PQ TechWatch

A product of the EPRI Power Quality Knowledge program

ELECTRIC POWER RESEARCH INSTITUTE

Power Quality in Internet Data Centers

November 2007

Contents

Brian Fortenbery, EPRI

SUMMARY

The building of Internet data centers (IDCs) is a growing industry that is pushing the limits of electric power and reliability requirements. As utilities must decide whether it is worth the cost to build new infrastructure to keep up with the present demand, facility operators are looking at power distribution designs that will improve efficiency and allow them to continue to expand their operations.

To meet customer expectations of "five nines" — or 99.999% — availability, IDC designers must improve power quality, reliability, and efficiency. In this quest, redundancy in the system becomes absolutely necessary, but also important are power quality issues such as mitigation of voltage fluctuations and harmonics and good techniques for grounding.

This *PQ TechWatch* provides information about the types of data centers being built and their design, along with new standards and certification processes that are being developed. Detailed descriptions are provided of power quality considerations and possible solutions. Grounding is given its own section, where electrical standards for safe and effective grounding are discussed and examples provided.

About the EPRI Power Quality Knowledge Program

The EPRI Power Quality Knowledge Program provides a wealth of resources in well-designed, readable, and accessible formats. Paramount among these resources are documents covering a wide range of PQ topics, written not only for use by busy PQ professionals, but also to be shared with important end-use customers and internal utility managers. The program's website, www.mypq.net, is the most comprehensive electronic PQ resource available, providing 24-7 access to proven expertise via the PQ Hotline, hundreds of PQ case studies, over 200 PQ technical documents, PQ standards references, indexes, conference presentations, and a wealth of other resources.

For more information, please visit www.mypq.net.

BACKGROUND

Internet data centers (IDCs) are buildings that house servers and all the associated Internet equipment for Web hosting and data-storage functions. IDC companies provide space, electric power, security, fire suppression, and connectivity for their customers. The connectivity requirement usually prompts IDC companies to locate their buildings near fiber backbones. The electric power requirement, which can be remarkably large for the building size, has traditionally been treated almost as an afterthought. For electric utilities, these power requirements can pose extremely serious problems. For example, when an IDC company approaches a utility and requests 100 MW in an urban downtown area, the request usually can not be honored without a significant investment in infrastructure.

Utilities are generally skeptical about investing in this infrastructure — and rightfully so. The IDCs represent a young industry without a long track record, which leaves utilities unsure about how much power might actually be used or even how long the IDC may stay in business. The data center industry enjoyed phenomenal growth in its early years. Until the spring of 2001, it seemed that there would be no end to the growth of this market segment. However, a market downturn in the spring resulted in a slowing of this growth rate. Instead of building new facilities as fast as possible, IDC companies operated for a few years in a wait-and-see mode. In 2001, it was estimated that only 45% of the colocation space in the United States was occupied.

Today, the market has rebounded and data centers all over the world are filling up, with new construction proceeding in many areas. In today's market, the power density of servers has increased to the point that the main concern of facility operators is how to remove enough heat to allow installation of more servers. Many data centers have literally reached their limit in terms of cooling capacity.

There are several types of data centers, to choose from depending on the needs and intent of the designers.

Shared-Facility Centers

Shared-facility IDCs are most likely the oldest type of IDC in existence. Indeed, they have their roots in and before the 1980s when large computer rooms inside office buildings or warehouses were the norm. Banks, universities, and the payroll departments of large companies were all likely candidates for owning such a facility. Of course, these early computer rooms were full of equipment that seems antique by today's standards, but the principle behind them was much the same.

By definition, a shared-facility IDC is an IDC where at least part of the building that houses the facility is devoted to other purposes. For instance, one example is a tall office building with a few floors devoted to the IDC and the rest of the building devoted to offices and other spaces. The advantage to this type of facility lies in general cost efficiency. Existing building space can often be converted into an IDC for much less than the cost of building a new facility.

Stand-Alone Centers

Stand-alone IDCs were probably the second type of data center to evolve. By natural extension, the stand-alone IDC eliminates many of the drawbacks to shared facilities by removing the relative unknown factors associated with having general-use space attached to the sensitive data-center space. In a stand-alone building, the data-center company has complete control over security, fire control systems, ventilation systems, and power systems, which allows for a greater degree of security and system stability.

Building a data center is a major undertaking, because the standards in materials and construction are higher out of necessity to provide the most stable, secure environment possible for clients.

This type of facility is somewhat newer than the shared data center because only recently has raised-floor space (the area inside the data center that holds computer and telecommunications equipment) been in such high demand. Obviously, the greatest drawback to this type of facility is cost. Building a data center is a major undertaking, because the standards in materials and construction are higher out of necessity to provide the most stable, secure environment possible for clients.

Modular Centers

Modular IDCs offer some of the advantages of stand-alone facilities while not having the disadvantage of being tied to a single location. In fact, such facilities do not actually have buildings at all. Instead, computer and telecommunications hardware is housed inside a trailer or other easily transportable housing with all the requisite environmental controls.

Generally, facilities of this type are not designed for long-term use, instead being designed to provide additional raised-floor space in areas where a sudden and temporary increase in demand is expected. When they are needed, the modular facilities are transported to the site where their use is required, and they are supplied with fiber optic and power connections.

Pre-Built Centers

Along the same thread as modular IDCs are pre-built IDCs. Simply put, a company builds a data center from the ground up and then rents, leases, or sells the facility to another company that wishes to run an IDC in it. The biggest advantage of this type of facility is that little or no lead-time is required in acquiring one and the initial outlay of capital is much less if the facility is not being purchased on a permanent basis.

In general, everything the IDC needs to run except for the telecommunications and computer gear is included in the pre-built facility. These features range from generation equipment, which often has a high lead-time of its own, to heavy-duty construction materials that would have to be added to an existing building if it were to be renovated into an IDC.

