Impact of Leading Power Factor on Data Center Generator Systems

White Paper 200

Revision 0

by Neil Rasmussen

Executive summary

IT devices may exhibit electrical input current with a characteristic called "leading power factor". This situation may cause back-up generators to become unstable and shut down. Furthermore, a data center that is operating correctly for a long time may suddenly develop a problem as the IT load changes over time, or during an unusual event. This means that it is important to understand the margin of safety and correct for this condition before it happens. This paper explains the problem, why and how it occurs, and how to detect and correct it.

Introduction

Diesel or gas back-up generators are commonly used to allow data centers to operate for extended periods during a mains power loss or for certain maintenance activities. While these are rugged machines that generally tolerate load steps, overloads, and other conditions well, in the data center application there are recorded cases where the generator trips off line under load conditions that appear to be well within the generator ratings. The situation where this occurs is under conditions where the IT load becomes directly supplied by the generator, bypassing the UPS. This occurs when the UPS is in a bypass mode, or if the system is of a tier 3 or other design where one of the power paths to the IT load does not have UPS protection.

The problem is that IT loads may exhibit a characteristic called “leading power factor”, which all generators tolerate poorly. A data center may never experience a problem if there is always an operating UPS in the power path between the generator and the IT load, because in general modern UPS do not exhibit leading power factor and effectively shield the generator from this condition. However, when the UPS is automatically or manually bypassed using either a static or maintenance bypass, the generator is forced to supply the leading power factor of the IT load causing a fault condition at the generator that causes it to shut down, often at the worst possible time, when it is needed most.

A generator can tolerate a small amount of leading power factor, so a data center may initially operate correctly but later develop a problem as IT equipment is slowly changed over time, affecting the power factor.

Unfortunately, there is a great deal of confusion about how to evaluate the risk of this problem, and a superficial analysis of a particular design or installation is likely to significantly overstate the risk. Furthermore, when a risk is identified, it can be difficult to identify the best solution.

This paper starts by explaining the nature of this problem, and how and when it occurs. Then methods for assessing existing facilities for susceptibility to this problem are described. Finally, techniques and equipment for correcting the problem are shown.

Power factor in the data center

To understand the nature of power factor and its various effects, and to provide a foundation for understanding how it impacts generators, some background on the complex issue of power factor is necessary.

In an AC system, the generation source establishes a sinusoidal ***voltage*** waveform, and the load establishes the ***current*** waveform. In the simplest case of a so-called “resistive” load like a heater or incandescent light bulb, the current waveform is the same shape as the voltage waveform and lines up precisely in time with it (**Figure 1**). The product of the voltage and current at every instant of time along **Figure 1** is always positive. By definition, the resistive load is said to have a power factor of 1, meaning that 100% of the load current contributes to the watts of power transferred to the load. In an ideal world, all loads would have a power factor of 1.

However, many types of electrical loads draw some currents that do not contribute to the watts of power transferred to the load. These generally undesirable currents do not transfer watts because they are not aligned in time with the voltage waveform (they are out-of-phase currents), or they are of a different frequency than the source voltage (they are harmonic currents). **Figure 2** shows two waveforms of the same shape (i.e. no harmonic currents) but the current waveform is 90 degrees ahead of the voltage waveform. In this case, the instantaneous product of the voltage and the current is positive half the time and negative half the time, meaning that the power delivered to the load is alternating between positive and negative, with an average watt value of zero[1]. When the current is out of phase with the voltage like this, no average power is delivered to the load.

An actual load current can be separated into three parts: a part that is in phase with the voltage (the part that transfers Watts); a part that is out-of-phase (a part that does not transfer watts); and a part that carries harmonics (also does not transfer watts).

Figure 1

Diagram of voltage and current waveform with the same shape and aligned in time.

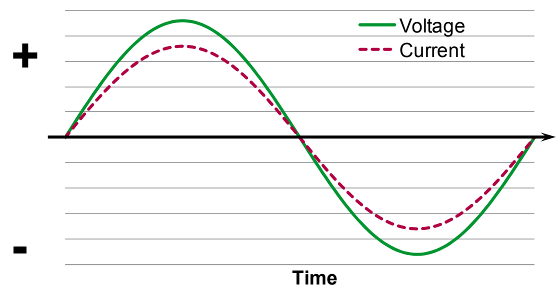

Figure 2

Diagram of voltage and current waveform with the same shape and current 90 degrees ahead of voltage in time.

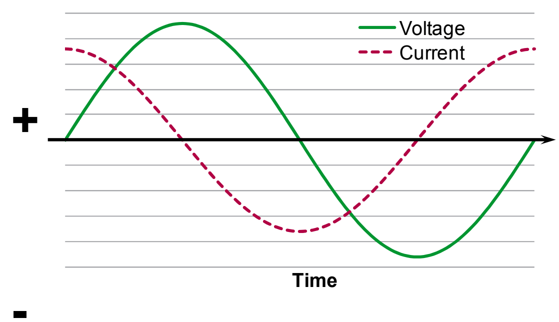

The ratio of the part of the load current that transfers watts to the load divided by the total current (including out of phase and harmonic currents) is called the power factor. Therefore, the power factor of a load is always between 0 and 1, with 1 power factor meaning that all the current is going towards delivery of power to the load, and 0 power factor meaning that no part of the current is going towards delivery of power to the load (i.e. it is all out-of-phase or harmonic current).

Both out-of-phase and harmonic currents contribute to reductions in power factor from the ideal value of 1, and can cause problems for data centers. Both cause wires, transformers, and circuit breakers to need to be oversized to handle the additional currents. However, harmonic currents create additional unique problems such as excessive heating in motors and transformers, voltage distortion, and can cause overheating of the neutral wire in three phase circuits. The out-of-phase part of the current, which is technically defined as the "reactive current", can affect voltage regulation in power systems and is the source of the problem with generators that is the focus of this paper. ***A great deal of confusion regarding the subject of power factor is the failure to recognize that the two different types of currents that contribute to reductions in power factor, namely the harmonic and the reactive currents, can cause different problems.*** Unfortunately, when a load is described as having a "power factor of 0.8" it really tells us nothing about whether the currents causing the value to be different than 1 are harmonic currents or reactive currents. Understanding the effect of power factor on generators requires us to have more specific information about these currents.

---

[1] Another way of thinking of this is that power is delivered to the load at one instant of time but is sent back from the load to the source at a later instant; the power is moving back and forth but, on average, no net power is delivered.

