HANDBOOK OF LOCAL AREA NETWORKS

SECOND EDITION

IX-1 Designing Power Distribution Systems for Fault-Tolerant Networks

DAVID FENCL

AUERBACH PUBLICATIONS
Boston and New York

HANDBOOK OF LOCAL AREA NETWORKS

SECOND EDITION

This article is reprinted from
Auerbach Publications'
Handbook of Local Area Networks
Second Edition
John Slone, Consulting Editor.

IX-1
Designing Power Distribution Systems for Fault-Tolerant Networks

DAVID FENCL

What does it mean to build a fault-tolerant network? Fault tolerance is usually taken to mean redundancy, and it is true that redundancy is an important part of any fault tolerant system plan. But, according to an old parable, "a wise man builds his house upon rock, a foolish man builds his house upon sand." The network administrator who looks only at redundancy as a measure of fault tolerance may be like someone who builds two houses without considering whether the foundation is like rock or like sand. The building power distribution and grounding system is a significant part of the foundation for the network. Its importance is often overlooked, but a marginal power distribution network is often the major cause of unexplained system crashes.

Those who have made a study of design techniques that enhance reliability may consider attending to the building power distribution network, and enhancing it with power conditioning devices, as a fault avoidance technique. Installation of back-up power systems (i.e., uninterruptible power supplies, or UPSs) might be considered a fault tolerance technique in that it introduces an element of redundancy in the power path. Whatever the label, both solutions make positive contributions to overall system reliability.

Deploying UPSs is certainly a good start toward improving a network's fault tolerance. Without them, if power goes away, even for a fraction of a second, many devices on the network will lock up or reset. The effect can range from slight inconvenience to business-stopping catastrophe. The severity depends on the role of the network in the business process and the characteristics of the devices effected. Today, most microcomputer networks use UPSs to protect the data on the file server.

Reliability expectations increase, however, as LANs become necessary to the minute-by-minute activities of a business. More and more LAN administrators, system planners, and support professionals are learning by experience that a UPS by itself does not always prevent system crashing abends caused by spurious hardware interrupts. These intermittent faults occur often because not all UPS architectures isolate the server from transient anomolies in the electrical environment.

Environmental concerns are not new to computing. Every mainframe and minicomputer vendor has requirements for the system's physical environment, in particular, cool ambient temperatures, dust-free air, and system power that is isolated from other building loads and which is supplied using distribution and grounding schemes that isolate computer circuitry from lightning-induced surges and high-frequency electrical interference.

For the most part, the independent network integrators that install and support most LANs have not had as much experience with the environmental disciplines as have vendors of traditional large computing systems. For this reason, the link between environmental factors and LAN device performance is often overlooked. This chapter is intended as a reference for good environmental practices translated for high-performance, high-reliability distributed computing systems.

Managing fault tolerance in a distributed systems environment is actually harder than in a conventional "big iron" environment. Networks are at a disadvantage because they are more complex, the electronics operates in much harsher environments, and staffing and staff training are less thorough. The situation is summarized in Exhibit IX-1-1.

If the goal of building a fault-tolerant system is to achieve higher system reliability through lower failure rates, all the components of the system should be designed with the goal of reducing the overall number of failures. Such a comprehensive approach includes concerns for the distributed operating environments, system administration, network management tools, device redundancy, disaster recovery, and contingency planning.

LAN SYSTEMS RELIABILITY NEEDS ASSESSMENTS

How much is system reliability worth? In one survey, 35% of respondents indicated that downtime costs exceeded $10,000 per hour, 42% said downtime cost their operations as much as $1,000 per hour, and 23% did not have any idea what downtime cost their organizations.

	Mainframe/Midrange	LAN
Operator Expertise	Full-time IS professional	Part-time system administrator
Support	Several professionals tending one installation	One professional tending many locations
System Complexity	Single location, single vendor	Multiple location, multiple vendor
Environment	Isolated conditioned room	Placed anywhere

Exhibit IX-1-1. Support Comparison of Mainframe/Midrange Systems and LANs

A standard model would be useful in assessing the potential value of reliability enhancement techniques. Such a model would consider the cost of downtime as well as the risk of downtime. The following is a simple equation for modeling the value of reliability:

Reliability Value = Cost of Downtime × System MTBF × Site Risk Probability

where

Cost of Downtime = (System Time Value × Mean Time to Repair) + Cost to Repair

MTBF = Mean Time between Failures

Understanding the Time Value of a System

One of the first steps in modeling the value of reliability is to assess the function of the planned network or network segment The network manager must understand the time value of the work group served by the system and whether the system is part of the critical path of a larger process. If the latter is the case, the time value of the system segment has to include the value of the total business effort served by the system's users.

Determing the Direct Costs of Downtime

The Time Value of Users. One method of determining the time value of users is to sum the burdened payroll costs of all system users. This could be refined by multiplying each user's payroll rate times an estimate of the user's system utilization rate. For example, users involved in continuous data entry activities would have a 100% system utilization rate—all the time they are working, they are working on the system. A user who uses the system for only one hour each standard work day would have a system utilization rate of only 12.5% ($^1/_8$). This method discounts the time value of users who could shift their system tasks into other time periods, presumably after the system has recovered from whatever fault has caused the downtime. This payroll cost valuation method might make sense for systems that are important but are not in the critical path. For systems in the critical path, a valuation method based on work group output makes more sense.

Time Valuation for Critical Path Systems. A critical path system is one in which a system fault affects not only the immediate users of the system but affects total output of the business effort dependent upon direct systems users.

Online Transaction Process (OLTP) Systems. OLTP systems include retail store systems, order entry for inside sales departments, automated banking systems and other systems where customers (i.e., clients) wait while system users process their transaction in real time.

Real-Time Decision Support Systems. Examples of real-time decision support systems include any financial trading desk where buy/sell decisions must be made in an environment where prices change minute by minute. Not making a decision or making an uninformed decision can be extremely costly. In some complex manufacturing environments, a continuous monitoring system is in operation that provides real-time information which, if unavailable, could result in costly scrapping of material in process. In an operating room or intensive-care unit, patient monitoring systems provide real time information that could cost a patient's life if it were unavailable to the attending physician.

In each of these critical path examples, the function of the network or network segment is of far greater value than the payroll cost of those directly using the system. The output valuation method only considers the value of the output of the user or work group versus the payroll cost of the user or work group. In the critical path example shown in Exhibit IX-1-2, a system development team is working on a branch office automation system that is expected to save $3.5 million in annual operating expenses for the company. Delays in system development will delay the realization of these anticipated cost savings, thus adding significantly to the direct cost of system downtime.

Indirect Costs of Downtime. Quantification of indirect costs is difficult if not impossible. Even so, lost future business and costs of work flow interruption need to be considered.

Image Loss Results in Lost Business. Customer frustration can be triggered by something as minor as waiting a little too long for processing of an order or something as significant as a supplier's losing track of a time sensitive shipment. One business model assumes that 0.1% of customers affected by a system failure will switch their business to a competitor. If those customers who switch to a competitor represent average customers, the image loss would be 0.1% of annual sales volume. If the frustrated customer is a large account, or if the company's business is in a highly competitive industry, lost future business could be much larger than 0.1% of sales.

