Analysis and Opinion on the Operation of Switch Mode Power Supplies under Switched Voltage Phase Conditions

James M. Galm, Ph.D., P.E., LayerZero Power Systems, Inc.

I. INTRODUCTION

This note presents an opinion related to the operation of switch mode power supplies (SMPS) as loads in data center power architectures where their input voltage may be subject to instantaneous phase shift. The opinion is motivated by the desire to use SMPS loads in data center architectures where the SMPS may be supplied by a power source capable of creating an instantaneous voltage phase shift. Such a data center architecture yields measurable performance and reliability advantages. Unfortunately, the reaction of the SMPS to said voltage input condition is usually not specified by the SMPS manufacturer. The objective of this note is to examine some general facts related to this combination of circuits and systems, and describe an informed opinion as to the anticipated consequences of their use.

Commercial SMPS circuit topologies have gravitated to a level of similarity due to continuous performance and cost optimization processes. The common traits that many commercial SMPS architectures share can be examined to understand the anticipated effect of voltage phase switching. It will be postulated in this note that the most common SMPS circuit topologies can be reasonably expected to operate through voltage phase switching events without deleterious effect on the SMPS or the circuits powering the SMPS.

There are important limitations to what can be asserted about the operation of a large class of potentially complex devices (SMPS in this case) when subjected to conditions for which the devices' designers do not make concrete specifications. It is important to state that this paper does not purport that any particular SMPS design or realization will operate correctly or reliably under voltage phase switching circumstances. To make that assertion without detailed knowledge of the particular SMPS design under consideration would be very irresponsible. Furthermore, this paper does not assert that all SMPS designs will operate in a reasonable manner under voltage phase switching, since SMPS design theory and implementation techniques are widely varied and constantly evolving. Such a blanket assertion regarding a wide and changing technical area such as SMPS design would be similarly irresponsible. With these limitations clearly understood, it is still possible to reach conclusions based on often used SMPS architectures and circuits that can help guide data center architectural decision-making.

II. MAGNETIC EFFECTS IN MISSION CRITICAL SYSTEMS

In the USA, bulk electrical power is distributed within most data centers using redundant 480V, 3P3W circuits. These bulk power distribution circuits may or may not be synchronized in phase, since their sources are usually combinations of 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 switching function is generally accomplished with a solid-state transfer switch, 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 STS equipment. The conversion of voltage from distribution to point of use voltage is almost universally accomplished using a simple transformer.

The high speed transfer provided by the STS results in a voltage waveform at the STS output that has the potential to upset connected loads when the phase of the sources feeding the STS are widely separated. The most common issue associated with loads that do not perform adequately under voltage phase switching conditions is loads that contain magnetic components such as transformers, induction motors and synchronous AC motors.

Transformers operate by converting electrical energy 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 a 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, as shown in Figure 1, showing that there is a finite limit to the number of volt-seconds of magnetic flux that the transformer is able to 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.