There is another measure of power factor that can help us separate the contributions of out-of-phase current and harmonic currents. The “displacement power factor”, also sometimes called the “fundamental power factor”, is the power factor when the harmonic currents are mathematically removed from the calculation or measurement. By removing the harmonic currents, this measure of power factor only reports the contribution of the reactive or out-of-phase currents. This will be useful in analyzing the impact of power factor on generators, as we will show. However, this particular measurement is only possible with limited types of equipment specifically designed for this purpose.

Reactive currents are caused by so-called “inductive” or “capacitive” loads. With inductive loads like motors, the current waveform is out-of-phase because it is delayed from the voltage waveform. For this reason, inductive loads are said to have “lagging” power factor. With capacitive loads, the current waveform is out-of-phase because it is actually earlier in time than the voltage waveform. For this reason, capacitive loads are said to have “leading” power factor. A power factor value, such as a power factor of 0.9, does not convey any information about the timing of the out-of-phase currents that caused it, so when discussing power factor related to reactive current it is often further clarified by describing the power factor as “0.9 leading” or “0.9 lagging”. This distinction is particularly important in describing how power factor affects generators, as we will explain.

Harmonic currents are caused by so-called “non-linear” loads that draw a current waveform of a different shape than the supplied mains voltage waveform. The most common non-linear loads are light dimmers or power supplies. Historically, early IT power supplies had a power factor of approximately 0.65. The harmonic currents drawn by with these power supplies were often **more than half** of the total current drawn and caused the reduced power factor (the out-of-phase or reactive current of these early IT power supplies was not significant). These harmonic currents began to be a problem when large numbers of IT devices were located together, such as in call centers or data centers, resulting is overheating neutral wiring and transformers as mentioned previously. In extreme cases, fires resulted. In response to this problem, international regulations were established in the 1990s to limit the allowable harmonic currents caused by IT power supplies, and all modern IT power supplies are equipped with a “power factor correction” function, greatly reducing harmonic currents to levels that do not impact power distribution systems. What is important to understand about this is that “power factor corrected” IT power supplies are intended to eliminate **harmonic** currents, but were not intended to fix **out-of-phase** currents. It is a remarkable fact that power factor corrected IT power supplies can actually create MORE out-of-phase current than the older uncorrected power supplies did. ***It is these out-of-phase reactive currents, that are present in power factor corrected power supplies*, that impact generators.** Before considering the impact on generators, it is useful to understand how power supplies create reactive currents.

Power factor characteristics of IT loads

Almost every power factor corrected IT power supply works on the same principle, shown in **Figure 3**. Basically, a power factor corrected supply is simply a traditional power supply with an additional “PFC boost module” that acts as a “filter” to shield the mains from the harmonic currents that a traditional power supply normally creates.

The PFC boost line module “filter” is actually a high frequency power converter, which shapes the input current to the form of a sine wave in much the same way that a modern audio amplifier does. The switching frequency of this power converter is typically in the range of 20 to 200 kilohertz, and these power factor correcting circuits are very effective at eliminating the harmonic currents that would otherwise be present. These converters by their nature have high voltage, high frequency switching circuits, and require a filter comprised of inductors and capacitors at the mains input in order to prevent the high frequencies, which are a form of radio frequency interference, from being passed from the power supply back into the power mains. A key part of this filter is a capacitor shown as C1 in **Figure 3**. It is this capacitor input filter that creates out-of-phase current, reducing the power factor from its ideal value of one. ***Even if a power factor corrected power supply is perfect at eliminating all the harmonic current, this capacitor still creates out-of-phase current and the power factor of the supply will be less than one.***

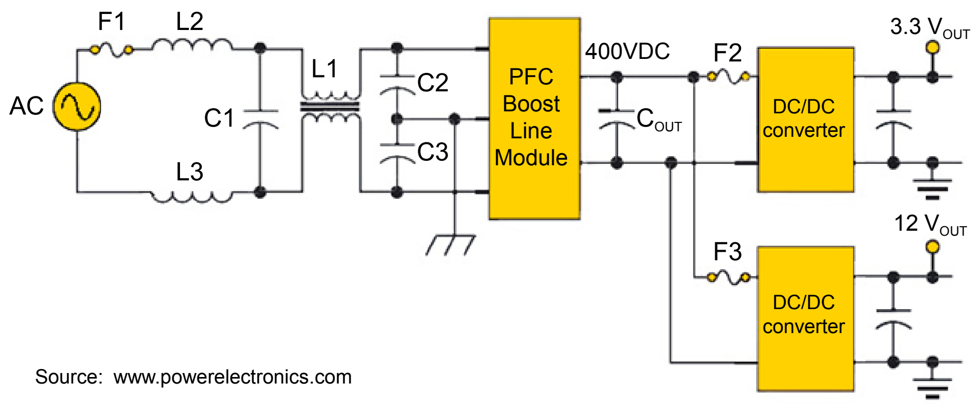

Source: www.powerelectronics.com

Figure 3

Diagram of a power factor corrected power supply, showing PFC boost module.

It is important to understand that the same capacitor is present at the input of a power factor corrected supply even if the supply is operating at light load, and in many cases even if the IT device is switched off. Therefore, the out-of-phase current caused by the capacitor is simply a fixed property of the IT device and is independent of the actual power consumed within the IT device. This means that the power factor of a specific IT device will actually get worse as the power consumption in the IT device falls; the current that contributes to the watt load declines but the capacitive out-of-phase current stays the same. This effect is shown in **Figure 4**.

Figure 4 shows that power factor gets worse as power supplies are lightly loaded. If a server were to be fed from two power paths via two internal 600W rated power supplies and the actual server watts were 300 watts (150W on each path), then it would be operating at the 25% utilization point in the figure, where you can see the power factor is much less than would be expected from the fully loaded power factor rating. In general, any time power supplies are over-provisioned or underutilized we end up with more capacitors per kW of IT load, creating more and more out-of-phase reactive currents as a percent of total load.

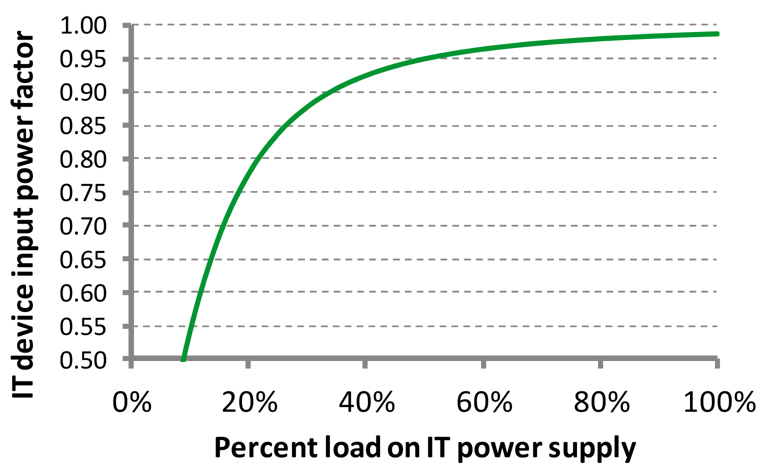

Figure 4

Mains power factor of the power supply of an IT device as a function of internal IT device power draw.