Creative Productivity Loss from the Interruption of Work Flow. The least likely scenario is the infinite value of the idea that got away. More likely is the amplification of interruption time that occurs among creative workers who break concentration and lose work flow momentum due to system faults that interrupt the creative process. For these workers, "re-contexting" time can extend the actual time cost of the system fault.

Determining Repair Costs in Time and Money

The direct cost of system downtime is the time value of the system multiplied by the time it takes to repair the system and recover from the fault (mean time to repair, or MTTR). Add to this, the actual costs to repair or recover from the fault.

Mean Time to Repair (MTTR). Mean time to repair is dependent on fault detection, service response time, and system recovery time.

RELIABILITY VALUATION MODEL

Components of Reliability Costs	Time	Important System (Product Engineering Group)		Critical Path System (MIS System Development Team)	
Cost of Down Time					
System Time Value Computations:					
Direct User Payroll (#users)		5		30	
Average Payroll Cost/Hr		$73		$65	
Average system use rate (use hrs/work hrs)		50%		95%	
—Subtotal (Pay $cost/minute)		$3		$31	
Direct Value of System Output					
Profit Contribution Forecast From Project Output ($/year) Reductions...		Profit on New Product Sales		From Process Cost	
		$500,000		$3,500,000	
—Subtotal (Cost of Delay: $/minute)		$4		$28	
Tangible System Time Value ($/minute)		$7		$59	
Mean Time To Repair	**(minutes)**	**Cost per Event**			
Scenario 1					
Soft Error/Intermittent System Crash					
Detection Time for SysAdmin	5	$35		$296	
System Reboot, no data loss	3	$21		$177	
Total Cost: "Soft" Fault/Quick Recovery			$56		$473
Scenario 2					
Hard Failure (Resulting in Lost Disk Files)					
Hardware Repair & Sys Recovery (4 hr)	240	$1692		$14185	
Data Restore (from Tape) (1 hr)	60	$423		$3546	
Data Recovery (manual) (i.e. time to enter data into system that accumulated during system down period — order entry, accounting, other on-line transaction processing)	0	N/A for these examples		N/A for these examples	
Cost To Repair (Parts and Labor)		720		720	
Total Cost: Hard Failure			$2,835		$18,451
Est. of Annual Costs based on Standard MTBF/MTTO & MTTR					
Assume system MTBF/MTTO yields 1 Hard Failure per Year					
Assume system MTBF/MTTO yields 1 "Soft" incident per Year			$2,892		$18,924
Site Risk Probability (1 = Average Site) *(subjective site weighting & factor interaction)*		4.32		0.84	
Out Of Building Risk Factors					
Utility infrastructure & Stability (Rural, Heavy Industry, High Growth = >1		1.2		1	
Geographic Characteristics of Lightning incidence *(Florida and South East = >1)*		5		0.7	
Within Building Risk Factors					
Quality/Integrity of Building Electrical Distribution *(Installed right and well maintained = <1, poor = > 1)*		0.8		1	
Electrical Active Building (Large Loads on/off) *(quiet building = <1, Active Build = >1)*		0.9		1.2	
Adjusted Annual Reliability Costs			$12,493		$15,896

Exhibit IX-1-2. Network System Valuation Calculations

Detection time can differ according to the user. A direct user may know immediately when the system is down. However, the person with the responsibility for system uptime and the authority to initiate any necessary recovery or repair procedures may not be aware of the fault until users report it.

Sophisticated network management systems can predict faults, making possible proactive response and preventing downtime. Some network management platforms have device specific management applications for remote monitoring and diagnosis of network devices. This facilitates rerouting around the fault or prediagnosing the fault in order to better prepare the field technician for a quick hardware repair.

Service response time is typically controlled by an agreement with a service provider. The agreement guarantees some response time (ranging from a few hours to days), and covers all parts and labor under a fixed pricing schedule. Some agreements include options for hot spares, which guarantee immediate availability of any needed replacement parts.

System recovery and data restoration time can be a brief as the time it takes to replace a failed hard drive and restore from backup tape or from a mirrored drive. Costs can include the labor costs and elapsed time costs of reentering the data from hard copy archives. The worst case is if the data is unrecoverable because no backup or other data security procedures were followed.

Risk Analysis

There are two elements of a system fault risk analysis. The first has to do with the failure rates of network hardware. The second has to do with the site dependent risk.

System Mean Time Between Failures. The probable failure rate of the system is a function of the combined failure probabilities of the components of the network, the file servers, disk systems, adapter cards, hubs, cable segments, workstations, and software. Each has its own probability of failure, usually represented by the MTBF statistic.

MTBF is a very difficult number to calculate. There is a military specification for calculating a device's MTBF by combining the failure probabilities of discrete components. Most manufacturers ignore this procedure because it is time-consuming and inaccurate. Other methods include laboratory stress tests and field observation.

MTBF versus MTTO. MTBF estimates the time between hard failures in physical devices. MTBF estimates do not record the incidence of intermittent "soft" failures. Even in sheltered environments (e.g., well-conditioned computer rooms) transient or intermittent (system) errors are 20 to 50 times more prevalent than hard failures.

In an analysis of system error logs and interviews with field service personnel at Carnegie Mellon University (the results are graphed in Exhibit IX-1-3), system fault data for a 13-server network providing mass storage for 5,000 nodes was tabulated. The data represents 21 workstation years of fault history. Permanent faults (i.e., hard failure) showed a mean time to occurrence (MTTO) of 6,552 hours. Intermittent faults (i.e., a device-specific, weekly recurring fault pattern) showed an MTTO of 58 hours, and transient faults (i.e., random

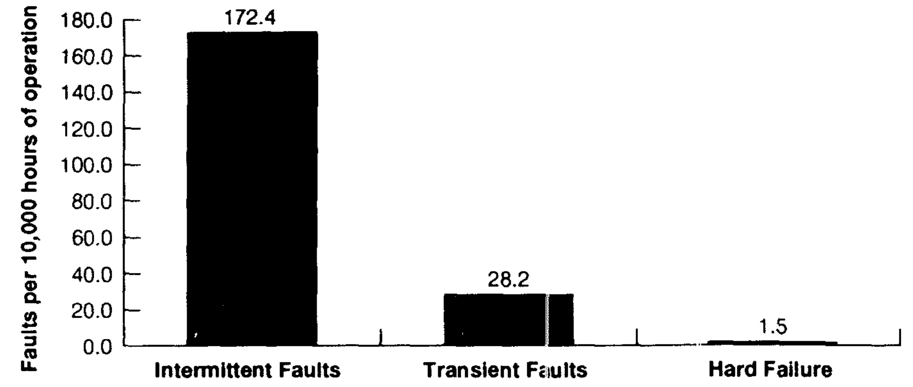

Exhibit IX-1-3. Fault Data for a 13-Server, 5,000-Node Network

nonspecific faults) were reported at 354 hours MTTO. Overall, system crashes recorded an MTTO of 689 hours. Hard failures represented only 10% of system crashes. Further analysis revealed that only about 25% of the transient or intermittent faults actually cause a system crash and that 75% of these faults resulted in recoverable situations.

