High-Density Computing:

The Path Forward 2006

A Special Report from the first High-Density Computing Symposium, April 2006

Double Your Power at the Rack...

...by Eliminating Inrush for Out-of-Phase Transfers

By Milind Bhanoo, President, LayerZero Power Systems

combined with...

"The Path Forward 2006" White Paper Executive Summary

By Ken Brill, Founder and Executive Director, The Uptime Institute

LAYERZERO POWER SYSTEMS, INC.

THE UPTIME INSTITUTE®

White Paper
(edited and approved by The Uptime Institute)

Double Your Power at the Rack

…by Eliminating Inrush for Out-of-Phase Transfers

by Milind Bhanoo

(Editor's Note: This paper has been reviewed for technical merit, accuracy and usefulness by the staff of The Uptime Institute. It was prepared for inclusion in the first edition of the Quarterly e-Journal on High-Density Computing: The Path Forward *2006 the official publication of The Institute's High-Density Computing e-Community of Practice and Symposium.)*

The focus of recent data center conferences has been on trying to remove additional heat being produced in computer rooms. The introduction of blade servers has not only increased the amount of heat being produced, but also increased the demand for AC power at the rack. No matter which cooling solution is chosen, the reality of the situation is that more power must be delivered to the racks. Here we address a power distribution architecture that can double the critical power available without increasing floor space.

This paper also presents an opinion related to the operation of *automatic static transfer switches (ASTS)* in data-center power architectures where the ASTS input voltages may be out-of-phase with each other during an emergency transfer. An out-of-phase transfer will cause an instantaneous phase shift to any and all transformers that may be downstream of the ASTS. This phase shift can result in saturation current from the transformer that will trip upstream circuit breakers and drop the critical load. For this reason, we must understand transformer physics. The objective of this paper is to examine some general facts related to a combination of circuits and systems, and describe an informed opinion as to the anticipated consequences of their use.

The Magnetic Effects in Mission Critical Systems

In the US, bulk electrical power is distributed within most data centers using redundant *480Volt (V), 3Phase (P) 3Wire (W)* circuits. These circuits may or may not be synchronized in phase, since their sources are usually combinations of *uninterruptible power supply (UPS),* engine-driven generators and utility feeds.

At the point of use, these redundant *480V, 3P3W* circuits must be converted to *208V, 3P4W* circuits for powering typical computer, communication and data storage equipment loads. Two specific functions are required at the load point: *(1)* selection of one active source, and *(2)* transformation of voltage.

The source-selection is generally accomplished with an ASTS, capable of extremely fast source quality monitoring and subsequent switching of the load from one source to another. The transfer is always accomplished in a small fraction of a cycle with modern ASTSs.

The conversion of voltage from distribution to point-of-use voltage is almost universally accomplished using a simple transformer (see *figure one).* The ASTS high-speed transfer-provided results in a voltage waveform at the output that has the potential to upset connected loads when the phases of the sources feeding the ASTS are widely separated. Transformers operate by converting electrical to magnetic energy using a process described by Faraday's law of induction and Ampere's circuital law.

The same transformer then converts magnetic energy back into electrical energy by reversing the primary conversion relationship. The constant of proportionality that relates the *time* integral of applied voltage (or magnetic flux) with *current* (or induction) is the permeability of the medium in which the magnetic circuit operates. In order to construct transformers that operate at power- circuit voltages and power-circuit frequencies with reasonable efficiency and size, a material with a very high permeability must be used for the magnetic circuit. Iron or materials primarily composed of iron are normally used for this purpose. These ferromagnetic materials have permeability values many orders of magnitude greater than free space.

While ferromagnetic materials have the desirable properly of very high magnetic permeability, they provide this benefit with the serious disadvantage of being subject to magnetic saturation. In ferromagnetic alloys, the permeability is not constant, but is a non-linear function of flux and induction.This non-linear relationship leads to the familiar magnetic B-H curve,

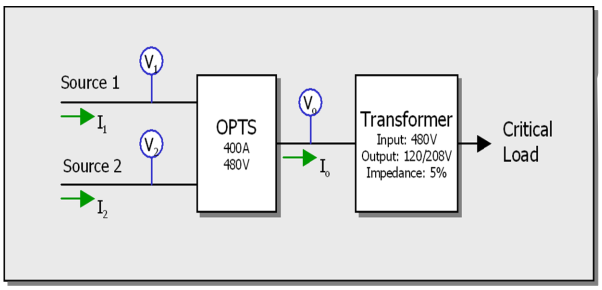

figure one — Type-P ePODs System. Primary-side STS feeds power

(see *figure two)*, showing that there's a finite limit to the number of volt-seconds of magnetic flux that the transformer can support.