The data from **Figure 4** are for a power supply exhibiting a capacitance of 10uF per kW of power supply capacity. This is at the high range of observed capacitance of real IT equipment. Therefore the power factor curve shown in the figure is about the lowest curve to be expected in real servers, and typical servers will exhibit a curve of the same shape but higher. A collection of data from real power supplies is shown in **Figure 5**:

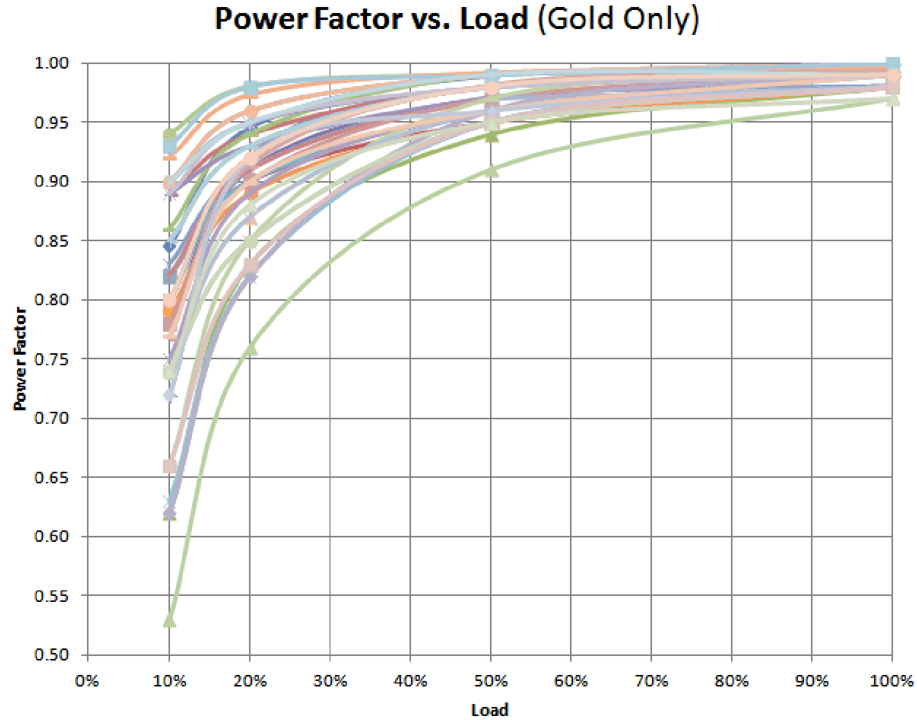

Figure 5

Measurements of actual power supplies showing the variation of power factor with load (Illustration from 80PLUS.org)

Note that **Figure 5** shows that all power supplies have a power factor that falls off at light load, exhibiting the curve shape expected for power supplies with different input capacitance per kW. The power supply in **Figure 5** that has the lowest power factor shows a curve corresponding to the curve in **Figure 4** of a power supply with a capacitance of 10uF per kW of power supply capacity. The higher curves in **Figure 5** exhibit lower capacitance. In the next sections we will show that power supplies exhibiting the higher curves in the figure cannot cause generator instability, but power supplies exhibiting the lower curves can result in instability.

The input capacitance is not a specified property of a typical IT device, and can vary considerably between devices as seen in **Figure 5**. Therefore, it is very difficult for an operator to determine how much out-of-phase current will result from a specific deployment of a mix of IT devices. Even measuring the actual power factor of an installation is not sufficient because this does not distinguish reactive currents from harmonic currents. To determine the out-of-phase current in a specific installation requires specific measurements that can distinguish between out-of-phase and harmonic currents; these will be discussed in a later section of this paper.

It is difficult to believe that these small capacitor filters in IT power supplies could disrupt a multi-megawatt generator. However, these capacitors add up and become quite significant, especially when the following conditions occur:

- when thousands of IT devices are placed together, as in a data center
- when the rated capacity of the IT device internal power supplies is significantly larger than the watt load of the device (under-configured IT devices, or over-provisioned power supplies)
- the total load of IT equipment in the data center approaches the full load design rating (more capacitors on the generator)

These situations are further compounded by the fact that generators are quite sensitive to out-of-phase currents, specifically to the "leading" currents caused by capacitors. This phenomenon will be explained next.

Effects of leading currents on generators

The physics of generators are complex and a complete description of how they work is beyond the scope of this paper. However, to understand how power factor affects generators requires at least a superficial understanding of some key principles.

A generator has is a spinning electromagnet that is constantly passing by the fixed armature wire coils as shown in **Figure 6**. The magnetic field passing by the stator coils creates a voltage at the output of the coils, which is the output voltage of the generator. The generator controls its output voltage by constantly measuring it and adjusting the current fed to the spinning electromagnet. If a load is suddenly applied to the generator and the voltage droops, the regulator acts to slightly increase the electromagnet current to maintain the desired output voltage.

Figure 6

Diagram showing the basic operating principle of a three phase synchronous generator

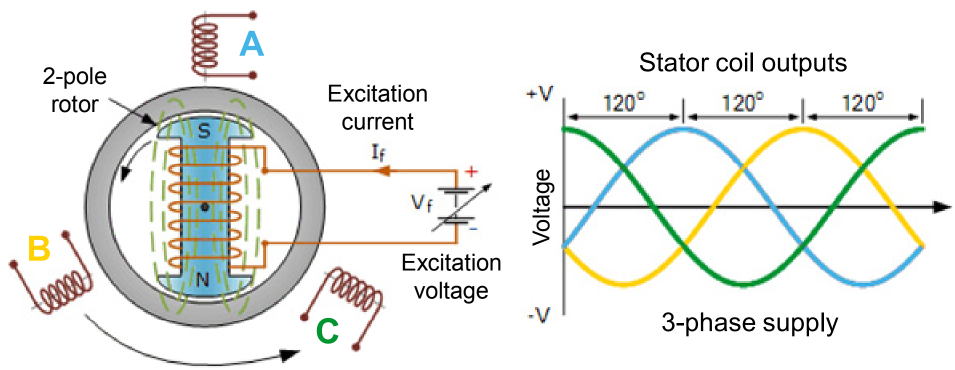

Three stator windings 120° apart

When the load on a generator has an output power factor of one, the magnetic field strength of the spinning electromagnet is determined solely by the excitation current, which is controlled by the generator regulator. However, when the load has some out-of-phase current, this current gives rise to a magnetic field within the generator rotor that either adds to or partially cancels the magnetic field of the spinning electromagnet. Out-of-phase current in the load therefore creates extra work for the regulator system; for example, if the load has a lagging out-of-phase current this partially **cancels** the field of the rotating electromagnet, and the regulator must increase the current to the electromagnet to compensate. Conversely, if the load has a leading out-of-phase current this actually **adds to** the electromagnet field, and the regulator must reduce the excitation current to the electromagnet to compensate.