The study's managers offer the following comments regarding system-level fault modeling: "The manifestations of intermittent and transient faults . . . are much harder to determine than permanent faults. . . . Because the fault is present only temporarily, and because most computer systems do not have online error detection, the normal manifestations of an intermittent fault are at the system level (such as system crash or I/O channel retry). Transient faults and incorrect designs do not have a well defined, bounded, basic fault model. Transient faults are a combination of local phenomenon (such as ground loops, static discharges, power lines, and thermal distributions) and universal phenomenon (such as cosmic rays, alpha particles, power supply characteristics, and mechanical design). Even if models could be developed for transient faults and incorrect designs, they would quickly become obsolete because of the rapid changes in technology."

Site Risk. Every field service manager or system support manager with multiple-site responsibilities will relate that some percentage of their sites are afflicted with an above average number of system faults. Some sites are chronic problem spots. Whatever method is used to estimate predicted system MTBF, actual MTBF performance for a given electronic device will be improved or worsened by the quality of the electrical operating environment. Quality of the power distribution network in the operating environment of a proposed system or device can be evaluated prior to system installation. A power quality evaluation can be used to determine whether the site electrical infrastructure will pose an above or below average risk to the reliable operation of electronic devices that make up the network.

Reliability needs assessments and site risk evaluations should be included as part of the general system requirements assessment process. Once the value of reliability is determined, and the site risk is evaluated, it is easier to design a level of fault tolerance or fault intolerance (i.e., avoidance) into the system that is cost-justified for the operational requirements.

MANAGING FOR MAXIMUM SYSTEM AVAILABILITY

Paying attention to the building power distribution network and enhancing it with power conditioning devices might be considered first- and second-stage fault avoidance techniques. Installation of a backup power system (i.e., UPSs) might be considered a third-stage technique. Bringing an intelligent UPS under the umbrella of a sophisticated network management utility could be considered a fourth-stage practice.

Comparisons of specific methods and architectures are covered later in this chapter. The balance of this section is an overview of each of the stages.

Stage 1: Auditing Building AC Distribution

Building power systems must be safe for people as well as adequate for supporting the basic power requirements of network devices. All branch circuits must be installed and grounded in accordance with the National Electrical Code (NEC) sponsored by the National Fire Protection Association.

Integrity. A well-managed facility has procedures for the regular inspection of the building's branch circuit connections, including wall receptacles, junction boxes, and distribution panels. Usually, however, electrical system inspections are rare. Left uncorrected, loose connections pose a fire and safety hazard, but they can also cause momentary interruptions and high-frequency electrical transients any time there is movement in the connection (e.g., when equipment is plugged in or power cords are bumped).

Capacity. Typical office building branch circuits are rated to carry either 15 amperes (A) or 20 A. Periodic inspections should include at least a visual review of the regular loads on each circuit. It is not a good idea to have such critical loads as file servers compete with harsh or heavy loads such as copiers, or laser printers. Loads for typical types of electronic office equipment are shown in Exhibit IX-1-4. Outlet types for 15 A and 20 A lines are shown in Exhibit IX-1-5.

Heavy loads can push a circuit over its limits if the combined loads exceed the rating. If a computer system is on that circuit, it will crash when the breaker trips. For this reason, no more than two file servers should be placed on each branch circuit.

Microcomputer workstation 1–2 A
File server 4–6 A
Laser printer 8–10 A
Photo copier 10–15 A
Electric space heater 5–6 A

Exhibit IX-1-4. Electrical Loads for Typical Office Equipment

In some retail store electrical plants, which are engineered to support sophisticated point-of-sale networks, an entire sub-panel is dedicated to computer circuits. In these plants, computer branch circuits are at least physically segregated from lighting and other building loads, even if they are not electrically isolated from noise created by those loads.

Several special wiring practices are supposed to improve computer system reliability by reducing electrical noise caused by ground connection effects. However, some of the practices are erroneous, and others are safe but fall short of the desired effect. In the final analysis, about all that can reasonably be expected from a building's power distribution infrastructure is that it satisfy the requirements of the National Electrical Code.

Stage 2: Power Conditioning to Reduce Random System Abends

Harsh loads actually create electrical transients as a normal by-product of operation. Whether these harsh loads are on the same branch circuit as electronic systems is, to some degree, irrelevant. All branch circuits are electrically connected to each other at the source or at the distribution panel. For this reason, specific power conditioning devices are recommended as a second-stage technique for improving reliability of electronic devices by isolating them from the natural building electrical environment.

For supporting distributed digital computing systems, conditioning the electrical supply has traditionally focused on two areas: maintaining a consistent AC line voltage (i.e., voltage regulation) and protecting the computer system from hard failure caused by lightning-induced power surges.

With respect to supporting distributed computing devices that are integral to a network and that have high reliability requirements, a third area of concern

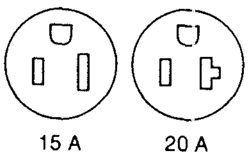

Exhibit IX-1-5. Outlets for 15-A and 20-A Electrical Branch Circuits

must be addressed, isolating the system from lower levels of high-frequency voltage transients (commonly referred to as electrical noise or ground noise).

Voltage Regulation. The power supplies in most microcomputers are very effective at blocking fluctuations in AC line voltage. In fact, these switching power supplies are typically more voltage tolerant than line voltage-regulating devices. Most system experts agree that supplemental voltage regulation for microcomputers is redundant.

Transient Protection. With respect to the issues of lightning protection and noise isolation, there is less consensus among experts. The debate can be distilled to two subtle issues: controlling peak voltage and the rise time of conducted transient voltage events, and providing this control in both normal (i.e., line-neutral) and common (i.e., neutral-ground) modes.

The most thorough and the simplest way to provide this control is through the use of a full output isolating transformer. This methodology was developed in the 1970s to provide the proper electrical environment for reliable operation of mainframe systems in raised floor computer room facilities.

This proven methodology is being applied in microcomputer-based LAN systems because the evolution of Novell's NetWare and the movement toward a standard UNIX have made microcomputers the platform of choice for many critical path business systems. In these environments, certain spurious hardware interrupts (e.g., the nonmaskable and general protection interrupt errors in NetWare) can cause an abend, or system crash. These errors are often symptoms of subtle power quality defects that are not corrected by surge suppression devices alone. A better choice is to use small power conditioning systems that employ the full output isolation and high-frequency grounding principles defined in Federal Information Processing Standards (FIPS) 94 for computer room facilities.

Stage 3: Deciding on an Uninterruptible Power System

Even one power failure–related system crash in a file server or other critical device can be damaging and costly enough to justify the ongoing expense of a UPS. Key issues in deciding on a UPS include how much time it will have to supply power, and what level of system interface should be implemented. If power conditioning was not implemented before UPS installation, power conditioning capabilities of the UPS become an important consideration. Some UPS designs combine effective power conditioning, based on full output transformer isolation, with reliable battery backup in one integrated device.

Backup Power for Simple System Shutdown. If the primary concern is to protect the integrity of data files during an uncontrolled system shutdown, the key question is how much time the system takes to shut down in a controlled fashion. In most NetWare environments, shutting down the server takes less than 2 minutes. In a busy UNIX system, this can take as long as 5 minutes.