Mobile Data Centers

Mobile data centers (MDCs) can be applied to a variety of potential applications:

- **Internet Start-Ups.** When a point of presence in a city or remote location is needed, an MDC may be used. Four hours after arriving on-site, information technology and gateway equipment can be installed and operational.
- **Media.** When a big event for broadcast over the Web is needed, an MDC may be used. In addition, when a network operations center for a news-gathering operation is needed “on the scene,” an MDC can be used to solve the problem. Fiber optics are connected or a satellite-dish base station is added and the Internet system is up and running within hours of arrival on-site.
- **Disaster Recovery.** When a place to house Internet equipment is needed after a disaster such as the World Trade Center attack, an MDC helps meet traditional demand but with a unique solution.

Other types of applications include military and healthcare. The MDC can be delivered by heavy air transport anywhere in the 50-Hz world. A 60-Hz design is available for the United States, Japan, and other countries using 60 Hz.

IDC FACILITY WIRING SYSTEMS

A one-line diagram for a typical stand-alone IDC facility is shown in the figure below. To increase reliability and availability, redundancy must be built into the power system.

One-Line Diagram of Data Center Power Distribution

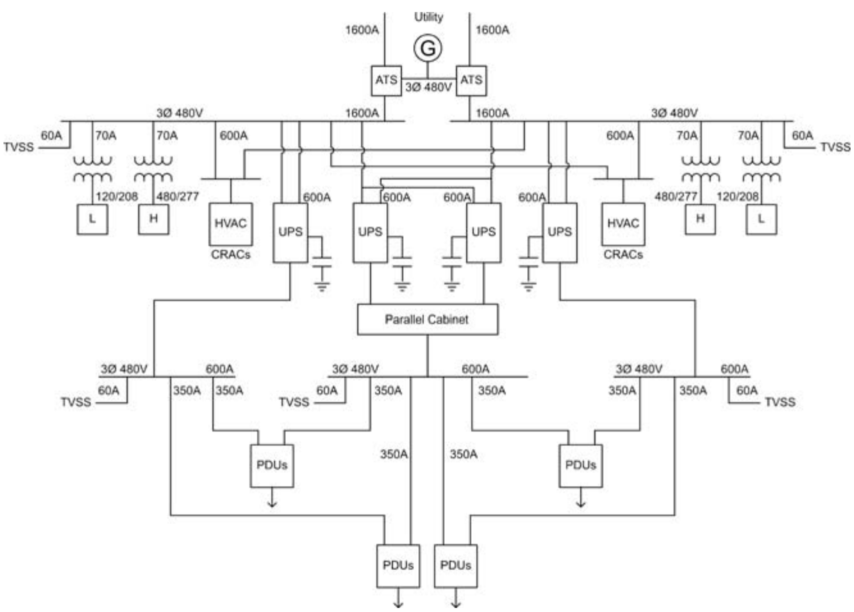

Redundancy is built into power distribution designs to increase reliability.

Automatic Transfer Systems

Two automatic transfer systems (ATSs) are in place at the 480-V level at the plant. One is designed to take the critical loads from the utility to the backup generator upon loss of power. The other is for a second backup generator that may be used in the future. At present the loading in the plant is such that a single generator is sufficient. However, if the occupancy ramps up to greater than half load, then a second generator will be necessary. The ATSs are rated 1600 A with a 65-kA interrupting rating.

Secondary Switchgear

The circuit breakers above and below the ATS are rated the same as the ATS. The ones above are in the main switchgear panel and the feeders to the ATS are three 500 MCM with a #4/0 ground. The ones downstream of the ATSs are located in the main switchboards, SBA and SBB, and control the input to those switchboards. The same size conductors are carried through from the ATSs to these switchboards.

AC Power Systems

From a facility manager's point of view, the main problems associated with power that affect how the IDC is operated are interruptions, power quality, and expandability. Great care and expense was taken to ensure that the interruptions and power quality at this site were addressed. The electrical arrangements to do so are described below. As for expandability, the utility has provided all the capacity requested at this site, so that should not be an issue here.

Switchboards SBA and SBB have various circuit breakers feeding loads. Among those are two 45-kVA transformers for office loads and lighting. One of these has a 120/208V secondary and the other has 480/277V. A third circuit is for HVAC units, one from each switchboard. The next loads are the uninterruptible power supply (UPS) systems. A dedicated feed goes from each switchboard to a UPS. A feeder also goes from each bus to one of two units in a dual UPS used for redundancy. This system is discussed in more detail in the UPS section. Downstream from the UPSs are the critical load distribution panels, which feed the power distribution units (PDUs). Two of these panels are fed by two of the UPSs. A third panel, fed by the third (dual) UPS provides an alternate source of power for the PDUs. This redundancy will ensure very high reliability of power in this plant.

Since the revitalization of the data-center industry, some much-needed formalization of the requirements and criteria for reliability and availability has occurred.

IDEAL CHARACTERISTICS AND REQUIREMENTS

The ideal characteristics and requirements for IDC facilities are based on contributions from industry experts. Each of the following are addressed in data-center design:

- Structure
- Power usage, availability, and reliability
- Ventilation and air-conditioning
- Floor space
- Telecommunications
- Networking
- Security
- Lighting
- Fire suppression
- Siting

For the purposes of this power quality document, the most important of these items will be reliability and availability, which will be explored in more detail. In the early years, criteria for these were market driven — time to market was the most important, with reliability and availability probably secondary. Today, a more formal set of guidelines is available.

Reliability and Availability

Since the revitalization of the data-center industry, some much-needed formalization of the requirements and criteria for the reliability and availability outlined above has occurred. Entities such as the 7X24 Exchange, AFCOM, and the Uptime Institute have played major roles in this endeavor. In particular, the Uptime Institute has taken on a lot of this effort, and although not a standards body, they are well-known and widely respected.

Customer expectations have grown over the years to the point that they now expect availability of "five nines," or 99.999%. This typically requires a substantial investment in hardware and infrastructure, which leads to a need for a common benchmarking standard to help define the level of availability and what is required to achieve it. The Uptime Institute has developed a tiered classification approach for this purpose. This system includes actual measured availability figures ranging from 99.67% to more than 99.99%. Recent surveys show that the Uptime Institute's tier classification system is used by the majority of respondents.

One major concept in the tier levels is that of dual-powered technology, which requires at least two completely independent electrical systems. The last point of electrical redundancy is moved from the UPS system downstream to a point within the computer hardware itself. One description of this type of design is system plus system, or S+S.