The voltage regulator in a generator is capable of maintaining the output voltage over a wide range of load conditions including various amounts of leading and lagging out-of-phase current. However, there is a special problem with leading out-of-phase current. As the leading out-of-phase current is increased, it adds to the strength of the field of the rotating electromagnet, and, as previously explained, the regulator decreases the current supplied to the electromagnet to compensate. This can continue until the leading out-of-phase current becomes large enough so that it boosts the electromagnet field to the point where the regulator no longer needs to supply ***any*** current to the electromagnet. It is a limitation of virtually all generators that the regulator cannot supply a negative current to the electromagnet, so it cannot compensate further when the leading out-of-phase current exceeds this point. If the leading out-of-phase current passes beyond the point where it is able to supply the entire field to the spinning electromagnet, the regulator shuts off. At this point, the field of the rotating electromagnet will increase beyond the desired value, and because the output voltage is controlled by the field, the output voltage of the generator will begin to rise in an uncontrolled way. The result is that the safety systems of the generator will detect an overvoltage condition and trigger an immediate shut down the generator.

This then is the so-called "power factor" problem with generators. This problem is not caused by harmonics. The problem is not caused by lagging reactive current. The problem is specifically caused by leading reactive current greater than a certain threshold value. When leading out-of-phase current passes a threshold value, the generator will lose control of its output voltage regulation and will trip off on overvoltage. It is important to understand that while reactive current does result in reduced power factor, it is the size of this out-of-phase current, **and NOT a specific value of power factor**, that causes the problem. We will now examine this it more detail.

Every generator has specifications regarding the acceptable operating conditions. The specification regarding acceptable out-of-phase current for a specific generator is given in a "capability curve" as shown in the example of **Figure 7**.

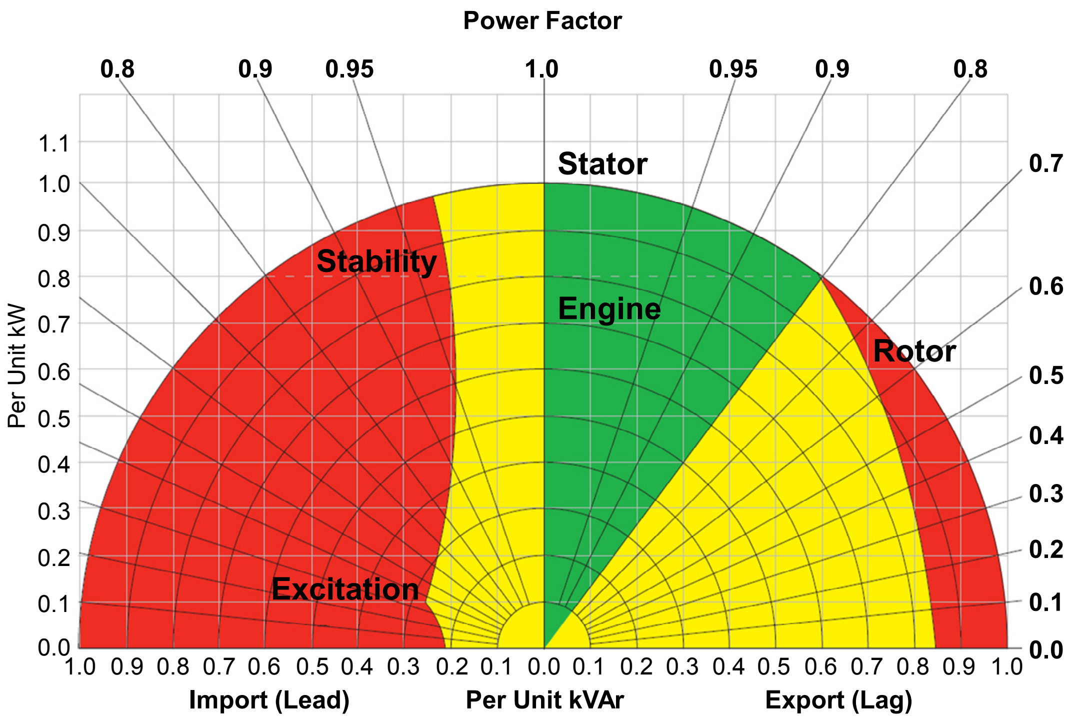

Figure 7

Sample of a generator capability curve showing acceptable operating conditions for load and power factor

The green area of this curve is considered the normal operating area, the yellow areas are acceptable but not normal, and the red areas are not possible. The horizontal axis is the magnitude of the reactive current of the generator load, with leading reactive current to the left. The vertical axis is the watts of generator output. The term "per unit" can be read as "percent of rated". The radial lines are lines of constant power factor. For this discussion, we will focus on the yellow area immediately to the left side of the center of the curve.

The important information in this curve is the shape of the yellow area on the left side. Note that the leftmost boundary of this area is crudely vertical at approximately 0.2 per unit, representing the condition where the leading reactive current is approximately 20% of the rated generator load current. Above this level of leading reactive current, the generator becomes unstable and shuts down. This curve is sufficiently representative of typical generators so that the following generalization is possible:

As a general rule, generators must be operated with leading reactive current of less than 20% of max rated current in order to avoid instability.

Note that this rule does not mention power factor. In fact, looking at the power factor lines in the compatibility curve above shows that the power factor at the boundary of instability varies widely, from 0.0 when there is no watt load, up to about 0.97 at full load! This is more clearly shown in **Figure 8**.

The curve clearly shows that lightly loaded generators are quite tolerant of almost any power factor, but as the load increases, the power factor must be maintained above a critical value, which reaches about 0.97 at full load. This suggests a key insight: **the instability problem is unlikely to occur on a lightly loaded data center.** It also explains why a data center generator system might be stable for many years and then become unstable when the IT load grows over time.