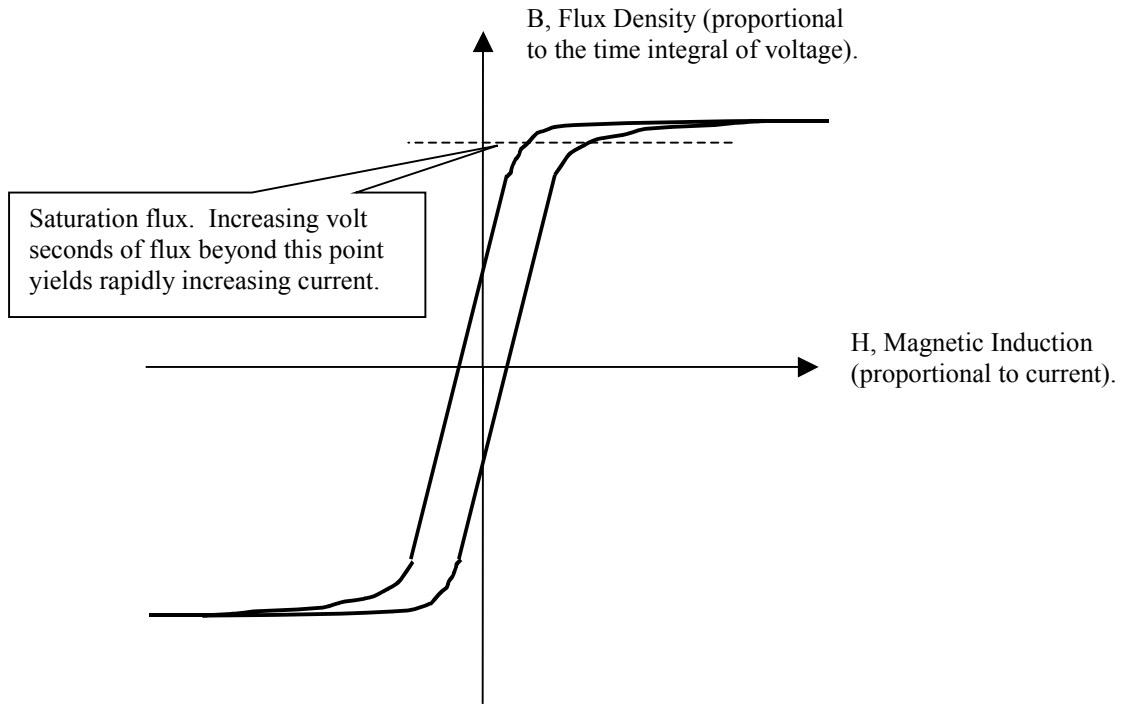

Figure 1. General form of B-H curve for ferromagnetic materials.

When driven by a steady sinusoidal voltage, properly designed transformers are never subject to saturation. This is a consequence of the fact that there 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 sine wave. The transformer 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 environment may present nearly instantaneous voltage phase changes of up to 180 degrees. If the STS 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 the old and new sources is sufficiently great, the additional volt seconds of flux may exceed the normal design margin and drive the transformer into saturation. The worst-case situation is where the STS 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. The saturation condition may result in a large flow of current from the source into the transformer primary, 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 startup, 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.

There are two common practices used to mitigate the problem of saturation current during switching. One practice, shown on the left in Figure 2, depends on using a low inrush transformer at the STS 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. The second practice, depicted on the right in Figure 2 is to place the STS between the secondary outputs of two transformers. In this case, the saturable magnetic elements are placed in the circuit where they will not 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.

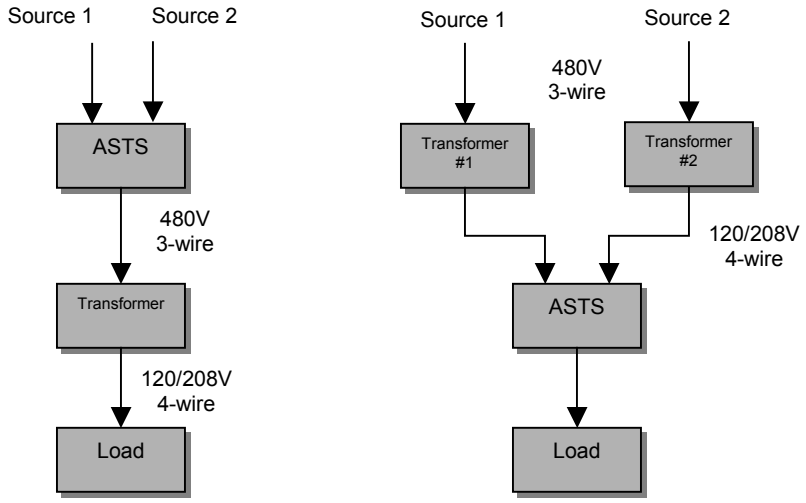

Figure 2. Block diagrams of load side and source side transformer placement.

Regardless of the location of the distribution transformer, either on the source or load side of the solid state switch, the ultimate mission critical loads are subjected to the nearly instantaneous voltage phase switch that appears at the output of the solid state transfer switch. The practices described to mitigate the effect of voltage phase switching on distribution transformers apply with equal efficacy to magnetic components located downstream of the solid state transfer switch and transformer sub-circuit. In the case of modern computing, communications and data storage equipment found downstream of the solid state transfer switch and transformer sub-circuit, SMPS circuits are used

in the overwhelming majority of cases to couple the basic AC input voltage to the internal circuitry of the mission critical equipment.