A site that can sustain at least one unplanned worst-case infrastructure failure with no critical load impact is considered fault tolerant. A site that is able to perform planned site infrastructure activity without shutting down critical load is considered concurrently maintainable (fault tolerance level may be reduced during concurrent maintenance). It should be noted that a typical data center is composed of at least 20 major mechanical, electrical, fire protection, security, and other systems, each of which has additional subsystems and components. All of these must be concurrently maintainable and/or fault tolerant to meet the requirements. In some cases, sites might be built with S+S electrical concepts but might not incorporate those into the mechanical systems. In such cases, the site may achieve Tier IV electrically, but only Tier II or III mechanically.[1]

The following is a description of these tiers, as revised in 2005. The figures below illustrate these tiers and the typical power-distribution topologies required.

Tier I

Tier I is composed of a single path for power and cooling distribution, without redundant components, providing 99.67% availability.

Description of Uptime Institute Tiers I and II

	Tier I Basic Data Center	Tier II Redundant Components
System Architecture	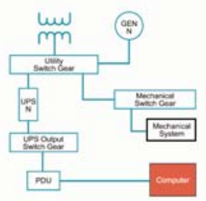	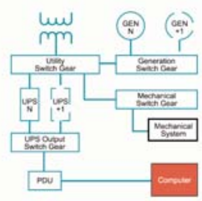
Industry Types	Manufacturing Equipment, Semiconductor Fabrication	Medical, Law Firm, Small Data Room
Downtime Cost/Hour	Less Than $1,000 per Hour	$1,000 – $10,000 per Hour
Annual Downtime	~ 28.8 Hours	~ 22.0 Hours
Availability	99.67%	99.75%

As the cost of downtime increases, systems must become more reliable by adding redundancy and increasing fault tolerance.

Description of Uptime Institute Tiers III and IV

	Tier III Concurrently Maintainable	Tier IV Fault Tolerant
System Architecture	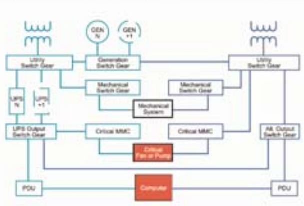	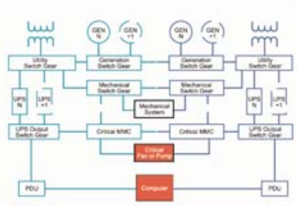
Industry Types	Airline, Retail, Banking, Insurance, Pharmaceutical	Wall St. Trading, Web Hosting, Broadcasting
Downtime Cost/Hour	$10,000 – $50,000 per Hour	Over $50,000 per Hour
Annual Downtime	~ 1.6 Hours	~ 0.8 Hours
Availability	99.98%	99.99%

At Tier IV, Uptime Institute adds specifications for required hardware performance.

Tier II

Tier II is composed of a single path for power and cooling distribution with redundant components, providing 99.74% availability.

Tier III

Tier III is composed of multiple active power and cooling distribution paths (but with only one path active), has redundant components, and is concurrently maintainable, providing 99.98% availability.

Tier IV

Tier IV is composed of multiple active power and cooling distribution paths, has redundant components, and is fault tolerant providing 99.995% availability.

The main addition for Tier IV is that all computer hardware must have dual power inputs. The Uptime Institute has a specification for that also, which includes thirteen points describing the required equipment performance. The requirements include two power sources to the equipment, wherein the loss of either will not affect operation. Return of lost power will not require any intervention. The other source may fail 10 seconds after return of the lost source with no impact on operation. Specifications regarding frequency, phase angle, and phase rotation are included, as well as internal terminations without transfer switches. Neutral-to-ground bonds inside the equipment are not allowed, to prevent circulating ground currents between the two sources. Other requirements include load sharing within 10% of average for multiple power sources, and internal ride-through designs are only allowable for safe orderly shutdown. Annunciation is to be provided for fault indication within 60 seconds, and the manufacturer supplies a written certification of compliance with the specification. The Uptime Institute tests

In data centers, large expenditures are made to ensure that power quality issues do not affect the operation of the business.

products to independently verify and certify compliance with this specification. A listing of certified products is maintained on its website.[2]

Power Quality Considerations

Power quality issues are important in data centers, but in a different way than most types of operations. In data centers, large expenditures are made to ensure that power quality issues do not affect the operation of the business. PQ issues that should be considered in the data center industry include:

- Voltage sags
- Momentary interruptions
- Voltage swells
- Voltage transients (surges)
- Harmonics
- Grounding

Ride-Through Curves for Servers

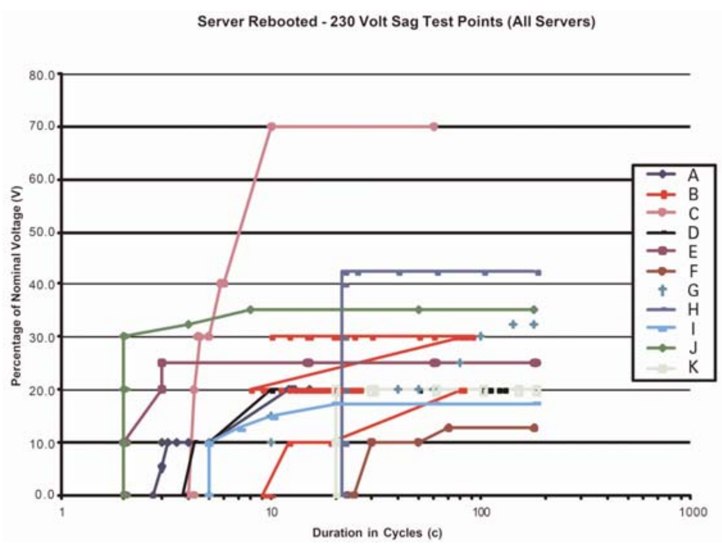

Most servers can ride through sags and interruptions that last from 30 to 300 ms.

Voltage Sags and Momentary Interruptions

By virtue of the nature of the power distribution system in data centers, sags and interruptions are not usually a cause of downtime. In fact, these events have been all but removed from consideration by the upstream power conditioning, using UPS and backup generators. Even a Tier I site will use UPS and generator sets, so sags and interruptions can only cause a problem if there is an equipment failure. This is why the tier definitions are so important-the power quality problem has become an infrastructure design and maintenance problem, so that the probability of downtime is significantly reduced. In fact, the most common cause of downtime in the data center industry is human error, quite often accidental activation of the emergency off switch. Other human errors can occur, which can be minimized through training and, often, installation and maintenance of the redundant systems.

