTALLATION IN NEW OR RETROFIT APPLICATIONS.9

SUMMARY.....9

February 8, 2007

Who tripped the circuit?

The clear-cut case for branch circuit monitoring in data centers

The large Internet service provider had spent millions of dollars on reliability measures for its state-of-the-art data center. After all, its data center was its life blood. No downtime was acceptable.

So the infrastructure was bolstered with redundant UPS solutions, redundant cooling, redundant generators and redundant power feeds into the AC servers.

However, in spite of all these protections, the ISP experienced a lot of unplanned downtime—at the worst times, when Internet activity was at its highest. The trigger was embarrassingly simple: branch circuit breakers were tripping.

For this ISP, so much money and focus had been directed at ways to ensure high uptime, but a key aspect had been overlooked: ensuring appropriate loading at the branch circuit level. Variable power consumption patterns caused circuits that were normally loaded at less than 50 percent to jump beyond 100 percent—and trip a breaker. To compound the problem, maintenance staff had a hard time locating the tripped breaker, so there would be unnecessary and costly delays before power was restored.

This is just one example, but it's not unusual. In fact, scenarios like this are surprisingly common. Ironically, as data centers become more sophisticated and technology-rich, the risks actually tend to worsen.

Why are branch circuits being overloaded?

There are several reasons...

Data centers are more power-hungry than ever.

Traditionally, data center managers could plan for about 60-100 watts of power consumption per U of rack space. With today's blade servers, this figure has escalated to more than 600-1000 watts per U and growing. Power consumption can easily double or triple during peak periods, and it fluctuates with every move, add or change. Adding a 1U or 2U server used to mean drawing 300 to 500 more watts from the branch circuit; now a new blade server consumes *10 times* as much current.




Figure 1.
Blade servers have dramatically increased power consumption in a rack.

February 8, 2007

IT equipment has more variable power consumption than before.

Traditionally, computer equipment required pretty much the same amount of power, no matter how much processing was taking place. Trouble was, computers consumed too much power overall, relative to the work they were doing. This was particularly true for laptops, because precious power was being drained from batteries even when the unit was fairly idle.

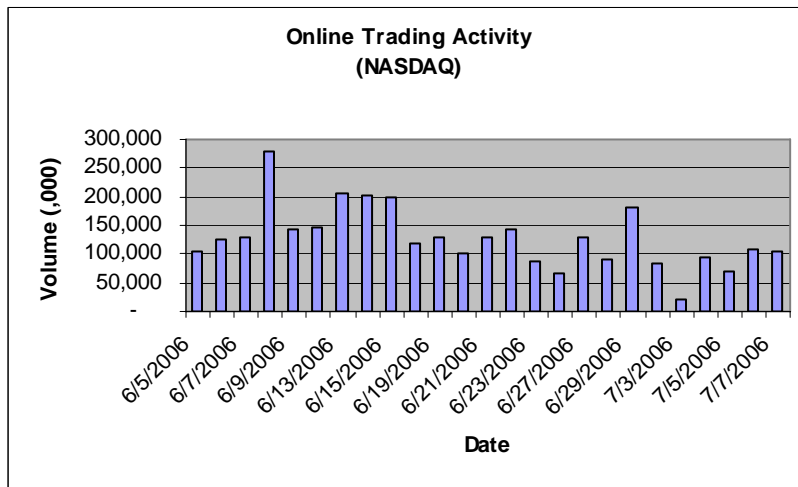
So in the last decade, the trend has been toward designs that consume power at a rate that matches the computing load—less when the PC isn't working very hard, more when needed. Power management technologies such as voltage and frequency scaling, low-power disks, and low-power power supplies have managed to reduce power utilization by more than 80 percent during periods of low computational burden. Not surprisingly, server manufacturers started adopting these technologies to reduce power consumption and heat output in high-density data centers and network rooms.