For a UPS, the minimum time acceptable is twice that required for controlled shutdown.

Twice the shutdown time is necessary because UPS batteries deteriorate with age. A new system that provides twice the minimum runtime can be expected to have a life of at least 12 months. After 12 months, batteries should be tested, and replaced when they can no longer deliver the minimum runtime required.

It may also be wise to buy enough battery backup capacity to keep the network functioning through brief power outages. If after 2 to 4 minutes, the AC is still off, 2 to 5 minutes of reserve time are still necessary for a controlled shutdown. The calculation for total reserve time under these circumstances is 2 × [(server shutdown time) + (reserve for operating through brief outages)].

Backup Power for Extended Runtime. For some systems, more backup time is needed than the minimum for operating through a brief power outage and a controlled shutdown. Critical path systems need more reserve time to keep the systems online. The following are a few examples.

- A transaction processing fileserver may need 30 minutes of uninterrupted time at night to finish a batch end-of-day process in order to have the system clear and available the next day.
- A retail store system that matches item codes to price files may have to function for 20 to 30 minutes to fully clear the store of customers in the event of a power outage.
- A telephone system may require 7 to 8 hours of reserve time in order to keep emergency communications lines open during an extended outage.
- A premises security system may have to run for days or weeks on emergency power. Generally, requirements for very long runtimes are better met through internal combustion engine generators. In these situations a battery based UPS provides a seamless transition from utility power to generator power and back again.

Power conditioning is especially important when generators are used for power backup because generators are tested periodically. Generator tests create significant transients in the building electrical system as the generator switches in and out of the AC supply line.

UPSs with Intelligent System Interfaces. A UPS supplies power until its batteries run down. If a system administrator is not present when the UPS's audible alarm sounds, the system will crash when battery power fails. To solve this problem, UPSs can communicate with the supported operating system and CPU and trigger an automated system shutdown routine. Most network operating system (NOS) vendors (e.g., Novell, IBM, Banyan) provide basic UPS monitoring modules as a standard part of the delivered NOS package. For UNIX and OS/2, third-party applications that monitor UPS communications and initiate automatic system shutdown are available.

All basic UPS automated monitoring systems can send two types of signals, one for the detection of a power failure and one a low battery (when the system is working on the battery) warning. Some basic monitoring systems support a third signal type, UPS inverter off, which shuts the UPS off after a controlled shutdown of the system. This saves any energy left in the battery, maintaining some reserve in the event that a second outage occurs before the UPS battery has had enough time to fully recharge.

The next level of sophistication for UPS status reporting is a UPS status display of information about the UPS (e.g., load level used, condition codes, service alerts) as well as more information about the operating environment.

Stage 4: UPSs as Manageable Network Devices

In each of the previous monitoring and reporting examples, a file server is the primary recipient of UPS communications. As UPSs are pressed into other networking applications (e.g., backing up bridges, routers, and hubs) other means of status monitoring will be required.

Inband through SNMP. During 1994, a standard management information base (MIB) for UPSs will become available to allow management through the Simple Network Management Protocol (SNMP). The goal is to provide a standard definition of attributes that might be of interest to developers of SNMP-based network management systems. In this mode, a UPS signals its condition directly over network communications pathways to a management console that can then shut down all dependent devices, call for service, reroute wide area traffic around the troubled hub, or take whatever action minimizes the effects of the outage.

An SNMP controlled UPS can also reboot a frozen system by allowing remote control of a UPS's output power. Output could be cycled off and on to force an attached system to go through a cold boot restart

Out of Band Through a Modem. For the less sophisticated network, UPSs can be configured to communicate with a modem that is programmed to call a predetermined number, a pager for example, that will alert a mobile system administrator that there is a fault in a particular location.

ELECTRONIC EQUIPMENT SENSITIVITIES

This section reviews equipment sensitivities from three perspectives. The first is a theoretical model of expected system sensitivities that is based on the physical characteristics of various system elements, from power supply to mother-board components. Next, the results of laboratory experiments are summarized; these experiments measured system responses to induced electrical transients. Last, a real-world example of the effects of interference from the electrical environment on system reliability is examined.

Basic Operating Principles of Digital Electronic Devices

Semiconductor Device Characteristics. The microprocessors at the heart of every computer are very large scale integrated circuits made by depositing trace amounts of various elements into the surface of treated silicon crystals. Software instructions and data (i.e., bits) are represented in these circuits by small electric charges. Instructions are processed by measuring patterns of these small charges.

As tiny and delicate as these circuits are, they can theoretically run reliably for centuries in the right operating environment. There is nothing to wear out because there are no moving parts and no measurable consumption of materials—just electrons trading places in molecules of impure silicon.

At the microcircuit level, very small transient voltages can be mistaken for a digital signal. The transient can change the characteristic of a bit or group of bits, and if it has just the right characteristics and happens at just the right moment, critical instructions or data become altered. These circuits usually operate at 5 volts (V). A transient as small as 0.5 V can change a logical 0 into a logical 1. In some systems, altered instructions cause spurious hardware interrupts that can produce system errors or abends.

Severe transient voltages can completely destroy a microcircuit. At the chip level, as little as 10 V can damage a circuit that is designed for 5 V. Degraded circuits become hypersensitive and eventually fail altogether. The cumulative nature of component degradation is one reason that damaged equipment functions poorly in some sites yet operates fine in others.

Power Supply Requirements: Design Considerations. A computer's digital logic circuits operate on +5 V direct current (DC). The purpose of the computer's power supply is to convert the alternating current (120 V or 230 V AC) supplied in the building branch circuits to a steady, finely regulated 5 V to 18 V DC.

During the 1970s the computer industry used a power conversion device commonly referred to as a linear power supply. The linear supply's first stage was a large 60 Hz transformer; its purpose was to step the incoming line voltage from 120 V or 230 V to some lower level between 7 V and 17 V. The low-voltage AC was then rectified, regulated, and filtered to DC. Regulating the DC output was a significant problem; in the end, supplemental AC line voltage regulation was often specified and as a holdover from these times many people today still assume power conditioning means voltage regulation.

During the early 1980s, a different power supply architecture came into broad use. The new design was smaller, lighter, inherently less expensive, and produced tightly regulated DC output over a wide range of AC input voltage. The advantage of these power supplies, called switching power supplies, made them a natural for use in microcomputers.

The new architecture makes use of a simple, inexpensive diode bridge as a first stage element, which converts the high voltage (120 V or 230 V) AC to high-voltage DC. A high-frequency inverter converts the direct current to high-frequency (10 kHz to 100 kHz) AC. High-frequency, high-voltage AC is then stepped down to lower voltage AC with a small high-frequency transformer, and rectified again to DC, to serve the power requirements of the electronic and mechanical devices inside the computer.

Exhibit IX-1-6 shows a block diagram representation of a typical switching type computer supply. Variations to the basic switching architecture are being introduced to make the switching power supplies easily adaptable to the range of line voltages and frequencies in worldwide use.