In addition to the above discussion, previous testing by EPRI has shown that most servers can ride through sags and interruptions for durations that vary from 30 to 300 ms. While there may be many sag events longer than 30 ms, few are longer than 300 ms, and so devices that survive 300 ms sags are far less likely to be affected by voltage sags in general. The plot to the left shows test data for a variety of servers tested at 230 V. It should be noted that the curves represent voltage tolerance envelopes, or "ride-through" curves. The horizontal axis is time in cycles, while the vertical axis is the value of the remaining voltage as a percentage of nominal. Each curve shows that if sags occur that fall under the curve, the server may trip offline. If the sags are above the curve, the server will continue to operate without interruption.

Voltage Swells

Many UPSs on the market do not protect loads from voltage swells, but there is not much evidence of damage or equipment misoperation as a result of swells. The IEEE definition states that swells may be up to 1.8 per unit (PU) for as long as 2 minutes. EPRI testing programs for temporary overvoltage up to 3.0 per unit have determined that most end-use equipment can sustain damage under these conditions. The tests performed for temporary overvoltage are outlined below.

Test Parameters for Temporary Overvoltage

Test Number	Simulated Condition	Magnitude	Duration
1	Poor voltage regulation	1.15 PU (138 V)	6 hours
2	During a fault	1.3 PU (156 V)	2 seconds
3	Loss of secondary neutral	1.5 PU (180 V)	4 hours
4	Ferroresonance	2.0 PU (240 V)	1 minute
5	Contact to high voltage circuits	3.0 PU (360 V)	1 second

Swell Rates from EPRI Survey

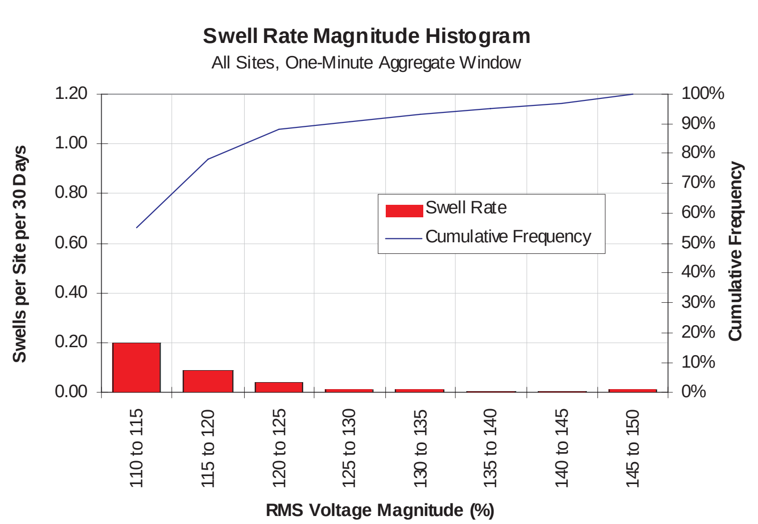

Voltage swells are infrequent and rarely cause damage.

However, based on EPRI survey data, these are not very common events. In fact, such events can be estimated to occur as seldom as once every six years in the United States. The chart on the lower left shows some of the data from the survey.

Voltage Transients (Surges)

Lightning is a well-known phenomenon that can cause equipment damage. Surge protection devices are widely used in commercial, industrial, and residential settings, and the data center is no exception. Most UPS manufacturers will offer surge protection for the critical loads, and in addition, the servers themselves will typically be supplied with surge protection.

When studying the power supply, the transient damage can affect the following:

- Diode rectifier bridge
- Filter capacitor
- Fuses
- Switching transistor

The figure below shows the points where surges are likely to enter the equipment.

Surge Entry Points for Equipment

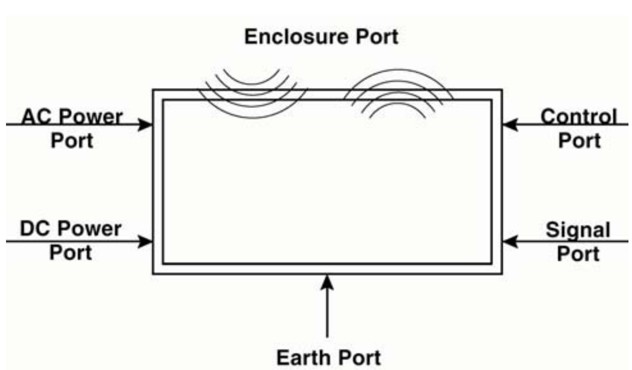

Surges can enter electrical equipment and numerous connection points.

The three common elements used in surge protection are

- Gas tubes
- Silicon avalanche diodes
- Metal oxide varistors

Various vendors supply devices based on these technologies, which can be applied inside the appliance, or at the plug outside the appliance. Surge arrestors can also be applied at the panel level, or the switchgear at the service entrance to the building. Utilities sometimes apply lightning arresters outside the building, as well.

The figure below shows an example in which the primary surge protection device is installed at the switchgear on the incoming bus from the transformers.

Secondary surge protection devices may also be installed on the PDU transformer primary and secondary windings. Protecting both primary and secondary is the ideal arrangement for protection at the transformer level.

Surge Protectors in PDU

An example of a surge protector installation for a power distribution unit.

Surge Protector on Switchboard

Caption: Surge protection is often installed along with other protected equipment in the same panel.

Also, as seen in the figure above, with this type of device the lead length issue cannot easily be addressed. Lead length should be kept as short as possible for optimum performance, but it is unclear how this installation could be improved. This installation method may affect the peak clamping performance by about 100 to 800 volts depending on the type of transient.

Another issue may be coordination of the protectors and the main switchgear surge protector. Depending upon the location of the impinging transient, the primary protector may allow a majority of the surge into the facility, which is undesirable.

Data lines must also be protected from surges, and plenty of devices are available to provide that protection.