III. SMPS ARCHITECTURE AND CIRCUITS

Modern SMPS architectures have evolved to a relatively few basic topologies. SMPS provide vastly improved power per unit volume and unit mass compared to linear supplies, and have been refined to high levels of performance. This refinement has lead to the convergence of topologies to a small number of basic layouts.

Most SMPS follow an overall layout as shown in Figure 3. Primary line fuses or circuit breakers are usually included, followed by an EMC filter. The EMC filter is composed C1 through C4 and L1 through L3. These parts comprise high frequency filters to prevent high frequency switching transients from being transmitted out of the supply to the input source. Bifilar inductor L1 is a common mode filter inductor. It is constructed in such a manner that equal magnitude current of opposite direction in each side of the inductor causes cancellation of magnetic flux in the core material. The common mode inductor thus adds no inductance to the circuit for differential currents such as the power carrying currents. The common mode inductor does add significant inductance for common mode current, i.e., current flowing with equal magnitude and direction in each side of the circuit. Switching noise current often appears as common mode current, thus the common mode inductor is a particularly important part of the EMC filter. This inductance, combined with capacitors C3 and C4 form the common mode low pass filter circuit. Inductors L2 and L3 do not share a common core, thus add inductance to the circuit to impede the flow of differential mode current. These inductors are necessarily of very low value, since any impedance added to the circuit by L2 and L3 serve to reduce the voltage available to the bridge rectifier. These components are often either air core inductors, or inductors wound on high frequency ferrite magnetic materials. The inductance added by L2 and L3, combined with capacitors C1, C2 and ½ of the total value of C3 and C4 form the differential mode low pass filter circuit.

After the EMC filter, a passive diode rectifier, D1 through D4, converts the AC voltage waveform to a relatively high DC voltage, which is stored in the filter capacitors, C5. Storing bulk energy at high voltage has the advantage that the amount of energy stored per unit capacitance is relatively high, since the energy stored in a capacitor,

$$E = \frac{1}{2} C V^2,$$

is proportional to the square of the voltage across the capacitor. Following the bulk energy storage capacitor, C5, is a set of one or more DC to DC converters. Most of the variation in SMPS designs occurs in the DC to DC converters, since these sub-systems determine the number of outputs, power capacity, voltage regulation, current limiting, transient response and input to output isolation characteristics of the overall SMPS.

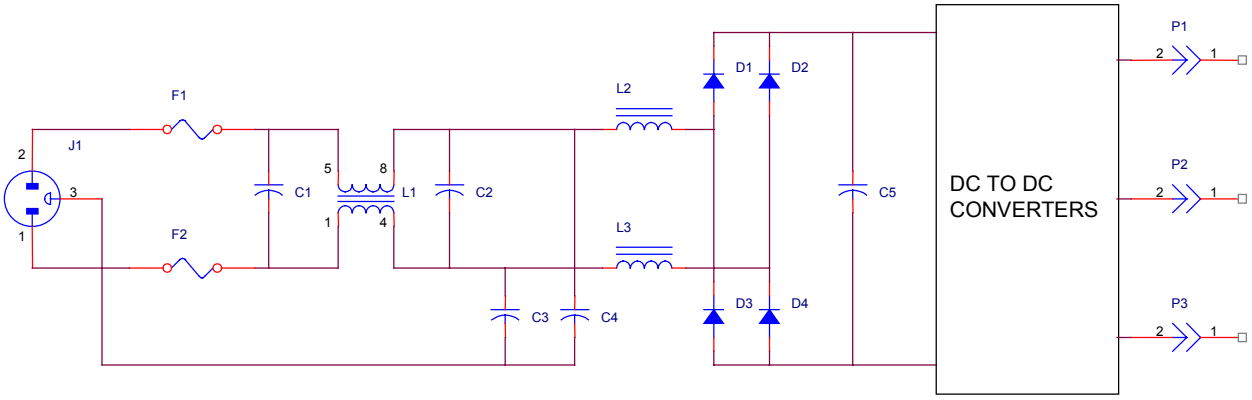

Figure 3. Typical SMPS front end circuitry.