Beyond the saturation point, at which maximum flux is reached, the permeability of the magnetic material drops rapidly. Since the permeability relates the integral of voltage to current, a reduction in permeability leads to a reduction in impedance as seen by the voltage source connected to the transformer. If the transformer is powered by an essentially constant voltage source, the result is a drastic increase in current flow through the transformer. This drastic increase in current is limited only by the source and conductor impedance, the transformer leakage impedance and what little impedance is provided by the saturated transformer material.

When driven by a steady sinusoidal voltage, properly designed transformers are never subject to saturation. This is a consequence there being are exactly the same number of volt-seconds of flux produced by the positive half sine wave as there are produced by the following negative half. The designer will have insured that the peak amount of flux produced in the transformer material will remain below the saturation point, with an acceptable margin to account for rated over-voltage conditions and material tolerances. The operation of the transformer in the region of high permeability is demonstrated by the small excitation current drawn by transformers under normal conditions.

As described above, the output voltage waveform of a solid-state transfer switch in a data-center may present nearly instantaneous voltage phase changes of up to 180 degrees. If the ASTS transfers to a new source that is out of phase with the previous source, the transformer may be subjected to additional volt seconds of flux of the *same polarity* as the previous half-cycle. If the phase separation between old and new sources is sufficiently great, additional volt seconds of flux may exceed normal design margin and drive the transformer into saturation. A worst-case is where the ASTS transfers near zero volts to another source that is 180 degrees out of phase.

No rationally designed transformer has a margin of two times excess flux capacity, so saturation is virtually assured in this case. This condition may result in a large flow of current from the source into the transformer primary (22 times nominal), lasting for enough cycles to clear upstream circuit protective devices. A similar phenomenon can occur when a transformer is initially energized. If the last half-cycle of voltage that powered the transformer at shut-down is of the same polarity as the initial half-cycle at start-up, the flux remaining in the core will add to the new flux being added at startup, often leading to an inrush current manifestation of material saturation.

Solving the Inrush Problem, Phase One

Until late 2003, there were only two ways to mitigate the problem of saturation current during switching. *Method one*, shown on the left in *figure three*, depends on using a low inrush transformer at the ASTS

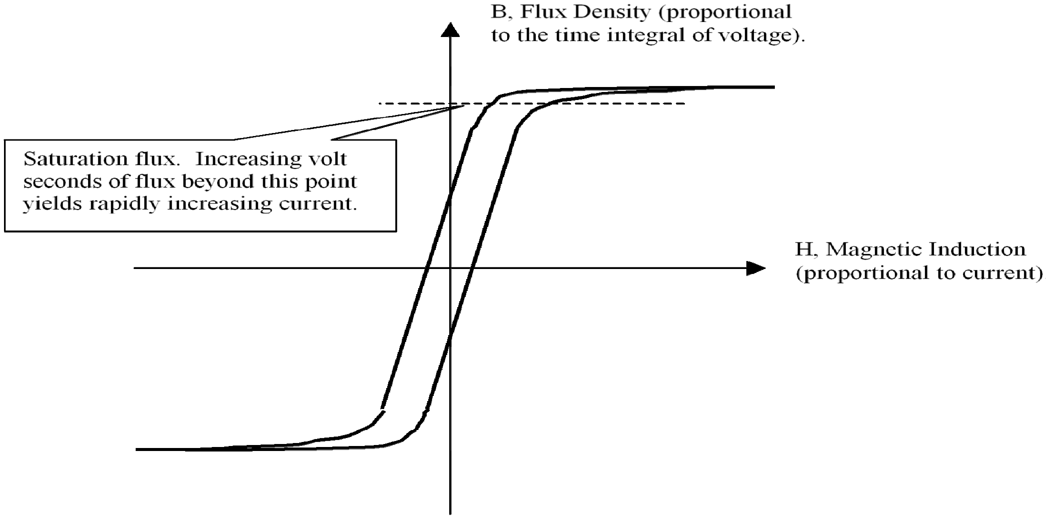

figure two — general form of B-H curve for ferromagnetic materials.

output. By designing a transformer with the appropriate balance between flux capacity and leakage impedance, saturation current can be managed to workable levels and coordinated properly with upstream circuit protective devices.

Method two, depicted on the right in *figure three* is to place the ASTS between the secondary outputs of two transformers. In this case, the saturable magnetic elements are placed in the circuit where they won't be subjected to voltage phase switching events. Although requiring an additional idle*, i.e.,* non-load bearing transformer, this method precludes any current due to distribution transformer saturation.