Data lines must also be protected, and plenty of devices are available to provide that protection. Lightning transients tend to cause currents to flow on ground conductors, which can induce voltages on data lines. These voltages can easily exceed the withstand levels for silicon devices in data ports. Exposure levels in data centers, however, will be lower because the lines never leave the building. In addition, most computing devices will have protection built in.

Harmonics

Many servers today are available with power factor-corrected (PFC) power supplies. IEC 61000-3-2 addresses specific harmonic limits for end-use devices, which often will require the PFC circuitry to be included. In addition, the UPSs serving those critical loads are often power factor corrected. Because of this, many people believe that data centers do not suffer from harmonics like industrial or typical commercial settings. It may be true that the levels are better because of the PFC, but that does not necessarily preclude the possibility of harmonics in the building. Some offending loads in data centers are as follows:

It may be true that harmonics levels are better because of the use of power factor correction, but that does not necessarily preclude the possibility of harmonics in the building.

- UPSs
- Auxiliary electronic loads (non-PFC)
- Rectifiers
- Fluorescent lights
- Transformer inrush (short duration)
- Drives (HVAC)
- Soft starters (short duration)
- Battery chargers

Some or all of these may exist and play a role in the harmonic content in the data-center wiring. In particular, when transferring to and from generator, due to the higher source impedance, harmonic currents can cause voltage distortion at the main bus, affecting operation of other connected equipment. For example, motors and capacitors can be affected by this voltage. Harmonic currents can cause heating in transformers and cables, including neutrals. Other issues may include misoperation of fuses and breakers. When power factor correction capacitors are added to the system, resonant conditions can occur. This happens when the magnitude of the capacitive reactance is close to the magnitude of the inductive reactance of the system, usually dominated by the source impedance.

Specific problems include:

- Heating effects in phase conductors as well as neutrals
- Heating in three-phase dry-type transformers
- Harmonic voltage-related heating in other equipment, such as motors
- Harmonic voltage stress on system capacitors and equipment capacitors
- Resonance with power factor correction capacitors
- Voltage distortion at the point of common coupling (PCC)

Solutions are discussed in the table on the following page.

Filters can be placed at various locations, including at bus or panel level or at loads filters at bus or panel level can be arranged as

- series-connected neutral current filters,
- parallel-connected zigzag filters, or
- parallel-connected active power filters.

Harmonic Solutions

Overloaded Phase Conductors	Overloaded Neutrals	Transformer Heating	Motor Heating	Capacitor Stress	Resonance Issues
Larger Conductors Filters at Various Locations	Larger Neutral Wire Neutral for Each Phase Double Number of Neutral Conductors Third Harmonic Filters at Loads Third Harmonic Filter in Neutral Zigzag Filter Close to Load Active Filter	Derating Transformer (Larger Transformer) Applying K-Rated Transformer	Derating Motor Filters at Various Locations	Proper Sizing of Capacitors Filters at Various Locations	Filters at Various Locations De-Tuning Capacitor Banks (Creating Filters) Judicious Choice of Capacitor Size and Location

Built-in filters at loads can be series inductors or can boost converter power factor correction. Branch circuit or cord connection filters can be applied as

- parallel-connected resonant filters,
- series-connected resonant filters,
- neutral current filters,
- zigzag filters,
- active filters,
- choke upstream of motor drives, or
- passive three-phase filters upstream of motor drives.

Generic Power System Grounding

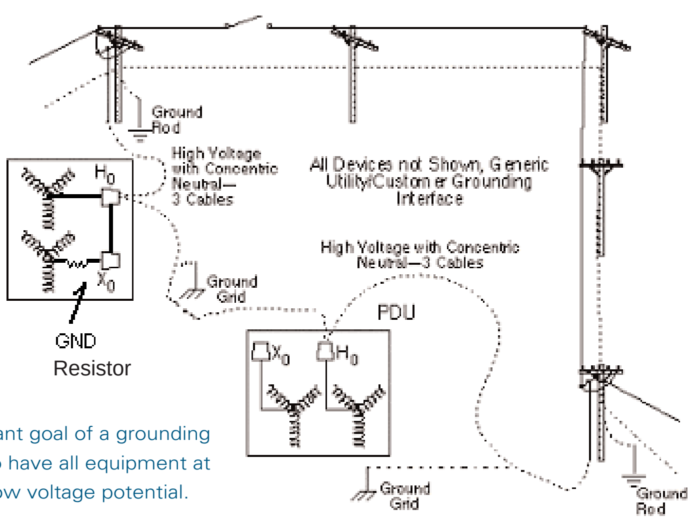

One important goal of a grounding system is to have all equipment at the same, low voltage potential.

Importance of Grounding

Grounding is required for safety foremost, and in the United States, sites should meet the National Electric Code (NEC) as a minimum to ensure proper safety. The secondary reason for grounding is for performance (noise control). Most equipment manufacturers will provide details on grounding of their equipment. The practice of installing separate grounding systems should be avoided and one should always look for "special grounds" that the manufacturer may want to install or has installed as they may violate safety codes. The NEC practices allow for safe and effective signal grounding; however, single point and star grounding may not be suitable for high frequencies. Article 645 "Information Technology Equipment" of the NEC provides guidance for the minimum wiring and grounding of such facilities. The practices described in IEEE Standard 1100-1999 provide guidance for noise control without violating the NEC. IEC 60364 also addresses earthing connections.

Power System Grounding

Grounding of the power system ties all equipment together and ensures that the equipment is at the same potential as the earth. This is accomplished by bonding the equipment to the grounding electrode system so that everything may be referenced to earth potential. This is critical during fault conditions to ensure that overcurrent protective devices operate in a timely manner to reduce or prevent equipment damage.

The figure to the left illustrates an example of a power grounding system including the PDU. As can be seen, the entire distribution system on the primary side of the PDU transformers provides a high impedance return path at any point during a line-to-ground fault condition. This limits fault current to a level that allows continuous operation. The grounding issues on the secondary side of these transformers are addressed in a later section. European practices use various connections, including the TT system, the TN system, and the TI system.

The TT system is designed for personnel protection, with the exposed conductive parts earthed and residual current devices used. The neutral is not bonded to the earth downstream of the transformer. This allows the first fault to occur without interrupting the service.