In operation, the primary sinusoidal voltage appears at the output of the EMC filter, where current flows from the primary circuit only when the voltage across the diode bridge exceeds the voltage across C5. Since the DC to DC converters represent a constant power load to C5, the voltage across C5 voltage will decrease as

$$V_{C5} = V_0\sqrt{t_0 - t}$$

whenever the diode bridge is not conducting. When the voltage across the diode bridge exceeds the voltage across C5, the diodes conduct and current is drawn from the primary source. This current charges the capacitors to their peak voltage, V_0. Since the current is drawn in relatively short bursts occurring near the peak of the voltage sine wave, the crest factor and harmonic content of the current drawn by most SMPS are relatively high.

SMPS designs with multiple AC inputs have been accomplished in various ways, however the most sensible architecture simply replicates the EMC filter and diode bridge for as many AC inputs as required. All of the diode bridges feed a common high voltage DC bus and bulk energy storage capacitor. All of the discussions in this paper apply with equal generality to multiple input SMPS that follow the described general architecture.

Most SMPS are designed to operate over a wide range of primary input voltage, often over a range greater than two to one. This results in a similar variation in DC voltage at C5. The DC to DC converters are designed to provide rated output and meet specifications over the entire range of input voltages and resultant DC bus voltages.

IV. USING VOLTAGE PHASE SWITCHED SOURCES WITH SMPS LOADS

Under normal operation from a steady sinusoidal voltage source, the circuit of Figure 3 reacts symmetrically to either positive or negative half cycles of voltage. It has been shown that when used in a data center architecture consisting of unsynchronized redundant sources, switched by a solid state transfer switches, SMPS loads may be subjected to voltage phase switching events. These voltage switching events may approach the worst case situation, from a magnetic component point of view, of a complete half cycle of voltage, followed by another complete half cycle of voltage of the same polarity.

Understanding the reaction of the general SMPS architecture of Figure 3 to the worst case voltage phase switching condition involves examination of each components' reactions to the worst case voltage phase switching condition. It is reasonable to assume that the DC to DC converters will operate correctly whenever the high voltage DC bus is within the limits established by the high frequency transformers, solid state switches and feedback control circuits comprising the converters. That assumption allows a reasonable boundary on SMPS components that are affected by voltage phase switching events to be placed at the connection between the bulk energy storage capacitor, C5, and the DC to DC converters.

The EMC filter components comprising the common mode filter, C3, C4 and L1, are not impacted by a worst case voltage phase switching event. In particular, inductor L1, although often constructed of ferromagnetic materials, is constructed such that the net flux in the inductor caused by SMPS operating current, i.e., current used to charge C5, is zero. This makes the number of volt seconds of flux in the ferromagnetic material zero under differential mode current flow, regardless of polarity. The capacitors C3 and C4 generally store insufficient energy to produce noticeable reactions to any voltage phase switching event.

The differential mode filter components, C1, C2, L2, L3 and to a less extent C3 and C4, are fundamentally insensitive to voltage polarity. The inductors are designed to have very low series impedance, to keep the overall SMPS efficiency high, so they have negligible saturation impact. If they were designed in such a manner that saturation would occur under some circumstances, the most deleterious reaction is the temporary loss of differential mode filtering efficiency of the circuit. Since the inductors are connected in series with the primary current flow, reduction of their impedance during saturation is invisible at power circuit frequencies.

The diode bridge and bulk energy storage capacitor circuit do not react differently to a worst case voltage phase switching event than they react to normal operation. To demonstrate this, two simple SPICE simulations were run. Figure 4 shows the results of the first SPICE simulation run, showing two of the circuit parameters of the general SMPS architecture, the line voltage and the bulk energy storage capacitor voltage. The plot shows that, as expected,

the capacitor voltage increases rapidly near line voltage peaks, and then decays between peaks as the DC to DC converters draw current from the bulk capacitor. Note that the capacitor voltage increases twice per AC line voltage cycle, since the full wave bridge rectifier charges the capacitor equally from both the positive and negative half cycles of line voltage.