So why is that a problem for branch circuits? With the new and more efficient technologies, server power consumption now varies dramatically as the workload varies—often more than doubling. Imagine the impact of seasonal peaks in retail sales for an online retailer, such as eBay, Amazon or Wal-Mart—or card vendors such as Visa and MasterCard. Christmas season volumes could stress circuits beyond their thresholds, bringing systems down at the worst possible time.

As new server-class processors are developed, consuming more power than ever to deliver higher performance, the problem is likely to worsen. Six years ago, a typical server consumed 100 watts per U space. Today's blade servers consume more than ten times that amount, more than 1 kilowatt. A 10 percent variation in power consumption today means more than a 50 percent variation in old servers.

Blade servers also use very few disk drives, which means processor power represents a larger percentage of overall power consumption—and variations due to changing computational loads become much bigger factors than ever.

Figure 3. Transaction volume can vary 200 to 500 percent in some enterprises, with big implications for data center power consumption.



February 8, 2007

Circuits get loaded for day-to-day averages, rather than potential maximums.

Server consolidation efforts have increased server utilization, but still the average server operates at 5 to 30 percent of processing capacity, typically less than 15 percent. That means most servers spend much of their time operating at light computational loads, drawing far less power than they potentially could.

Under normal operating conditions, data center managers could mistakenly assume there's plenty of capacity on a branch breaker or panel board. After all, day-to-day power consumption seems well within limits. But there's always the potential for one or more servers to increase computational load, draw more power to match the workload, and cause a branch breaker to trip. With the increased use of denser, more space-saving equipment, the risk of overloading a branch breaker is higher than ever.

Since average server utilization tends to be light, the risky situation can go undetected for a long time. The breaker then trips and shuts down power to critical systems at the most damaging time, when the equipment is handling the highest transaction volumes.

Many IT components are served by two power supplies.

The equipment has two independent yet paralleled power supplies and accepts two separate power feeds. Either power supply is sufficient to take over if the other fails.

On the surface of it, that doesn't seem like a problem, does it? When power requirements are distributed among two sources, you gain redundancy and fault tolerance for greater system availability. If a power supply or power source fails, for whatever reason, the other can take over and the equipment keeps running. However, consider that scenario from the branch circuit perspective. The power supply inside the IT equipment may be sized to handle the full power draw, but does the circuit have the capacity? Often, it does not.

Each breaker in a dual power path must be loaded at less than 50 percent of trip rating during normal conditions. That way, a failure of one path won't overload the alternate path and cause breakers to trip. However, if you've loaded each circuit conservatively at less than 50 percent, a failure condition can go unnoticed for a long time, because the alternate power path has picked up the slack, and servers are running normally.

Then, what if the unthinkable happens? What if the alternate path, now the sole power path, fails? It trips the breaker and brings the system down, possibly even producing a cascading effect that jeopardizes upstream power feeds and puts the entire system at risk.

Moves, additions and changes are constant.

If your data center is typical, flux is continuous. Managers of large data centers reported that moves, additions and changes take place several times a week. Even in the smallest computer rooms, changes take place at least once a month (see chart below).

Study - July 2006 (474 IT/Data Center Managers)			
Frequency of M.A.C. in a Data Center	Small Computer Rooms <20 Racks	Medium Data Centers 20- 100 Racks	Large Data Centers 100+ Racks
Several times a week	11.3%	29.5%	42.0%
Once a week	15.7%	26.4%	18.0%
Once a month	40.7%	28.7%	38.0%
Once a year	28.8%	10.1%	20.0%
Never	3.6%	5.4%	-

February 8, 2007

Changes in data center configuration can put a lot of pressure on branch breakers. In constantly changing environments, it's easy to create misbalanced loads. Because you don't have current or historic data on the current draw per breaker, you might overload one breaker while under-loading another.

The risks are intensified by the more power-hungry nature of new servers. Chances are the breaker loads are optimized based on the average power consumption of servers when the data center was designed. But new servers consume a far greater percentage of the power available on a breaker. That means there's a much higher likelihood that adding new equipment to a branch circuit will result in overload.