Switching Power Supply Performance Characteristics

Switching power supplies inherently are able to provide finely regulated low-voltage DC output over a wide range of high-voltage AC input. In fact, a standard switching supply operating in a nominal 120 V AC environment can usually sustain a well-regulated DC output when AC input falls well below 90 V for extended periods of time. Operation in such "brownout" conditions hampers the supply's ability to compensate for momentary power failures. Operation at extremely low voltage may also stress the rectifier diodes. By and large, however, this power supply architecture does not need any supplemental voltage regulation. In fact, nearly all the available commercial voltage regulators have less tolerance for changes in voltage than a switching supply.

The diodes used in the first stage rectifier of a typical switching type power supply have an upper end rating of approximately 400 V. These diodes are attached directly to the incoming AC line. If there is no protection against transient AC line voltages this relatively delicate electronic front end will not survive long even in a normal operating environment. For this reason nearly all switching supplies incorporate at least some form of protection from transient over-voltages.

Electrical Noise Considerations. Any electronic device that generates signals at a rate in excess of 10,000 Hz and is intended for use in the United States is subject to regulation by the Federal Communications Commission (FCC). Similar requirements are in effect in Europe. The purpose is to keep radio communications free from interference. Most computer power supplies comply by having some form of radio frequency (RF) filter on the input of the power supply. The primary purpose of these input RF filters is not to protect the computer system from conducted AC line noise, but to prevent high-frequency noise that is created in the computer or power supply being conducted into the AC lines. These filters typically block noise from either direction, but their effect on inbound noise is predictable only when the noise is of similarly low energy content and frequency as the noise generated by the computer power supply and digital circuits.

Conducted AC voltage transients, or power line noise, can easily fall outside this range. The energy content of power line transients is typically significantly

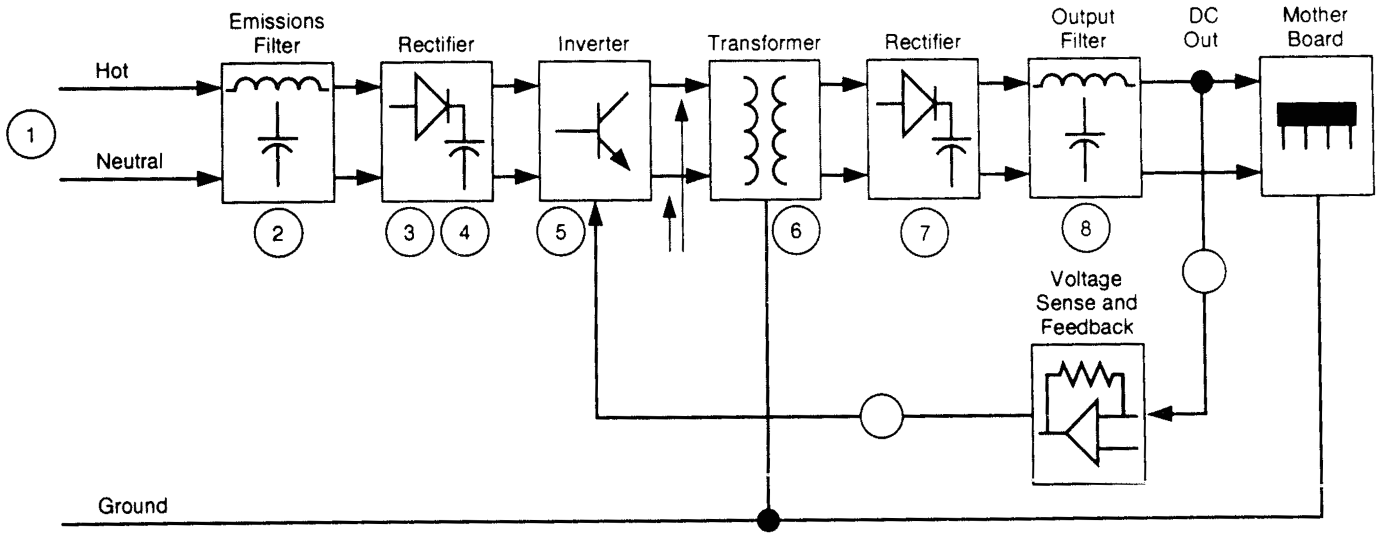

1. AC line in: hot and neutral
2. Surge and noise suppression: diode surge protection and RF emissions filter keeps switching frequency noises from being conducted onto building branch circuits (these would act as broadcast antennas for radio frequency noise that would interfere with other communications devices).
3. Diode bridge rectifier: output is rectified 170/(340) V DC pulses
4. Storage capacitors; output = 170 (340 V DC (smooths AC pulses to an acceptable ripple)
5. Switching circuit (high-frequency inverter): output = 10 kHz to 100 kHZ square wave AC, 170 (340) V peak
6. High-frequency step-down transformer: output = 10 kHz to 100 kHz AC, 5/12 V peak
7. Diode rectifier: output is pulsed DC
8. Output filter: eliminates switching frequencies to provide clean, "noise-free" DC to output (regulated 5 V DC to mother-board logic circuits)

Exhibit IX-1-6. Block Diagram of a Switch Mode Power Supply

greater than the computer's internal noise. Radio frequency rated filter circuits have no measurable effect on large incoming transients.

Noise can cause a loss of output regulation or failure of the power supply. High-frequency transients can also couple through the power supply circuit elements and find their way onto the low-voltage bus of the system mother board. Once on the mother board, these noise impulses can corrupt data or damage electronic components.

Measured System Sensitivities of Typical Microcomputers

The previous section reviewed the theoretical sensitivities of microcomputers at both the chip and system levels. The chip level sensitivities are hard facts. They represent the physical characteristics of semiconductor devices, tolerances published by designers and vendors of high-density integrated circuits. System sensitivities can never be described as precisely. At the system level, the design and environmental variables are hard enough to measure in a laboratory setting let alone predict over the full dynamic range present in the real world.

It is generally accepted that capacitive coupling is the mechanism by which transient voltage disturbances are transferred to a computer's internal components. Even so, many experts in the fields of computer system design and power protection design disagree on just how sensitive systems are to various types of transient electrical events. The basis of the disagreement centers on assumptions regarding the filtering and attenuation properties of power supplies and other in-system, DC circuit protection schemes.

During 1992, research on system-level sensitivity to electrical phenomena was published by PowerCET Corp. (Santa Clara CA), an independent consulting organization. The study compared the electrical sensitivities of two microcomputers, one a well-designed model from a major computer vendor and the other a cost-leader 286 IBM clone.

Summary of PowerCET Research. PowerCET's research demonstrated in a controlled laboratory environment that the well-designed microcomputer was less sensitive to electrical-power events than the budget machine, but on both computers interference currents were measured on the +5-V DC bus that supplies the system mother board. These interference currents caused functional problems in both machines, even when the external noise voltage was under 50 V.

The report concluded that system sensitivity is a function of:

- Frequency response and input impedance of the emissions (i.e., FCC) filter.
- Bypass capacitance for common mode interference.
- Data interconnection, filtering, and protection of data I/O ports.
- Wiring and circuit board layout.

- Lateral system grounding schemes.
- The type and speed of active processing.