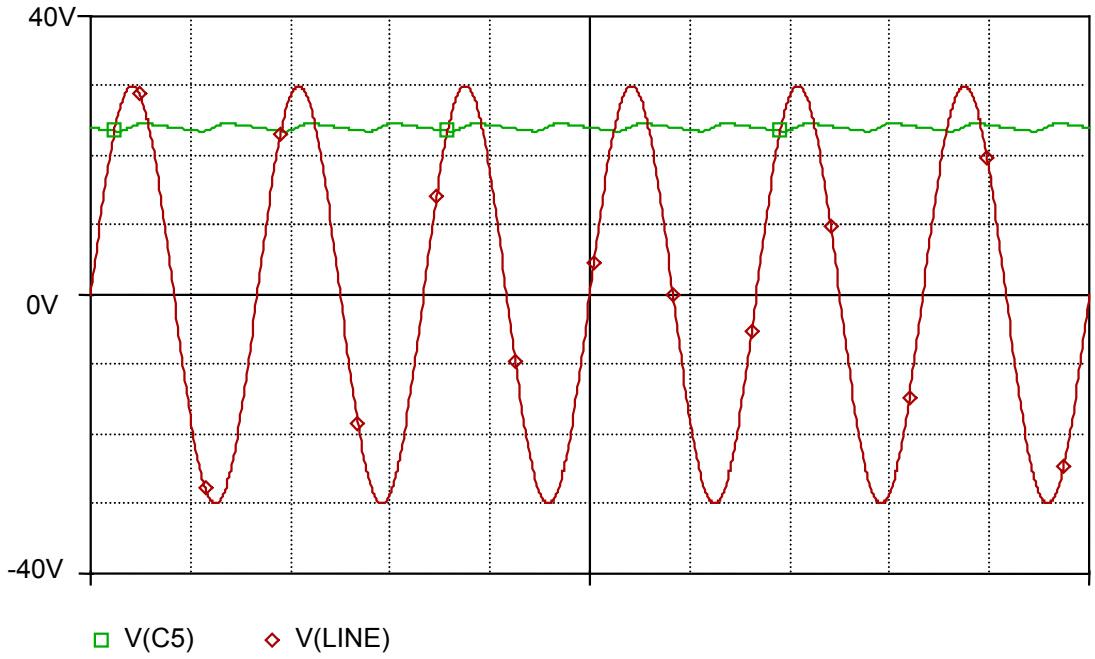

Figure 4. SPICE simulation of line voltage and C5 voltage, no phase switching.