Method one reduces inrush to five times nominal (down from 22). *Method two* completely eliminates inrush because there is no downstream transformer. In 1995, consulting engineers gravitated to *method two* and the pendulum swung from using static switches on the primary side of transformers to the secondary side. The upside was the elimination of the inrush problem and the downside was *(1)* the ASTS footprint increased, *(2)* the amperage of the ASTS doubled, and *(3)* the price of the system increased significantly.

Switch Mode Power Supplies ***(****SMPS)* circuits are used to couple the basic AC input voltage to the internal circuitry of the mission critical equipment. In 1980, these components were experiencing problems that caused computers to crash at an alarming rate. That was the year the Computer Business Equipment Manufactures Association (CBEMA) was founded. Members included IBM, Amdahl and Wang. Their goal was to find out why their systems were going down.

The CBEMA Curve

After much testing, it was discovered that the problem was voltage deviations. They published their findings in the form of the *CBEMA curve* which represents the *withstand* capabilities of computers in terms of the magnitude and duration of a voltage disturbance *(see figure four).* The under-voltage portion showed that the SMPS internal to the computers could survive without voltage for *eight ms (ms)* or *0.5 cycles*. After *eight ms*, increased voltage is required.

This is why electrical consultants required all ASTS supplying power to computers transfer within *four ms* (well inside of the *8ms* threshold). Today, after 20 years of improvement and testing, SMPS can survive without voltage for *20ms (1.25 cycles).* In 2000, the CBEMA curve was updated by *The Information Technology Industry (ITI)* to reflect the *20ms* change. The new curve can be found at http://www.itic.org.

The ITI/CBEMA curve allowed for a control algorithm to be designed for *480V* out-of-phase transfers.

Named Dynamic Phase Compensated Transfers, it's intended to automatically eliminate inrush from downstream transformers *(Patent Pending, 2003, LayerZero Power Systems, Inc.).*

The answer was a simple time delay added to the transfer procedure. The duration of the delay is determined by the phase angle between the two sources at the instant of transfer. The delay is a consequence of there being exactly the same number of volt-seconds of flux produced by the positive-half sine wave as by the following negative-half sine wave.

If the ASTS transfers to a new source that's out of phase with the previous source, the transformer may be subjected to additional volt-seconds of flux of the *same polarity* as the previous half-cycle. A time delay eliminates this possibility. If the sources are 90 degrees (or less) apart, the transfer will be *4ms* in duration. If the sources are 135 degrees apart, the transfer will take *6ms*. If the sources are 180 degrees apart, the transfer will take *8ms*. A sample waveform capture with a phase angle variance of 180 degrees is shown in *figure five*. It's important to note that there's no inrush *(current excursion)* in the Output Amps waveform; and the output voltage (critical load voltage) is well within the ITI/CEBEMA bounds (maximum allowable outage = *20 ms*).

Phase Angle Difference = 180 degrees

480V, 400A ASTS; 300kVA Transformer. *300A* Resistive Load

Source one = Generator *(1200kW);* Source two = Utility

Return of Primary-Side ASTS — Double the Power

With the advent of this new control algorithm, consultants have been slow to change their power distribution designs and return to primary-side switching. In 2006, several Fortune 100 Companies took the initiative. By doing so, the user can either save money and space, or produce twice the power in the same footprint. Let's investigate both options:

Save money and space -- With a secondary-side ASTS, a typical configuration is two *216kVA* transformers supplying power to a *600A* static switch. This equates to *600A* of current available to the critical loads. By changing to a primary side switch, the static switch becomes a *400A* unit feeding a *288kVA* transformer. This equates to *800A* of current available to the critical load. The owner has reduced the footprint by getting rid of one transformer (saving space) and saving money by reducing the ASTS amperage to *400A* and eliminating a transformer.