Symptoms of Electrical Interference. PowerCET's research demonstrated that no commercially available computer hardware is completely immune to electrical interference. Electrical interference can cause spurious hardware interrupts leading to abends in operating systems (e.g., NetWare or OS/2) that typically run in protected mode.

Novell's engineers recognize the consequences of hardware susceptibilities to electrical transients. In the system messages manual for NetWare 2.2, "poor power line conditioning" is cited as a probable cause of GPI, NMI, and at least 15 other system errors. Novell's *AppNotes on 386 NetWare System Messages* say, "The majority of NetWare operating system messages are of the fatal/abend type. Fatal/Abend messages are usually caused by consistency checks. Not all consistency check errors are caused by software anomalies. These errors might also be related to corrupted OS files, defective memory chips, static discharges, faulty power supplies or power surges and spikes."

This author's experiences with trouble-prone NetWare installations suggest a strong correlation between fatal/abend occurrences and the presence of electrical noise. The following is a representative scenario.

Case Study Situation. A team of consultants was working to develop a client/server application with a graphic user interface that would automate processing of home and auto insurance applications in regional offices across the country. The system was being implemented on a NetWare 3.11 platform running on an IBM 8595 server. A prototype system was scheduled for rollout. One of the critical milestones that needed to be reached before moving ahead with development was to verify the stability of the system.

Problem. Random hardware interrupt errors that crashed the server occurred four times in a three-week period. These unexplained system faults caused delays in development and were a source of much frustration. Literally, tens of consulting hours had been spent trying to identify the cause of the problem.

The systems development team assumed that the power supply was satisfactory. The electrical supply in the development lab consisted of dedicated, isolated-ground circuits with surge-suppressed outlets. The server was powered by a 900-volt-ampere smart UPS of a suppressor/filter type architecture. A system engineer from Novell was brought in on the problem. Based on the nature of the hardware interrupt symptoms, the system engineer suspected the cause to be related to power and grounding anomalies.

Diagnosis. An evaluation of the electrical service supplying the server room revealed no irregularities. All circuits were correctly wired and nominal voltage levels were between 115 V and 122 V. A power quality recording instrument was installed to measure and record the quality of power to the server. Transient impulses of 70 V to 80 V were recorded during the period that the system reported a series of spurious hardware interrupt errors that the system engineer from Novell believed were related to power problems.

Solution. The existing suppressor-type UPS was replaced with a UPS that incorporated full output transformer isolation and high-frequency filtering circuits. Random hardware interrupts no longer occurred. System stability was demonstrated and the project was put back on schedule.

CHARACTERISTICS OF BUILDING ELECTRICAL DISTRIBUTION

Typical Building AC Distribution Circuits

Typical branch circuits in offices or factories are the same as home electrical circuits. They carry 100-V to 120-V AC power from a building source (or subsource), and distribute it safely throughout the building.

The three primary parts of a branch circuit are the receptacle (wall outlet), the current carrying conductors (the wires), and the overcurrent protection device located in the main panel or in a subpanel (typically a resettable circuit breaker).

These circuits are sized to carry either 12 A or 15 A continuous and are protected from overload by either a 15-A or 20-A circuit breaker. When electricians refer to a circuit as being either a 15-A or 20-A circuit, they are referring to the maximum current capacity of the circuit. Circuit capacity is determined by the gauge of the current carrying conductors and the current rating of the receptacle. Conductors or receptacles rated to carry 12 A continuous would become hot and thus pose a fire hazard if not protected by a circuit breaker. A short circuit anywhere in a branch circuit will immediately cause the circuit breaker to open. Loading a circuit to near, or just beyond, its capacity can cause the breaker to trip from time to time in response to random overload conditions resulting from a combination of load interaction and low line voltage.

Midrange and mainframe computers, high-capacity printers and duplicators, and large PBXs require more power than can be delivered by a 15-A or 20-A branch circuit. These high-capacity circuits have special receptacles so that loads that need more power cannot be inadvertently connected to standard circuits. There are three wires in a branch circuit. One wire (commonly referred to as the live, hot, or line) is energized at the source. It is always live, ready to supply current the moment it is connected to a return path that completes the circuit. The second wire (neutral) serves as the proper return path. The neutral wire is connected to an earth ground at the building power source. The neutral-earth connection is typically made at the main transformer but is sometimes allowed at the main panel.

Safety Ground. The connection of the electrical power system to earth-ground is a safety measure, intended to reduce the risk of shock to people using electrical devices and to reduce the risk of fires caused by a breakdown in the electrical system. The National Electrical Code prescribes standards which govern the diameter of wire, type and thickness of insulation, circuit breaker characteristics, receptacle construction, maximum distance of the circuit from the breaker panel, and specific connection practices designed to assure safe grounding of the building power distribution system.

Classes and Causes of Power Quality Defects. Power loss is a visible defect. Electrical transients are not. Yet as has already been shown, even relatively low-level electrical transients can reset a file server running an operating system in protected mode. Experts are still developing definitive labels for different types of power quality defects, but it is probably sufficient for network managers to be familiar with the following two key concepts:

- Power quality defects are conducted to the system through two pathways—normal mode and common mode.
- Power quality defects can be put into one of two classes—fast and slow.

Typically, electronic systems are more sensitive to fast transient events than slow ones and are more susceptible to interference from common mode impulses than from normal mode impulses. Fast defects typically occur much more often than slow ones. Exhibit IX-1-7 summarizes the characteristics of these problems, which are described in more detail in the following paragraphs.

Normal Mode and Common Mode. These terms describe specific pathways that conduct electrical energy in branch circuits.

Normal Mode. In normal mode, defects occur between the current carrying conductors in a branch circuit (i.e., the wires connected to the two flat

Normal AC Power

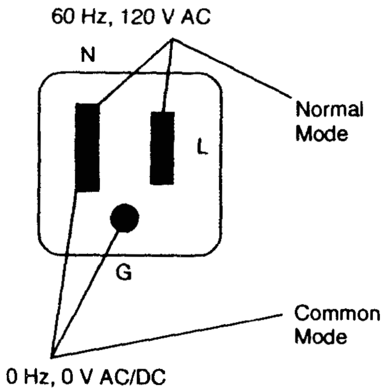

Defects in AC Power

Fast	Slow
Surge, Spike, Transient, Blinks, Noise	Low Line High Line No Line
Surge, Spike, Transient Noise	"Harmonic" or periodic offset

Exhibit IX-1-7. Power Line Defect Summary

slots on a typical wall socket). Normal mode impulses must pass through the system's power supply to contaminate the low-voltage DC that computers use internally.

Common Mode. In common mode, defects occur between either conductor and the ground connection (measured between neutral and ground). Common mode defects bypass the system's power supply and more easily contaminate the computer's internal electric environment.

System designers usually consider only normal mode power problems. Common mode events are not even acknowledged as real in some circles. But both modes are real and should be considered when planning power conditioning implementations.

Power quality defects occur over a frequency continuum. This continuum extends from power transmission frequencies (60 or 50 Hz, depending on country) to approximately 5 MHz. Beyond 5 MHz the energy content of conducted transients is quite low; also, noise above 5 MHz is as likely to be radiated into the computer system as it is to enter through the power line.