The TN system is more closely related to the U.S. system, with interconnection and earthing of exposed conductive devices and the neutral required. This necessitates protection through circuit-interrupting devices such as fuses and circuit breakers. This allows the first fault to interrupt the service.

The TI system is similar to the TT system, with indication of the first fault through the use of an insulation monitoring device. The second fault is interrupted using fuses and breakers.

Some practices in the United States are described below.

To keep information technology equipment noise-free, the ground is often removed; however, a better practice is the use of computer signal reference grids.

High Impedance Grounding

High impedance grounding usually is used in circuits rated at 1000 V or higher and is covered by Article 250-153 “Impedance Grounded Neutral Systems” of the NEC. However, as may be seen in the figure on top right, high impedance grounding is located within the facility for 480-V operations and is supported by NEC Article 250-5 “Alternating-Current Circuits and Systems To Be Grounded,” exception number 5.

High Impedance Grounding Resistor

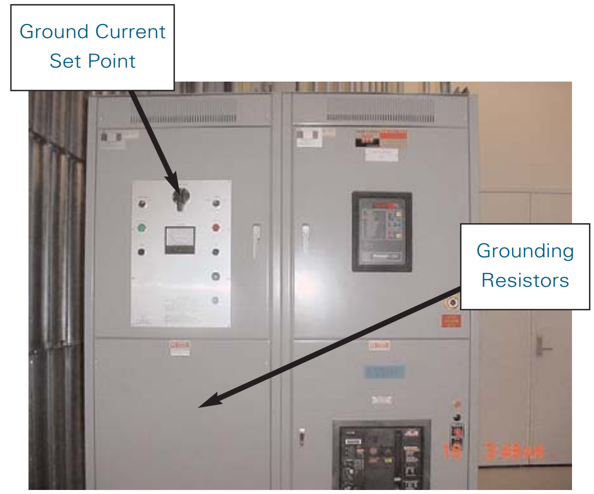

High-impedance ground resistors for a rotary UPS system.

The impedance ground-current set points were found to be set at reasonable levels, less than a 5 V drop across the grounding resistor.

Computer Signal Reference Ground

In an attempt to keep the information technology equipment noise-free, the ground is often removed in attempts to provide isolation and a separate grounding system is provided. However, a better practice, which does not violate the NEC, is the use of computer signal reference grids, which in turn are bonded to the grounding electrode system.

One such system is described in the IEEE Standard 1100-1999 and is illustrated on the following page.

Raised-Floor Grounding System

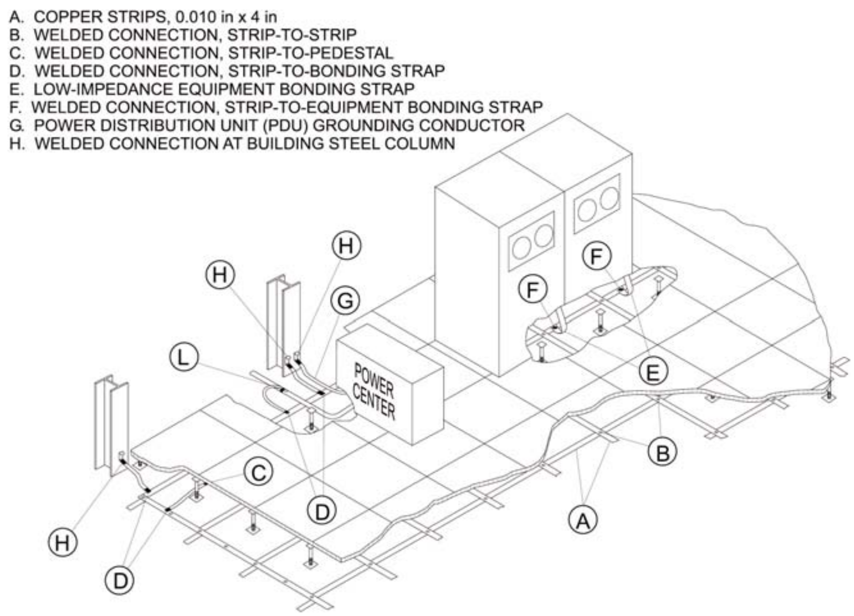

Caption: IEEE provided guidance for the grounding of raised-floor grounding systems.

Source: IEEE Standard 1100

However, the methods outlined in the IEEE 1100 standard or FIPS Publication 94 may not be practical or cost-effective in all cases. An alternative method is bonding of pedestal supports using a crisscrossing serpentine arrangement of cable, instead of flat ribbons, connected at every other pedestal, which results in a grid with 1.2-m squares as opposed to less than one-third of a meter. At each crisscrossing, the grounding conductors must be connected in grid fashion. The figure at top right illustrates the proper bonding method of the crisscrossing conductors.

The signal reference grid follows this practice of bonding to establish a low impedance signal reference grid. However, these practices only cover frequencies up to about 3 MHz and provide a reduction in noise level by a factor of 10.

Ground Grid Bonds Under Floor

A compression grounding bond.

As technology is improved and density increased, and with the onslaught of lower logic voltage levels of 3.3 V and less, noise-to-signal ratios may become a problem, and the signal ground reference grids may have to be improved. The figure below illustrates a central ground plate commonly used for such improved systems. The central ground plate is designed for high frequencies. Not only is the serpentine crisscrossing and bonding of the pedestals used, but 4-inch-wide copper straps are also necessary to form ground grids of areas smaller than 0.6 m square to ensure reduction of signal noise in some information technology equipment rooms.

Central Ground Plate Example

A central ground plate can serve as the common bonding point for many grounds.

Raceways should be bonded by direct connections, which improve the ground return loops of stray EMI currents.

However, the improvements may only increase the frequency range to 10 MHz. Problems associated with frequencies above 30 MHz may have to be handled differently (for example, use of a shielded enclosure).

Equipment Rack and Metallic Raceway Bonding

Raceways should be bonded by direct connections. These connections improve the ground return loops of stray EMI currents. Also, the mutual inductance can provide up to 75% reduction of the initial loop current by cancellation. This reduces the susceptibility and emissions problems related to common mode voltage in large sites. Any discontinuities will reduce the benefit of using metallic raceways for common mode noise reduction.