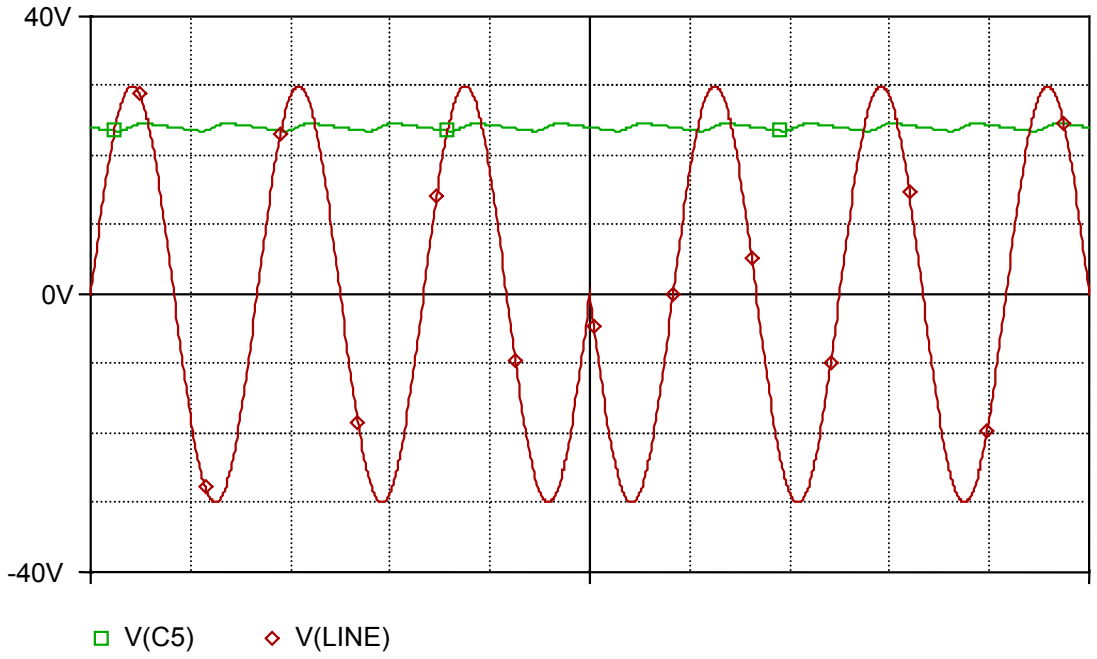

Figure 5. SPICE simulation of line voltage and C5 voltage, worst case voltage phase switch event.

Figure 5 was generated from the same SPICE model as Figure 4, with the single change that a worst case voltage phase switch event was inserted at the center of the data collection segment of the simulation run. The V(LINE) trace clearly shows a 180 degree phase reversal at the center of the plot. The bulk energy storage capacitor voltage, V(C5), is indistinguishable from that of Figure 4. This is the expected result, since the diode bridge masks effect of the voltage phase switch from the bulk capacitor.

At this point in the left to right analysis of Figure 3, the boundary has been reached where further downstream circuits, the DC to DC converters, may be assumed to operate correctly as long as V(C5) remains within their input design bounds. Each of the elements of the general SMPS architecture between the input line and the DC to DC converter boundary has been examined and seen to operate properly under the worst case voltage phase switching situation. From this, it is reasonable to conclude that the general SMPS architectural model of Figure 3 will operate without deleterious results when driven from a voltage phase switched source.

V. LIMITATIONS

It was stressed in the introduction of this paper, and it bears reiteration, that these results are intended to constitute a generalized opinion of how an SMPS designed using typical circuits and standard techniques is anticipated to react when used in a redundant switched power system. This opinion is based on industry norms, standard practices and commonly described circuits and their applications. As such, the conclusions should be useful to decision makers at the data center architectural decision making level. However, no assurance is given that any particular SMPS will react in a specific way. Such assurance can be achieved only by analysis and testing of the specific SMPS under consideration. Additionally, no implication is made that all SMPS will operate correctly or reliably under conditions of switched input voltage phase. Such a blanket assertion would be unsupportable in fact and irresponsible to propose.

VI. SUMMARY

In this paper, the essential problem of switched voltage phase sources was framed and established as a consequence of high reliability, redundant mission critical power system architectures utilizing potentially unsynchronized sources. Magnetic circuit components, such as distribution power transformers, play a major role in architectural decisions due to their characteristic large current draw when subjected to switched phase voltage. The physical mechanism responsible for current draw during magnetic saturation was outlined, along with two circuit strategies commonly used to mitigate the effects of transformer saturation.

Given that solutions are available for mitigation of distribution transformer saturation current, the paper discussed a generalized model of SMPS circuits that interface with their primary power source. A boundary was established in the generalized SMPS topology beyond which the SMPS circuitry was unaware of the phase of the primary voltage and incapable of reacting to it. The generalized SMPS model was then subjected to a walk-through analysis from its primary input to the phase immunity boundary. Each major component was determined to be able to accomplish its function, as well as not engage in disruptive electrical behavior, when subjected to a worst case switched voltage phase event. A simple SPICE model was used to demonstrate that the most important operational parameter of the generalized SMPS model was not affected in a measurable way by a worst case switched voltage phase event. Based on the analysis of the generalized SMPS model, it is reasonable to conclude that SMPS devices operating similarly to the general SMPS model, based on common industry circuits and practices, are generally capable of correct operation under worst-case switched voltage phase conditions, without degradation of the SMPS or adverse effects on surrounding sub-systems.