Figure 8

Lowest acceptable leading power factor for a generator as a function of load, below which generator will shut down

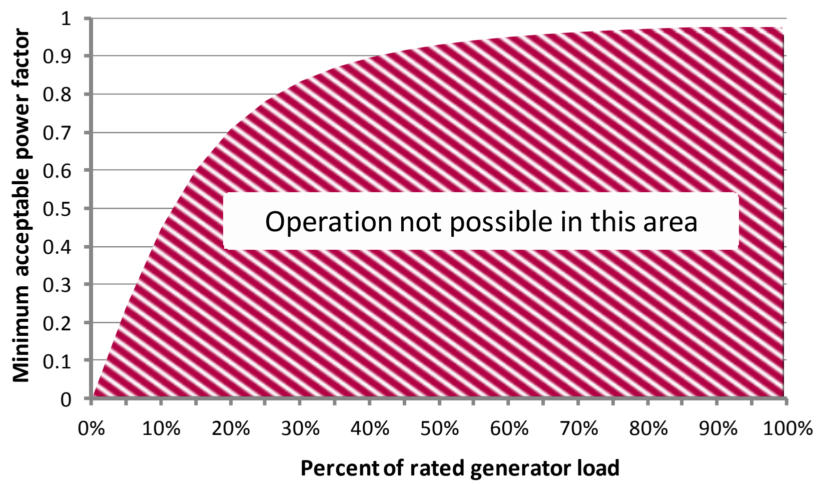

Note that these acceptability curves do not include the contribution of harmonics currents to the power factor. If there are harmonic currents, the power factor lines on the curves of both **Figures 7** and **8** are invalid and cannot be used, and only the percent leading out-of-phase current determines the instability boundary.

This leads to the important insight that **the instability problem SHOULD NOT be quantitatively described in terms of a specific power factor**, but rather in terms of percent of leading reactive current. *Many operators experiencing this problem (or concerned regarding it) attempt to use power factor as a measurement or design criteria and reach mistaken conclusions or are otherwise confused by this issue*. For the remainder of this paper, we will therefore discuss the problem quantitatively in terms of percent out-of-phase leading reactive current.

Example case

The issue of generator instability is only of practical importance if the leading reactive current from the IT load can become sufficiently large to exceed the reactive current limit of the generator. Therefore, it is instructive to consider examples of real IT equipment to see if and when this situation might occur.

For a test case, we will assume a data center with an IT load rating of 1MW. We will assume that this data center is equipped with a generator rated at 2MVA. Then we will populate this data center with increasing percent load of IT equipment and see when instability might occur. We will assume that the IT loads are directly supplied by the generator, without an intervening UPS (or if there is a UPS, it is in bypass). We will assume there are no other significant loads other than the IT loads for this example (the contribution of these loads will be discussed later in the paper).

For real IT equipment, there is significant device-to-device variation in the amount of out-of-phase leading current; even two servers rated at 1kW each may draw out-of-phase current that differs by a factor of 5 or more. Recalling the earlier discussion, the servers have input filter capacitors that are the primary source of this out-of-phase current. We find that typical IT devices exhibit a capacitance range of 1-10 microfarads per kW of power supply rating and will use this for our example case[2]. The power supplies are assumed to operate at 230VAC. We can compute the out-of-phase current resulting from this capacitance, which adds up as more kW of IT load is added to the data center. We also need to consider that IT devices are almost always configured to draw less power than their power supply rating, and we will assume IT device power supplies are loaded to 40% of rating for the test case.

Figure 9 shows how the out-of-phase leading load current increases as IT devices are added to increase the IT load in the example data center. When the reactive leading current reaches 20% of generator rating, the generator will become unstable.

Note that in this example case, the curves for 2 and 5 uF per kW never reach the 20% leading current limit, and so will always be stable. However, the curve for 10uF per kW becomes unstable as the IT deployment reaches 80% of rated data center load.

The above example was for a data center with a single power path. If the same data center with a 1MW IT load rating has two power paths and dual corded IT devices, there are up to twice as many power supplies in the system, and in a dual path system there are conditions where all of the supplies in both paths can end up powered from the generator source for one path[3]. Under this condition, the same IT configuration results in additional leading current as shown in **Figure 10**.

Figure 9

*Leading current applied to a generator in a **single path** data center as a function of IT load, for different IT device capacitance characteristics.*

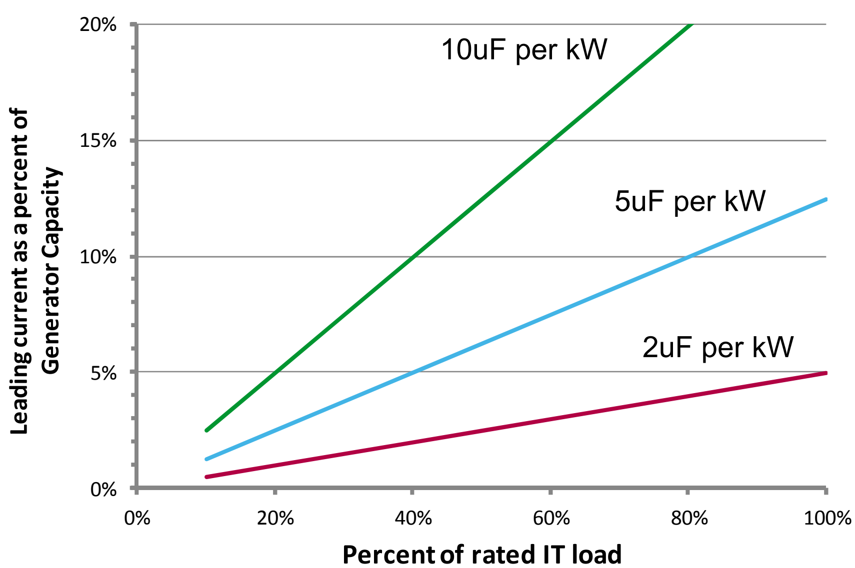

[2] Although we believe based on sampling that 98%+ of IT power supplies exhibit capacitance within this range, power supplies with larger or smaller capacitance can exist. We speculate that a few models of very small servers may exhibit unusual or larger capacitance. Therefore, if large numbers of small servers of the same type are deployed, it is recommended that testing be performed to determine actual capacitance values, rather than assuming typical values. Methods for testing are described later in this paper.

[3] It is possible to create data center architectures where the power supplies on both paths can never become connected to the generator on one path. This is discussed in the later section on mitigation.