Is branch-level power distribution a weak link?

It often is. When power consumption is greater than ever, and varies minute by minute, branch circuits are at risk. It is easy to overload a circuit, trip a breaker, and cause costly, unplanned downtime. The problem has grown significantly in the last few years and continues to grow—reaching a point where it cannot be ignored.

Seem improbable? Even the architecture of the typical panelboard represents a point of vulnerability. Consider that a standard panelboard generally has a main breaker rated at 225 amps with 42 branch circuits. Based on a derating factor of 80 percent for all breakers, 42 single-pole 20-amp breakers can potentially draw up to 672 amps—nearly *four times* as much as the main breaker. Without proper coordination and monitoring, a branch breaker could trip and actually cause the panelboard main breaker to trip as well.

Furthermore, many organizations rely on so-called “lights out” or “dim” data centers that are not staffed. In these environments, a tripped breaker can be particularly problematic. Some form of automated monitoring must be implemented to replace the human presence.

Why traditional approaches have fallen short

Clearly, changing conditions and new IT equipment predispose today's data centers to the costly and damaging effects of tripped circuit breakers. At the branch circuit level, you might not be able to see trouble coming until the breaker trips, and that's too late. Systems go down. Valuable data is lost, and business comes to a standstill. It can take hours to recover.

The best way to avoid such an undesirable situation is by monitoring the power consumption of each branch breaker.

Some organizations address this need by periodically checking current loads with hand-held meters. Unfortunately, this occasional sampling can easily miss problem conditions, which can appear and disappear in microseconds. The sampling might also take place during lightly loaded periods, and therefore produces unreliable data. And of course, it would be cost-prohibitive to have a permanently installed meter on every branch circuit.

Another common approach is to monitor power at the rack power strip level. However, a single branch circuit may support multiple rack power strips or racks. Conditions could be fine at the power strip level but approaching overload at the branch circuit level.

Furthermore, not all equipment in a typical data center is powered from power strips. Storage devices, mainframes, and other large elements are usually powered directly from a branch circuit. Strip-level monitoring would miss these systems altogether. And finally, monitoring at the power strip level requires the organization to reserve a lot of IP addresses and connections, which complicates the management environment.

February 8, 2007

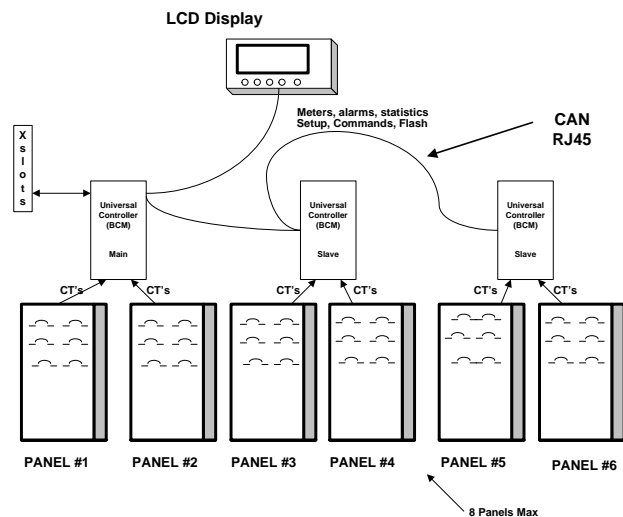
The only dependable way to monitor appropriate loading for each breaker is by installing monitoring devices that measure the current draw 7x24. However, branch circuit monitoring systems that have been on the market until now were seen as too expensive and did not provide enough data points or features to ensure complete information. Facility managers needed a better way to ensure reliable, continuous power at the branch circuit level.

Powerware® Energy Management System

The Powerware Energy Management System continuously measures the current on branch circuits and warns you of impending trouble, so you can take proactive action. Armed with these insights, data center and facilities managers can more effectively manage energy consumption to prevent overload conditions, optimize power distribution, and, when applicable, accurately bill internal customers for power usage.