Produce twice the power – A typical secondary-side configuration of two 288kVA transformers feeding an *800A* static switch equates to *800A* of current avail-

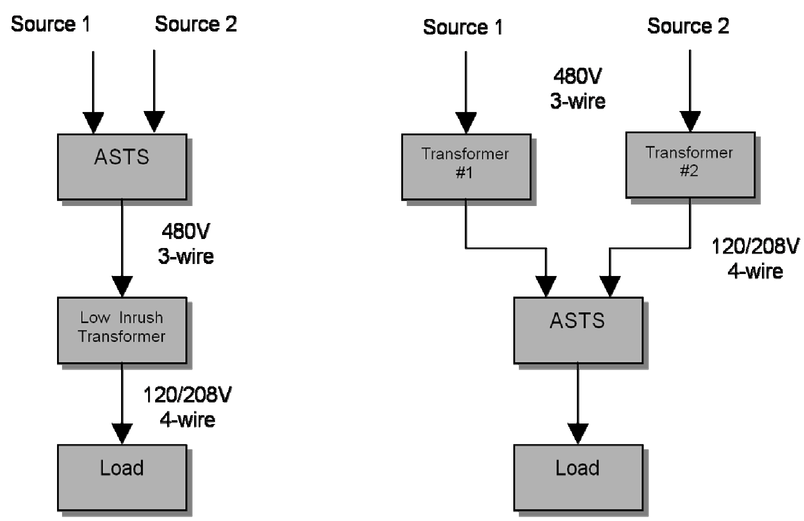

figure three — block diagrams of "method one" (left):load side and "method two": source side transformer placement.