Figure 10

*Leading current applied to a generator in a **dual path** data center as a function of IT load, for different IT device capacitance characteristics.*

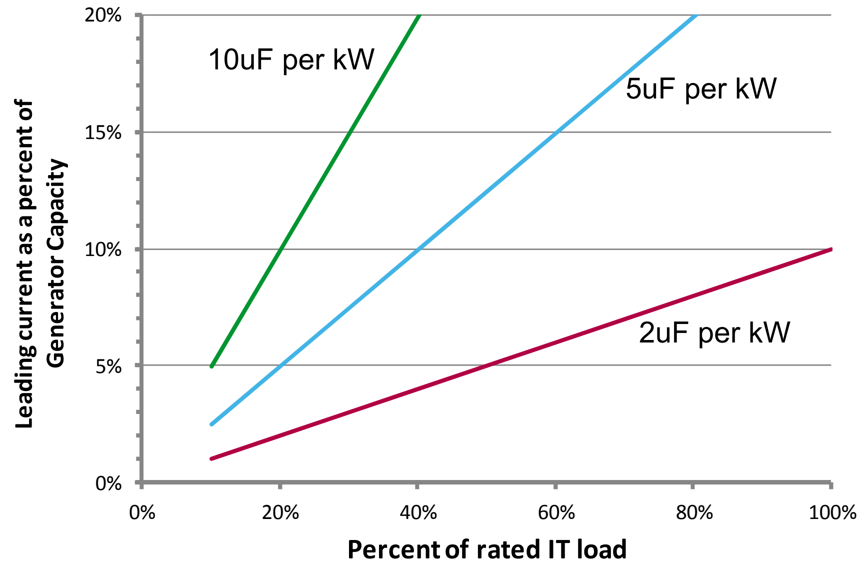

In the case of a dual path system, the possibility for the leading current to pass 20% of the rated generator current and cause unstable operation occurs in all cases except where the capacitance of the IT devices is 2uF per kW. If the IT devices exhibit 10uF per kW, then the data center cannot power more than 40% of its design load without becoming unstable, if that IT load becomes powered by the generator.

If different assumptions are made for the test case, the possibility of instability is increased. Specifically, if servers have less than 40% utilization of their power capacity, the problem gets worse. If the generator is closer to the IT rating, the problem gets worse. However, if other loads that don't exhibit leading power factor such as air conditioners and lighting are powered from the generator, the situation is improved.

The following important conclusions can be drawn from this example:

- Generator instability due to leading power factor is a real problem that all data centers should consider and take action to prevent
- A data center loaded below 20% of IT load will almost always be stable
- A dual path data center loaded to nearly 100% capacity has a high likelihood of becoming unstable if the entire IT load becomes connected to the generator, with unexpected generator shutdown
- To determine whether a generator will be stable requires information about the input capacitance of the IT devices, either through measurement or by manufacturer's specification (if available).
- If all IT devices (or even the power weighted average of IT devices) have capacitance of 2uF per kW or less, then there is no instability problem in any realistic design

Determining capacitance of IT loads

The input capacitance of an IT power supply is not normally a specified parameter. Even though IT device manufacturers do not publish input capacitance specifications, it is often possible to obtain measured power factor specifications of IT device power supplies from the manufacturer or from an independent data-reporting source such as 80PLUS.org. If we assume that all of the current that contributes to reducing the power factor from one is caused by the input capacitor (the safe, worst case assumption), then it is possible to establish the worst case (largest) value for the capacitance per kW. It is desirable to obtain the power factor reading at a point such as 50% load as this is more representative of typical operating points, and the resolution for establishing the capacitance value is increased. Given the reported or measured power factor of a power supply at 50% of rated load, such as those provided by 80PLUS®, the capacitance per kW of rated IT load power can be determined from **Figure 11**.

Figure 11

IT device power supply capacitance per rated kW as a function of reported or measured power factor at 50% load

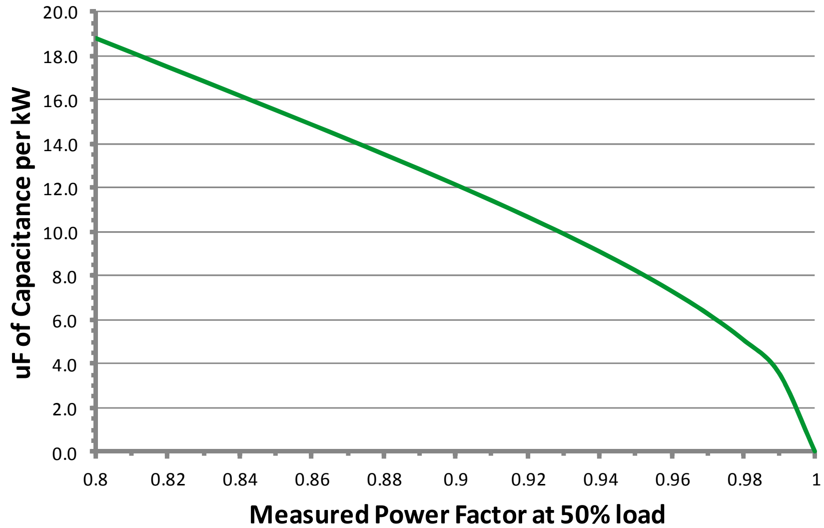

Expressing the capacitance in uF per kW normalizes the capacitance value and allows the percent out-of-phase current to be conveniently established for a deployment of any kW size, using previous **Figures 9** and **10**. For example, **Figure 11** above shows that a server with a 0.92 power factor at 50% load has a capacitance of 10uF per kW; if an entire dual path data center is populated with these servers, then previous **Figure 10** shows that the generator is may become unstable if that data center exceeds 40% load.

It should be noted that for the recent data reported by manufacturers for high efficiency power supplies, the typical power factor at 50% load is around 0.93 to 0.99, suggesting an average value of about 2-10uF per kW for recent generations of power supplies. The published data also suggests that smaller power supplies generally have a higher capacitance per kW. In general, this suggests that a data center with a very large number of 1U servers is more likely to be unstable on generator than a data center that has the same watt load of fewer larger servers.

Knowing the capacitance per kW of a power supply is an important part of predicting how much out-of-phase current will be created per unit of power supply capacity installed. Power supply capacity can then be translated to other convenient metrics, such as reactive current per server, or reactive current per IT kW. The mathematics of these computations will not be presented here.

Assessing risk of instability

The risk of instability can be assessed for a current data center under existing and expected future conditions, and it can be assessed for data centers that are still in the planning stages. The process is different for each.