Wiring Practices

Connections in panels should be tight, and panels should be inspected once a year using thermal imaging as a minimum preventive maintenance action.

Bonding

In general, proper bonding in accordance with NEC article 250 part G should be carried out throughout the facility. The figure below illustrates one such required bond for flexible conduit joints, thereby establishing acontinuous and permanent ground requirement as stipulated by the NEC.

Telecom Grounding

A typical telecom system design uses a three ground rod system bonded to the main signal grounding grid and an additional grounding electrode bonded at the service entrance. The NEC allows for this additional electrode as it becomes part of the overall grounding electrode system. This will prevent the telecom system from surge-induced overvoltages by allowing it to rise to the same potential as the building grounding system during a fault or surge current condition.

Ground Currents

A reasonable amount of ground current may be expected to be flowing at ground plates throughout the facility. However, NEC Article 250-21 "Objectionable Current Over the Grounding Connections" states that measures should be taken to reduce the currents to levels satisfactory to the authority having jurisdiction. Ground currents should not be present on the grounding electrode conductors or the counterpoise grounding electrodes. These currents are associated with induction and are contained within the feeder conduits and circuit loops between the interconnected equipment and their associated grounds. A typical finding for a data center is shown below.

Grounding of Conduit Flex Joint and Transformer

Flexible conduit often requires special accommodations for maintaining good ground continuity.

Grounding Plate in Data Center

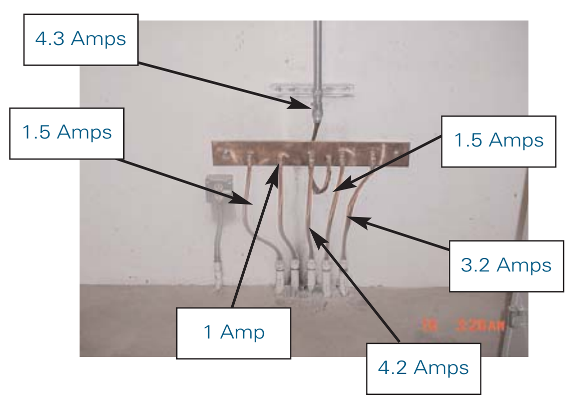

Grounding plates provide a common bonding point for different grounds

Parallel paths created by the insulated ground conductor and the conduit grounding provide paths for circulating ground current. Unbalanced spacing between phase conductors and the equipment-grounding conductor along the feeder path allows current to be induced in the grounded conductor, as seen in the top figure below. When several parallel feeders are involved and the conduits and conductors are bonded at both ends, as is the case at the service entrance and the bus, sizable currents may be introduced. This does not necessarily constitute a problem, but if problems do occur, such as noise from ground loops, the situation should be addressed. Below, at bottom, is an example of induced currents in a typical communication/computer equipment facility.

If some sizable ground currents exist on the conductors at a site, it would be prudent to document and explain them. Specifically, the system documentation should describe the cause, effect, and appropriate levels of these leakage currents. Also updating of all site drawings to show the as-built condition of the grounding systems is suggested. This would be helpful later if ground loops created enough noise to require some mitigation.

The PDU Grounding System "Isolation Transformer"

The PDUs are basically step-down transformers and provide a localized source for each colocation remote power panel. The figures below show the details of a PDU, including the grounding arrangement.

Arrangement of Conductors

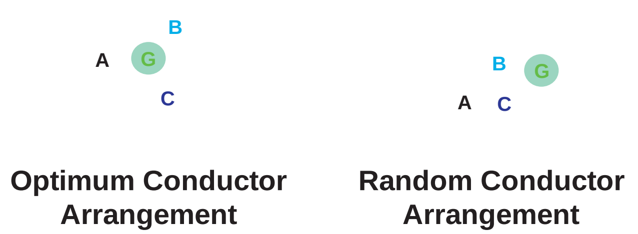

Induced ground currents come from asymmetry distribution of conductors within a conduit. This can be exacerbated by long cable runs.

Induced Currents in an Information Technology Facility

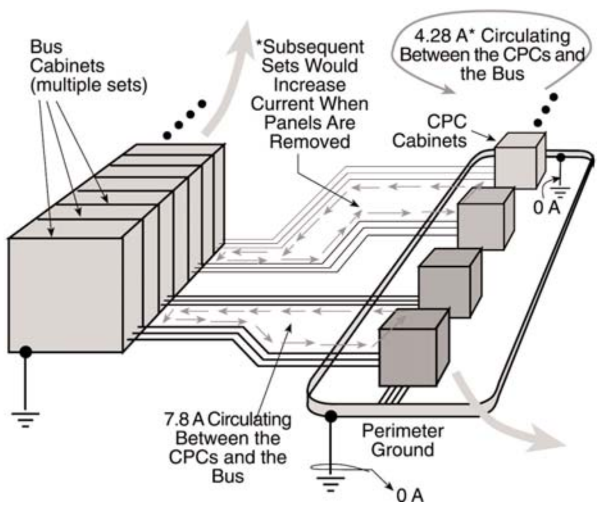

Sizable ground currents should be documented in case they require mitigation later.

PDU Front View

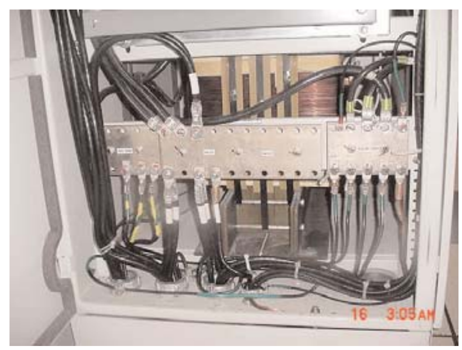

The interior view of a power distribution unit.

PDU Grounding Detail

A robust PDU grounding lug.

Article 250-26 "Grounding Separately Derived Alternating-Current Systems" in the NEC covers the requirements for the PDU. Article 250 is met satisfactorily when the PDU's neutral ground bond is grounded to the signal reference grid, which in turn is bonded to the electrode grounding system.