The architecture is straightforward, which makes it also stable and reliable. Branch circuit wires are threaded through current transducers (CTs), which sense the current level going through the wire and send the information to a universal controller. The controller assimilates the information from CTs at multiple panelboards and then relays it to a local display or remote management system for analysis and reporting.

A single system can monitor up to eight 42-circuit panelboards (up to 336 pole positions and/or 16 three-pole sub-feeds), in addition to monitoring the main input/output parameters of the distribution equipment, such as a power distribution unit (PDU) or remote power panel (RPP).



The Powerware Energy Management System sends data from up to eight panelboards, up to 336 pole positions, to a single display—and sends warning alarms if conditions could potentially trip any branch or main circuit breaker.

Detailed insights into circuit-level power conditions.

The system assesses circuit activity 24/7 and provides time-stamped metering, alarm and statistical information for each branch circuit. You get the significant information necessary to effectively manage the edge of the power distribution system.

For example, you can see minimum, maximum and average current, voltage and frequency; power quality metrics such as total harmonic distortion and power factor; time-stamped event records and more. This information is shown for individual circuits and summarized for each panelboard and at the equipment level, such as PDU or RPP level. You get visibility at all levels within one unit.

February 8, 2007

With this information, you always know how close a circuit is to exceeding its overall rating—and whether or not a device can be added to a branch circuit or panelboard. This insight enables you to operate at maximum efficiency, using data center assets, energy and space wisely. Without this insight, you wouldn't know if you're pushing the limits of a breaker or any safety factor.

Early warning of problem conditions.

When current on a branch circuit or main circuit breaker reaches a user-defined threshold, the system sends a warning alarm to the local display and the defined communications network, identifying the specific circuit and the condition. The data center or facility manager is alerted, with time-stamped information, *before* the current approaches the breaker's trip point. An on-board, real-time clock accurately time-stamps the min/max data with date and time of occurrence.

Local and remote access to power monitoring data.

Other branch circuit monitoring systems on the market require that you purchase a separate display if you want to view circuit activity and alarms locally. With the Eaton system, a multi-line LCD local display is included in the base system package. Technicians always have an up-to-the-minute view of power conditions right at the equipment.

With an optional X-Slot™ communication card—the same type of card used in our UPS solutions—the universal controller sends data to your building management system or power management system (such as those incorporating PowerVision® or FORESEER® software) using industry-standard Modbus and Web/SNMP protocols.

All the available circuit-level and system-level information can be accessed any place in the world via most common communication interfaces. Similarly, notifications of warnings and alarms are displayed locally and sent remotely. This is especially helpful while managing large data centers with hundreds of panels and thousands of branch breakers.

Monitoring of environmental conditions.

Generally, most of the electrical power consumed by computing equipment is dissipated as heat. When the power consumption varies (as it does), so does heat output, potentially creating undesirable hot spots.

To mitigate that risk, the Energy Management System also monitors key environmental conditions, such as temperature and humidity of the distribution equipment. For a PDU, the system can measure the temperature of the unit's transformer. These capabilities enable data center managers and facility managers to identify trouble conditions and then design and monitor layouts that avoid adverse hot spots.

Trending and load profiling for more accurate management and planning

It is important to be able to assess the historical power consumption of a branch circuit or panelboard, to consider past trends before adding new loads.

A unique feature of the Powerware Energy Management System is its ability to archive data and perform load profiling across many months. The system captures the highest and lowest readings of current, frequency, power level and total harmonic distortion for the main panelboard breaker and branch breakers.

Users can view data for the previous 23 months to gain valuable insight for managing power consumption and diagnosing issues. With this innovative feature, you can assess long-term load trends and optimize utilization of available power at all levels—which in turn will help extend the useful life of existing infrastructure and delay the need for new investment.

February 8, 2007

Easy installation in new or retrofit applications.