Risk in operating data centers

To assess the risk of generator instability in a currently operating data center, it is necessary to determine the leading out-of-phase current that the generator may be called on to supply, and then compare that to the current limit (which is typically 20% of the generator full load current rating). The best way to do this is to actually measure this current, but that is typically not practical, because there may not be instrumentation in place to make the measurement, and instrumentation often is not capable of directly making the required measurement.

If a measurement is proposed to be made during a generator test, remember that the generator load during a generator test is typically the UPS, and not the IT load during that test. No known UPS system has enough leading out-of-phase current to affect a generator, so generators will always pass this test without instability. As explained earlier, the problem we are concerned with is when the UPS is in bypass, so that the IT load is directly powered from the generator. In some data centers, this mode may never have actually occurred before. Therefore, forcing the UPS into bypass when on generator in order to connect the IT load to the generator may trigger instability and is not a recommended procedure. In an operating data center, we should try to measure the out-of-phase current of the IT load at the UPS output and then check to make sure it is less than the generator limit, before ever applying the IT load to the generator without an intervening UPS. This measurement may be complicated if there are multiple separate UPS with separate output busses connected to a generator.

The problem with most measurement devices is that they include both harmonic currents and out-of-phase currents combined, and therefore do not accurately report the reactive currents unless the harmonic currents are zero. Unfortunately, many data center do have enough harmonic currents to affect the measurement. Meters that report VARs (Volt-Amps Reactive) or Reactive Amps are **not** typically reporting the needed value of the out-of-phase current[4] unless they specifically specify that they are reporting only the value of the "Fundamental" or base mains frequency or they are certified to implement the official definition of reactive power as defined in IEEE Standard Dictionary 100-1996. If power instrumentation reports "displacement power factor" then this can be used to compute the reactive current. The key is to ensure that measured values exclude harmonics so that the true fundamental frequency reactive power can be either measured or computed. Again, it is important to re-emphasize that most equipment does not correctly make this measurement, so it is critical to ensure that any measurements are made using equipment double-checked to verify that they are not including harmonics and can correctly distinguish leading reactive current.

Examples of meters that correctly report displacement power factor are PowerLogic® PM820, PM850, and PM870 (**Figure 12**). Meters may be installed either temporarily or permanently to make the required measurements.

Figure 12

Power meter capable of measuring displacement power factor

It is often not possible to find a single measurement point that consolidates all the loads. The reactive currents can be measured separately on different parts of the data center and added. If displacement power factors are measured, those measurements must first be converted to reactive currents and then added.

[4] The technically correct implementation of the reactive power computation is difficult to implement in electronic systems at reasonable cost. It typically requires complex DSP processing. Most meters use simplified methods to compute reactive power or current that overstate the out-of-phase current when harmonics are present.

Risk in planned data centers

In a data center that is in the planning stage, no measurements are possible. Nevertheless, it is possible to draw conclusions regarding the leading reactive power susceptibility of the generator system.

First of all, it is possible to greatly reduce or eliminate the risk of instability via design choices. For example, if a dual path system is designed so that all the IT supplies on both paths cannot end up connected to and fully loading a generator, the risk of instability is much lower as described previously. For example, a dual path system could be designed so that this combination is not possible in normal or maintenance operations. For example, if UPS systems on both paths are bypassed either automatically or manually, the system could be interlocked so that an upstream maintenance cross-tie cannot connect both paths to a single generator. This is just one example, and different architectures may have different approaches.

Furthermore, a planned data center can ensure that meters are installed to report and sum all of the leading reactive IT load currents (using the correct type of meters) and set an alarm if it exceeds a value set slightly below the generator instability threshold.

In addition, a planned data center can compute the acceptable level of IT device capacitance, and establish policies that either exclude high capacitance devices, or limit the penetration of high capacitance devices. For example, a policy that states that any single device type comprising, for example, 10% or more of the total data center load must have a capacitance of less than 5uF per kW, and the power supplies in these devices must not be grossly oversized or underutilized.

Mitigation

It may be determined that action must be taken to correct a currently unstable condition, or to ensure that the data center does not enter a future unstable condition. The options to address a current condition are examined first, and then some additional options to prevent future problems are described. It is very important to note that not all data centers require mitigation, so mitigation is not necessary unless a risk analysis finds a current or impending situation requiring correction. The following are the options for correcting a situation of excessive leading out-of-phase current:

Remove load from the generator

When load is removed from a generator, sources of out-of-phase current are removed. The generator stability is assured when the out-of-phase current is less than the chosen target value (less than the instability threshold of 20%, such as 15% of rated generator current). Ideally, the IT devices with the highest capacitance per kW would be removed first. Note that IT virtualization projects typically reduce the load and almost always eliminate large numbers of small servers, which are often the worst contributors of out-of-phase current. ***Therefore, a virtualization project may actually be an opportunity to not only reduce load, but to change the nature of the load to IT equipment with less capacitance per kW. Such projects may temporarily or even permanently eliminate the problem.***

Install an inductive load bank

An inductive load bank is simply a group of large inductor coils placed on the generator bus. This provides a fixed amount of lagging out-of-phase current that can cancel the leading current due to the IT loads. The load bank could be applied either at the generator output or after the UPS. Because the load bank has losses, it should be at the output of the generator (which is normally de-energized) in order to maximize energy efficiency.

It is also possible to install an inductive load bank that has automatic switches, which adds or removes inductors as needed to balance the leading current. The idea is that if the leading current may change over time, then inductors are switched in and out in order to hold the power factor close to a desired target value, such as 1. However, this degree of fine tuning is not necessary and the cost and complexity are difficult to justify.

A practical plan is to have two basic load banks fed by manually switched breakers, rated so that when used together they assure stability of the system. For example, two inductive loads, one rated at 5% and one at 10% of the generator current rating can be used to apply 0, 5, 10, or 15% lagging out-of-phase current to the generators, which would overcome leading power factor in almost any imaginable case. The units could be manually switched based on the operating load of the data center. Smaller rated inductive load banks may be possible depending on the risk analysis. Such switched inductive loads are not a standard product but can be built to specification by a number of manufacturers[5].

A concern with inductive load banks is their behavior when they are energized. In general, an inductive load bank may have an inrush current due to magnetic saturation, unless it is specifically designed with an oversized core to never saturate (which adds significantly to the cost and weight). Therefore, it is desirable to attach a load bank upstream of any circuits that might be automatically switched. The ideal location for a load bank is right at the generator output.