To simplify the analysis of power quality defects that affect electronic systems, they are classified as either low frequency or high frequency. Low-frequency events are in the range from 50 Hz to 20 kHz. High-frequency events occur above 20 kHz.

Low-Frequency Events. The following types of power quality defects are classified as low frequency. Exhibit IX-1-8 summarizes the causes and consequences of these irregularities.

High Line. These slow, normal mode surges can damage power supplies.

Low Line. Momentary sags are caused when other large loads in a building turn on and pull a lot of current from the electrical system. Extended periods of low voltage on a circuit are the sign of a utility feed problem, a phase imbalance, or poor wiring within the building.

Power Failure. Normal mode outages, both momentary and extended, shut down the system and corrupt data, but the loss of power in itself does not damage hardware. Power interruptions of short duration create transients, which do damage computer systems. Utility failure statistics indicate that utility outages are rare events. The total outage time as a percentage of time a utility is up and ready is approximately 0.02%. Utilities also maintain that when outages do occur, they are usually resolved within 15 minutes. Utility company statistics do not include momentary "blinks"—typically site-specific phenomena caused by other loads within the building or such normal utility activities as fault clearing or grid switching.

Harmonic Disturbances. 60-Hz hum and higher-order harmonics of 60 Hz can produce a repeating voltage offset on data communications lines. This leads to increased error checking activity or, in extreme cases, completely interferes with data communications.

Common Name	What It Is	What It Looks Like	Caused By	Why It Matters
Sags Brownouts Low Line High Line Over-voltage	Slow deviations from, or aberrations in power transmission wave form standards. (Normal mode only) Normal peak = ± 170 V when rms is 120 V.	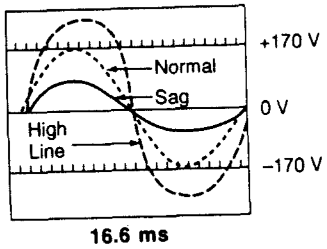	Over loaded circuit, unbalanced phases, high impedance neutral conductor or utility feed fault.	• Too high (more than 200 V peak, more than 140 V rms) can stress power supply. • Too Low (less than 108 V rms) lessens system ride-through capabilities if sustained.
Blackout Outage Dropout	Visible loss of power transmission wave form (normal mode).	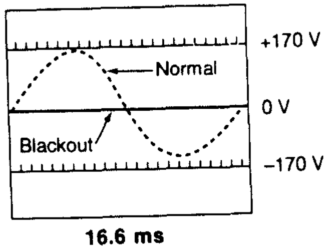	Overloaded circuit causes breaker trip, or utility feed fault.	No power, no system. Data at risk. Productivity suffers.
Harmonics 3rd Harmonic	Sinusoidal or complex wave form that repeats hundreds of times per second. Can occur in normal or common modes.	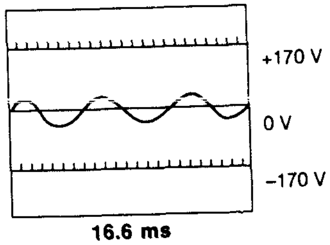	Nonlinear loads on multiple phases sharing common neutral conductor. Illegal neutral-ground bond in loaded branch circuit.	Neutral-ground (common mode) voltage wave forms can cause communications problems on networks running unbalanced communications lines (i.e., ground as signal return path).

Exhibit IX-1-8. Low-Frequency AC Power Line Disturbances

Fast Events. Fast or high-frequency defects (those with an electrical frequency above 20 kHz) include spikes, transients, and other defects collectively referred to as noise. Random, fast-edged voltage events occur frequently in normal and common modes—most often both. Electrical noise can confuse system logic and damage electronic components, resulting in random system lock-ups and premature board failure. The causes and consequences of fast defects are summarized in Exhibit IX-1-9.

Grounding-Related Issues

Separate from but related to power quality defects is the subject of grounding. Four areas of interest with respect to grounding are relevant to computer systems:

- Human safety
- Electronic signal reference
- Controlling electrostatic discharge (ESD)
- Ground skew, ground loops, and ground offset

Grounding details can become complex but there are two simple key points:

- Ground does not mean the electrical potential of the ground or earth; to refer to an electrical ground is to refer to a point of reference.
- When any decision about grounding has to be made, the first priority is human safety.

Making safety the first consideration means putting electronic reliability second. The needs of one can create an environment that is less than ideal for the other, and the grounding compromise leads to power quality defects that affect system reliability.

Safety Ground. Safety requires that in issues of distribution of electric power, the earth be used as the ground reference. Inside buildings, grounding all wiring to earth gives the electrical distribution system and humans the same electrical reference point. The National Electrical Code requires low-resistance ground connections to earth in order to protect people from electric shocks. One published reference to ground requirements, taken from FIPS 94, reads as follows:

"Touch voltage is the voltage between any two conductive surfaces that can simultaneously be touched by an individual. The earth may be one of these surfaces. In the event of insulation failure, any (exposed) electric charge . . . must be drained to 'ground' or to an object that is reasonably grounded."

Signal Ground. The ground for an electronic signal does not have to be earth. The purpose of this type of ground connection is to give digital logic circuits an unambiguous point of reference from which to distinguish the electrical difference between a logical zero and a logical one. Inside the computer the logical 1s and 0s are represented as 5-V and 0-V pulses.

Common Name	What It Is	What It Looks Like	Caused By	Why It Matters
Surge Spike Transient Noise Blinks	Sudden, fast voltage changes independent from the 60 Hz power transmission voltage wave.	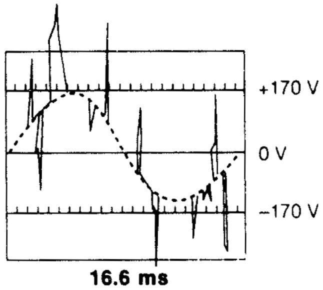	• *External:* Very large events typically originate outside a building on the utility grid. Natural events like lightning, or utility system maintenance practices produce very large events. These events typically occur infrequently. • *Internal:* Other loads in the building—elevators, copy machines, air conditioners, space heaters, dimmer switches and other computer systems can produce hundreds of events a day.	• Fast edged transients are high frequency energy. Normally high impedance pathways act as low impedance pathways at high frequency. Fast-edged transients couple into the low voltage digital world and alter memory states or damage VLSI circuits. *Comments:* Most systems have good low-level noise filtering and over voltage protection built into the system power supply—normal mode. System sensitivities are typically more acute for common mode

Exhibit IX-1-9. High-Frequency (Fast) AC Power Line Defects

The 5-V and 0-V pulses must be measured relative to some reference point. The 0-V reference point can be established anywhere: 5 V or 5,000,000 V above an earth ground. Whatever the reference, it is important that it remain free of transient voltages that can momentarily alter the electrical potential of the reference ground.

The aim is to ground or short-circuit all differences in electrical potential that might otherwise be detected by the electronic system and misinterpreted as a logic instruction. As with ground connections intended for electrical safety, it is essential that the path to an electronic reference ground be of low resistance. Very often, the safety ground is used to provide the signal reference point. This design is illustrated in Exhibit IX-1-10.