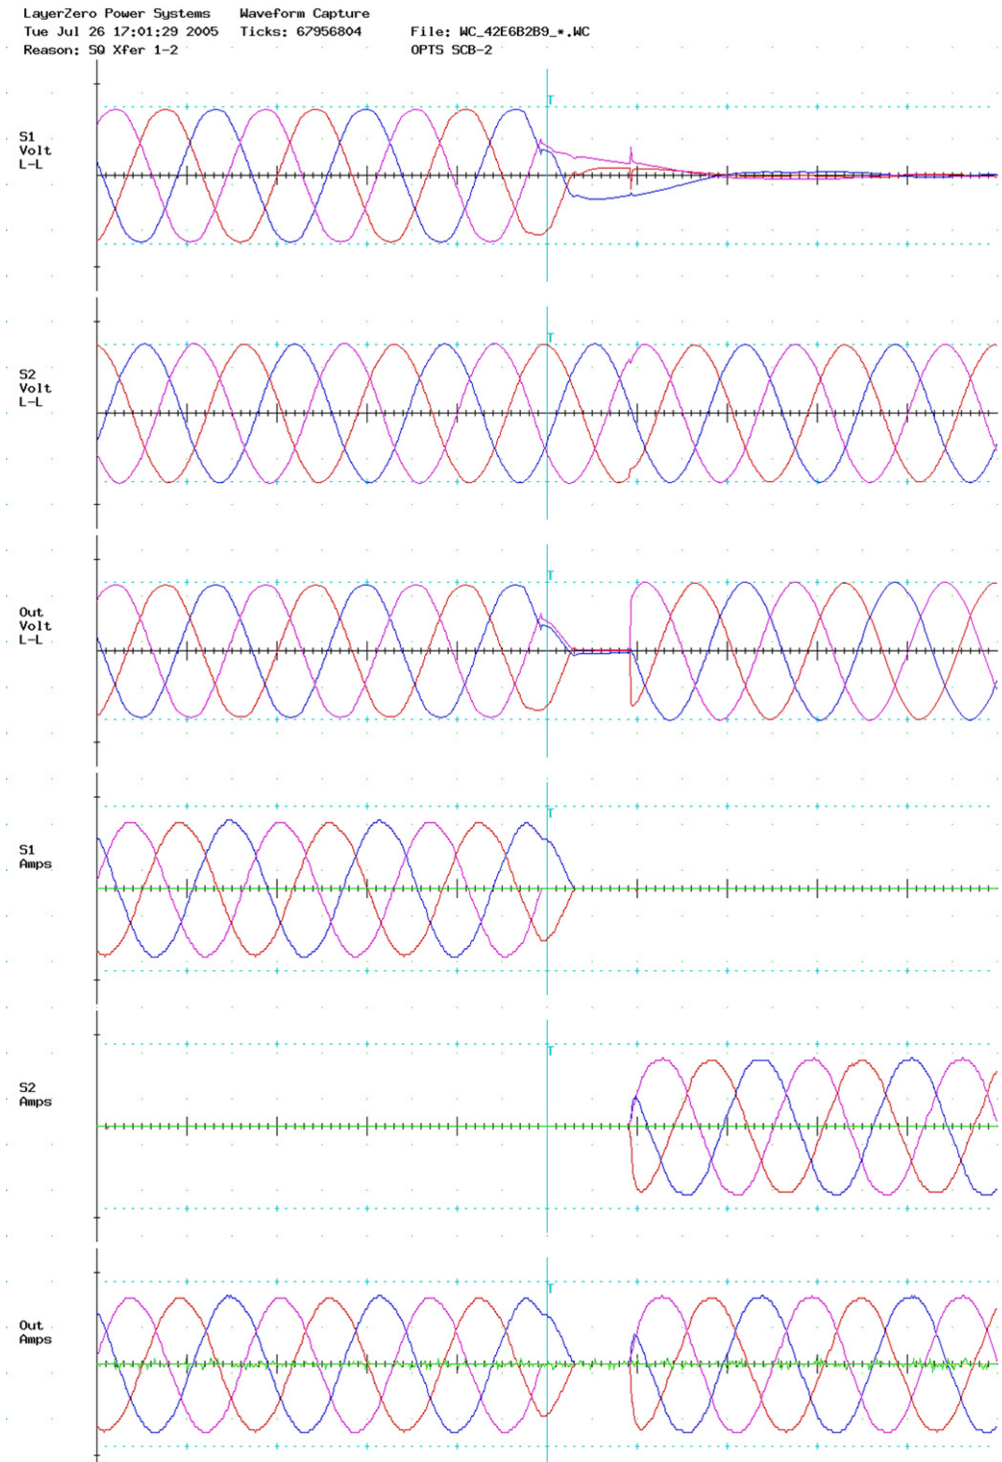

figure five — an out-of-phase transfer where the two sources are 180 degrees apart. The Output Amps show no inrush.

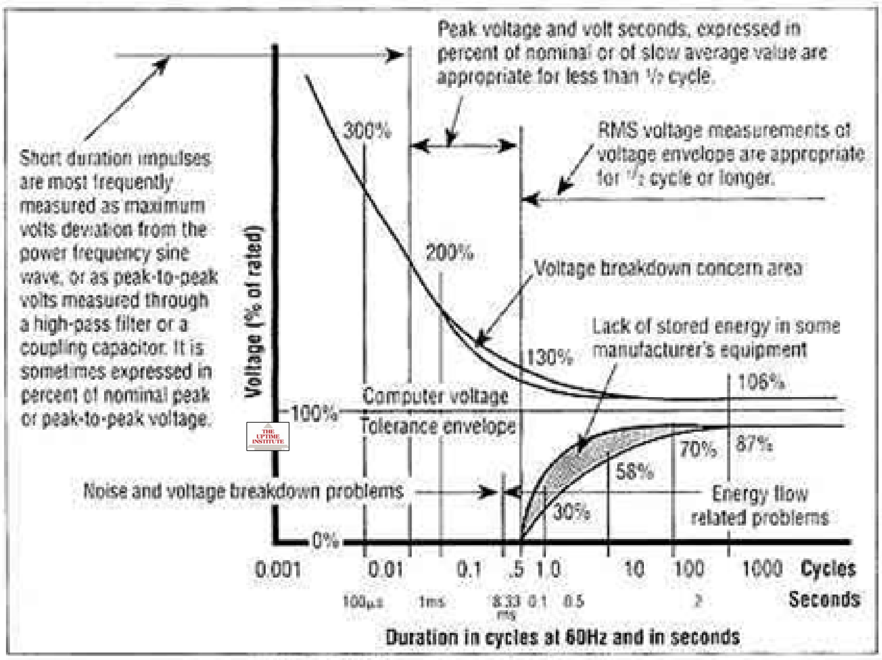

The CBEMA Curve from FIPS No. 94.

figure four—the CBEMA Curve as it first appeared in 1980.

able to the critical load. By changing to a *480V* ASTS and by moving the two transformers to the output, the effective current available to the critical load is *1600A.* That is twice the power! The footprint is the same.

Conclusion

Magnetic circuit components, such as power distribution transformers, play a major role in architectural decisions because of their characteristic large current draw when subjected to switched-phase voltage. It has been proven that servers are not affected by zero voltage for up to *20ms*. The use of a time-delay algorithm enables a *480V* static transfer switch to be used to provide many improvements over the secondary-side alternative. In ten years, the static switch pendulum has swung from primary-side to secondary-side and, now, back to the primary side.

About the Author: Milind M. Bhanoo is President and Co-Founder of LayerZero Power Systems. During his 14-year tenure in the critical power distribution industry, he's held various roles in product, sales, marketing and general management positions. He holds a BSEE, MS in science in systems and control engineering and an MBA, all from Case Western Reserve Univ. **Contact:** mbhanoo@layerzero.com

LayerZero Power Systems, Inc.

The Uptime Institute

2904 Rodeo Park Drive East, Building 100 • Santa Fe, NM 87505
Fax (505) 982-8484 • Phone (505) 986-3900
info@uptimeinstitute.org • www.uptimeinstitute.org