It is sometimes suggested that the inductive load bank be placed downstream of the UPS so that both the generator and the UPS see less leading out-of-phase current. This presumes that the load bank will solve problems that a UPS may have with leading power factor. It must be noted that some problems that UPS systems may have with leading power factor are actually caused by other effects caused by the existence of the capacitors in the load[6], and these capacitors still exist and can still affect the UPS even when an inductive load bank is installed. Placing the load banks downstream of the UPS subjects the load bank to various switching conditions, such as going to bypass, that can cause undesirable transient conditions, unless the load bank is specifically designed not to have inrush due to saturation.

Install an electronic power factor correction system

This is a device, known as an active filter, that is capable of correcting power factor and preventing lower order harmonics from passing upstream to the power source. Many models offer various modes of operation, and some offer the ability to cancel leading out-of-phase current. This device continuously attempts to hold the power factor at 1, by cancelling out-of-phase current and harmonics at the same time. There are no switching transients caused by these systems. Because they continuously correct the power factor, no ongoing checking is required. They are also much smaller and lighter than inductive load banks. An example of such a device typically used in data centers is shown in **Figure 13**.

---

[5] An example of a company capable of providing passive inductive load banks to a specification is Simplex, Inc of Springfield, Ill, USA

[6] The problem of arcing or pitting of switch or breaker contacts can be caused by high capacitance in an IT load. This is due to the capacitance, and not the power factor. Therefore, correcting the power factor does not solve this problem.

Figure 13

Electronic power factor correction system capable of correcting leading power factor and rated 300A (Schneider Electric Accusine)

The device shown in the figure is capable of providing sufficient cancellation of leading out-of-phase current for most 2MVA generator applications. Electronic power factor correction devices like the one in the figure can also correct certain undesirable conditions due to harmonics at the same time as it corrects out-of-phase current. The electronic power factor correction device is an effective solution when it is desirable to correct both leading power factor and harmonic problems[7]. Active filters offer more benefits compared to passive inductive load banks but cost more; about $180/kVAR compared to $70/kVAR for a passive solution. If sized to offset a 15% reactive current, this cost translates to less than 5% of the generator cost.

Identify worst offending IT devices and replace them

The capacitance of IT devices varies significantly. In some data centers, a large deployment of a specific device with high capacitance may be responsible for a considerable fraction of the total out-of-phase current. If that device type can be identified, consider replacing that device with a newer different device with lower capacitance. This may work conveniently into an IT refresh or virtualization program.

The above options comprise the list of alternatives for correcting excessive leading out-of-phase current that can lead to generator instability. When the data center is not yet designed, or when the data center has not yet reached a critical condition, the following additional options are available:

Select IT equipment vendors based on power supply capacitance

When selecting vendors for any IT device that may be deployed in high quantities in the data center, consider input capacitance as a selection criteria. If large server deployments have input capacitance that is near the low end of typical values, then generator stability can be assured under any operating conditions. Because the capacitance is not a typically specified parameter, low capacitance is established through evaluation and acceptance testing or by performance guarantee by the server vendor. A maximum value of 4uF per kW of rated power supply load, which would be assured by a measured input power factor of 0.99 at 50% rated load, would be a practical target requirement.

Avoid using very large quantities of small servers

Smaller IT devices like 1U servers typically have higher capacitance per rated kW of IT load. Therefore, an IT design that uses fewer larger servers generally (but not always) has a lower capacitance.

[7] See White Paper 38, *Harmonic Currents in the Data Center: A Case Study* for more information about harmonic issues in data centers

Design for increased power supply utilization

As explained before, the leading reactive power is proportional to the IT power supply rating, and independent of the actual IT watt load. Therefore, an underutilized or oversized power supply has excess input capacitance, compared to what the IT load requires. When a large-scale deployment of identical servers is planned, it is worthwhile to study the power supply utilization of any alternative server configurations, and try to avoid unnecessary oversizing of power supplies.

Do not allow two power paths on one generator

There are a variety of dual power path architectures used in data center design. In most designs, dual-corded IT devices have internal power supplies that share the IT load, and are sized so that if one path goes down the second path can pick up the entire IT load. Often, this means that one path may be suddenly subjected to a 2X load step as it picks up the load, but each path still is only physically connected to the power supplies connected to that path. Therefore, under this mode of operation the load capacitance is unchanged, even though the load power doubles. However, there are modes in many designs where all of the power supplies on both paths may be combined onto one of the paths; the most common cause of this is the operation of a “cross-tie” in the power paths. When a cross-tie is activated, the doubling of the connected power supplies causes the doubling of capacitance on the supply path, which can drive the system into an unstable condition when on generator. This can be a good reason to consider designs that do not have cross-ties.

Conclusion

The problem of generator instability in data centers has been explained, and shown to be caused by leading reactive currents generated by IT devices. In normal operating conditions, instability almost never occurs because there is a UPS interposed between the generator and the IT devices, which shields the generator from the leading reactive current. The problem occurs when the IT devices become directly connected to the generator, such as when the UPS is in bypass or for various maintenance conditions. Therefore, instability may not reveal itself until a transient event or during maintenance operations, which can be a highly unwelcome surprise.

This paper provides practical guidance regarding how to assess an existing or planned data center for potential generator instability, and strategies to prevent the possibility of instability. Analysis of potential instability should be an expected part of any effective data center design, and also part of an effective data center management plan.

About the author

Neil Rasmussen is a Senior VP of Innovation for Schneider Electric. He establishes the technology direction for the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks.

Neil holds 25 patents related to high-efficiency and high-density data center power and cooling infrastructure, and has published over 50 white papers related to power and cooling systems, many published in more than 10 languages, most recently with a focus on the improvement of energy efficiency. He is an internationally recognized keynote speaker on the subject of high-efficiency data centers. Neil is currently working to advance the science of high-efficiency, high-density, scalable data center infrastructure solutions and is a principal architect of the APC InfraStruXure system.

Prior to founding APC in 1981, Neil received his bachelors and masters degrees from MIT in electrical engineering, where he did his thesis on the analysis of a 200MW power supply for a tokamak fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.

Resources

Watts and Volt-Amps: Powerful Confusion
White Paper 15

Fundamental Principles of Generators for Information Technology
White Paper 93

Essential Elements of Data Center Facility Operations
White Paper 196

Harmonic Currents in the Data Center: A Case Study
White Paper 38

Browse all white papers
whitepapers.apc.com

Browse all TradeOff Tools™
tools.apc.com

Contact us

For feedback and comments about the content of this white paper:

Data Center Science Center
dcsc@schneider-electric.com

If you are a customer and have questions specific to your data center project:

Contact your Schneider Electric representative at
www.apc.com/support/contact/index.cfm