DC-Powered Data Centers

There has been a significant amount of research in recent years on the possibility of powering data centers with direct current instead of alternating current. Several publications have come from Asia, Europe, and the United States on this issue. Some data centers in Japan are operating with power distribution throughout the building at 300 VDC. In Europe, Cenelec and other groups have explored the possibilities for several years. In the United States, a project was undertaken by Lawrence Berkeley National Laboratory to demonstrate the implementation of 380-VDC power distribution in a data center. This involved supplying the servers directly with 380 VDC, instead of an AC supply.

The reason for this interest is that the system may achieve higher reliability, as evidenced by the telecom industry, as well as the possibility of higher operating efficiency. The efficiency gains come about as a result of the removal of the UPS and the transformer that usually accompanies it, as well as a higher efficiency in the server power supply due to the removal of the rectifier and power factor-corrected circuit at the input. With the DC concept, the only components necessary are a main rectifier and some kind of energy storage (flywheel, battery, etc.). The transformer and the rectifier section in the front end of the UPS, and their respective losses, are removed, as shown in the figures at right.

continued on following page

AC Power Distribution

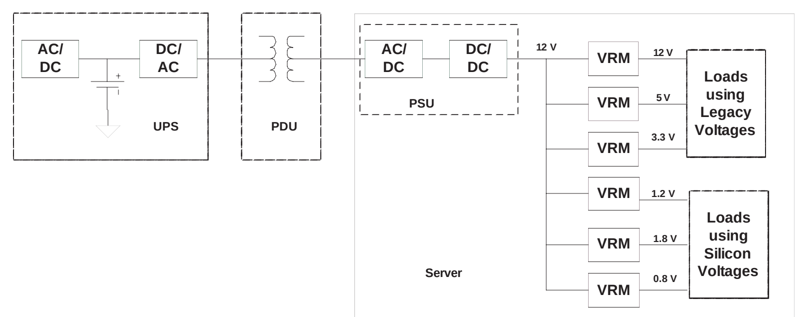

DC Power Distribution

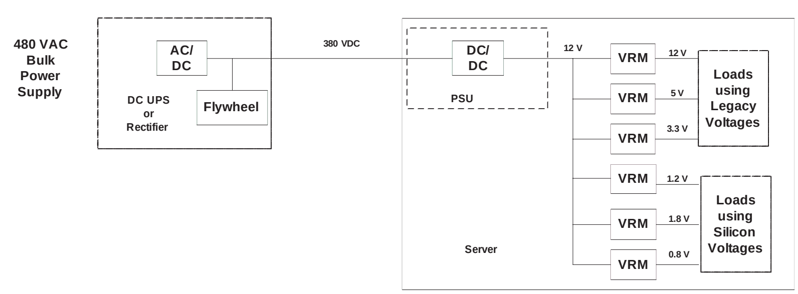

Losses from the UPS transformer and the power supply unit (PSU) rectifier are eliminated in a DC system.

VRM = voltage regulator module

DC-Powered Data Centers (*continued*)

The system was compared with two AC systems using two different UPS manufacturers, A and B. It should be noted that the DC system output load is less due to the improvement in the power supply efficiency. The basic results from this experiment are shown in the table to the right.

Test Results from DC Data Center Demonstration

	UPS	Transformer	Power supply	Total Efficiency	
Efficiency Measured with UPS A	89.6%	98.0%	90.0%	79.0%	
Efficiency Measured with UPS B	89.9%	98.0%	90.0%	79.3%	
Efficiency Measured with DC System	94.2%	100.0%	92.0%	86.7%	
Typical Efficiency	85.0%	98.0%	73.0%	60.8%	
	Output Load (kWh)		**Input Load (kWh)**		
Energy Measured with UPS A	23.3		26.0		
Energy Measured with UPS B	23.3		25.9		
Energy Measured with DC System	22.7		24.1		
Energy Improvement over UPS A					7.3%
Energy Improvement over UPS B					7.0%

Hardware Used in DC Data Center Demonstration

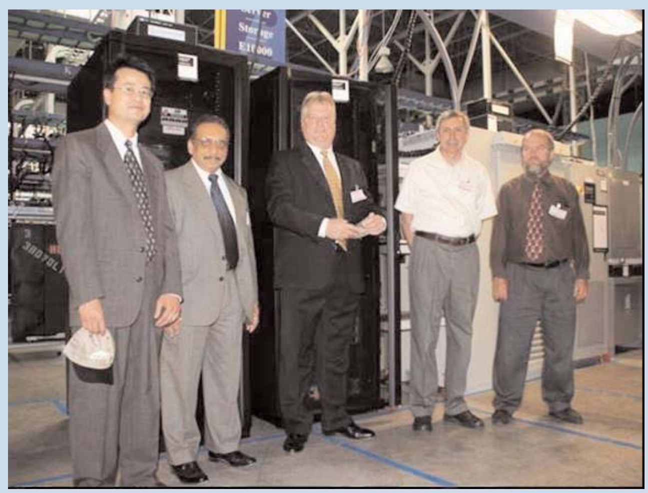

Representatives from Ecos Consulting, the California Energy Commission, EPRI, and the Lawrence Berkeley National Laboratory are shown in front of the DC power system at Sun Microsystems.

FURTHER READING

European practices can be easily found in the following sources:

Schneider Electric, *Cahier Techniques*, available from www.schneider-electric.com.

Merlin Gerin, *Electrical Installation Guide According to IEC 60364*, available from www.merlin-gerin.com.

British Standards Institution, BS 7671:2001, "Requirements for Electrical Installations, IEE Wiring Regulations," 16th Edition.

J. F. Whitfield, *The Electrician's Guide to the 16th Edition of the IEE Wiring Regulations, BS7671 and Part P of the Building Regulations* (EPA Press, 2005).

And several Leonardo Power Quality Institute documents (available from www.lpqi.org):

"A Systems Approach to Earthing"

"Fundamentals of Electromagnetic Compatibility"

"Earthing Systems—Fundamentals of Calculation and Design"

"Earthing Systems—Basic Constructional Aspects"

NOTES

1. The Uptime Institute, *Industry Standard Tier Classifications Define Site Infrastructure Performance*, White Paper (Santa Fe: The Uptime Institute, 2000).

2. The Uptime Institute, "Fault Tolerant Power Compliance Specification Version 2.0" (2002), available from www.upsite.com/TUIpages/tui_certification.html.