The Energy Management System is designed to complement and effectively work within a broad array of data center and facilities applications around the globe. It can be used in a variety of products, such as Powerware PDUs, remote power panels and rack-based distribution products, as well as third-party power distribution products. Packaged kits are available for use with any manufacturer's panelboards and breakers.

Retrofitting is a simple process, given that the current sensors do not require any calibration in the field, even if the configuration of branch circuits is not known.

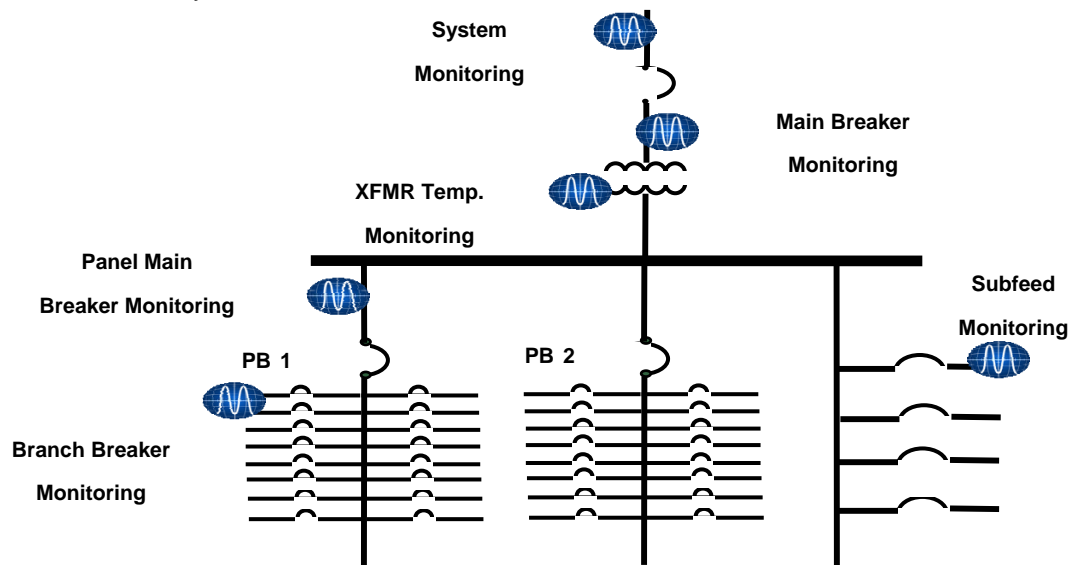
Summary

Branch circuit monitoring technology has been field-proven for years, and recent advances in technology and design have made the solution more affordable than ever. It definitely provides cost-effective insurance against tripped circuits and unplanned shutdowns.

The cost of downtime is difficult to quantify, but substantial. Besides the obvious out-of-pocket losses, there are hidden costs, such as loss of customers, reputation and image. The average cost of downtime ranges from approximately \$90,000 per hour in the transportation industry (airline reservations) to \$6.5 million for large financial brokerages (according to Contingency Planning Research, [2001]). Without a branch circuit monitoring system, it can take more than an hour to discover and rectify the problems created by a single trip.

In contrast, the average price for solid prevention ranges from \$1,000 to \$2500 for a 42 circuit panelboard.

When comparing features and technical specifications of different systems, you will quickly see that our system provides more all-in-one functionality than other vendors' offerings. It takes more of a system approach and is capable of monitoring all necessary parameters of your power distribution system, including branch level, sub-feed level, panel main breaker and system input level. It has the broadest list of measured and reported critical parameters in the market. Features such as local display, main circuit breaker monitoring and coverage for 336 circuits in a single display — which are add-on options or unavailable with other systems — are standard on ours.



February 8, 2007

Find out more about how Powerware Energy Management Systems can add an essential, extra layer of reliability to the management of your electrical distribution system. Contact Eaton at 800.356.5794 or visit us on the Web at www.Powerware.com].

February 8, 2007