Electrical noise is a high-frequency phenomenon. Resistance, or more correctly, impedance (i.e., resistance to alternating current) increases with the length of the conductor and with the frequency of the transient impulse. Low impedance at high frequencies can be provided only if the ground point used as the system's reference ground is physically very close to the system's logic circuits. The ideal system grounding arrangement is shown in Exhibit IX-1-11. For distributed systems this is not, unfortunately, practical, because the elements of such a network are sited throughout a building, physically distant from a building power source and the neutral-ground relationship established there.

Controlling Electrostatic Discharge (Static Electricity). Static discharge occurs everywhere. Charges build up on people as they walk across carpets or even move around in their chairs. A static discharge is usually a

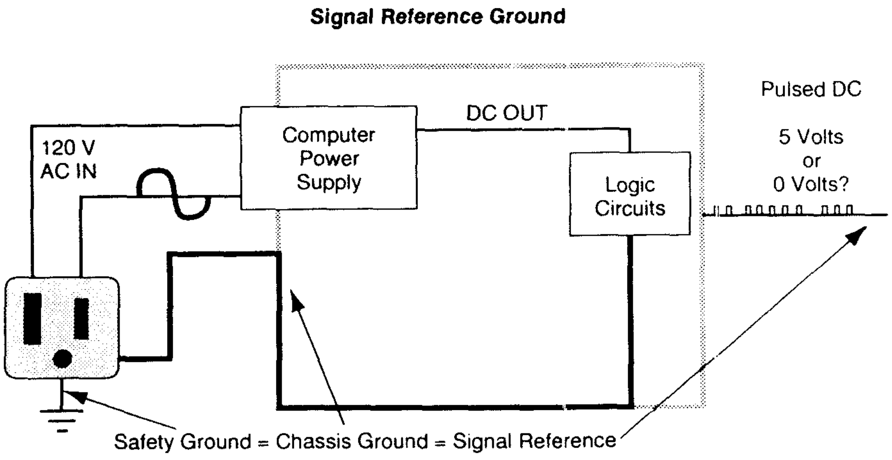

Exhibit IX-1-10. A Properly Grounded Electrical Branch Circuit

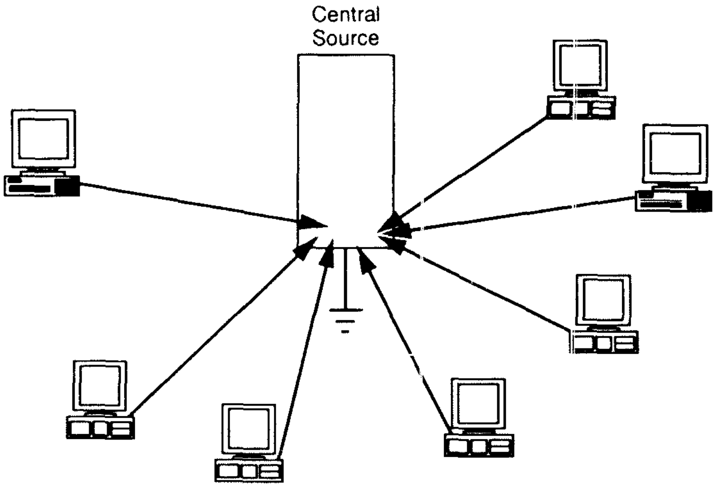

Exhibit IX-1-11. An Ideal Computer Sy:stem Ground Ar:angement

high voltage that is released very fast. Computer circuits can be especially susceptible to static electricity. To protect circuit cards from static effects, many repair technicians wear static straps while handling circuit cards, and new cards are packed in special antistatic bags to protect them during shipment.

Touchplates and static straps help drain static electrical charges before they build up enough to harm circuits. These can be helpful in office environments and they should be grounded to the system's reference ground point. Use of antistatic sprays, specially treated materials, and maintaining proper building humidity are all practices that minimize static buildup.

Ground Skew. Ground skew occurs when a surge of transient energy is applied to a power line to which two or more pieces of networked equipment are attached. The ground skew voltage is the instantaneous voltage difference between any two pieces of equipment. This voltage develops when transient ground currents occur and the impedances of the equipment and power wiring are not identical at all frequencies (illustrated in Exhibit IX-1-12).

In a local network segment that connects several pieces of equipment even slight length differences in the wires between the point of origin of the surge and the equipment will cause a voltage skew between the equipment grounds. This voltage stress shows up across the data communications cabling. In severe cases it can damage network adapter cards. Ground skew related voltage stress is aggravated by the widespread use of individual surge suppressors. At present, the only way to avoid ground skew problems is to use electrically isolated data communications connections.

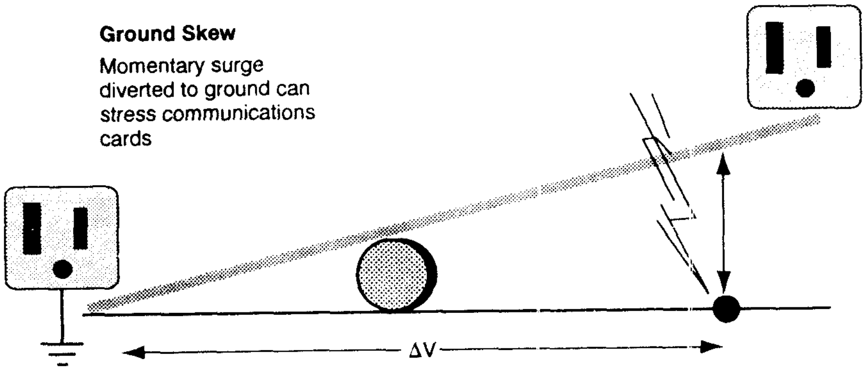
Ground Skew
Momentary surge diverted to ground can stress communications cards
ΔV

Exhibit IX-1-12. Ground Skew in Electrical Power Line

The Physical Model: Typical Building Communications Cabling

Communications lines between two distant elements of a network can complete an electrical connection between two buildings' electrical grounds. This circuit is called a ground loop and is susceptible to transient ground skew or continuous ground offset.

These occur when different ground locations are at different electrical potentials (either continuously or momentarily) because of wiring problems or transients. The continuous or momentary potentials can cause continuous or momentary currents in the communications lines that are unrelated to the communications signal. Whether this presents a real problem or not depends on the communications topology, a site's characteristics, and the design of the power treatment devices used throughout the network.

Data Communications Topologies. Communications topologies can be classified as balanced or unbalanced. If the ground connection is used as a return path or if either signal wire connects to ground the system is unbalanced. The two arrangements are illustrated in Exhibit IX-1-13.

Ground skew and ground loops are typically not a problem in networks connected with a properly isolated, balanced communications topology. Telephone circuits are an example of balanced communications lines. Balanced communications topologies use a pair of wires for each signal and these are not connected to earth ground. No ground loops occur in networks that use a balanced topology as long as adequate isolation is provided by the network interface card or an external balun.

Standards for Ethernet LANs call for significant levels of AC and DC isolation (2,250 V DC and 1,500 V AC). This is accomplished either through optical isolation or by using baluns. But even balanced communications topologies can absorb transients by means of external coupling if the electrical

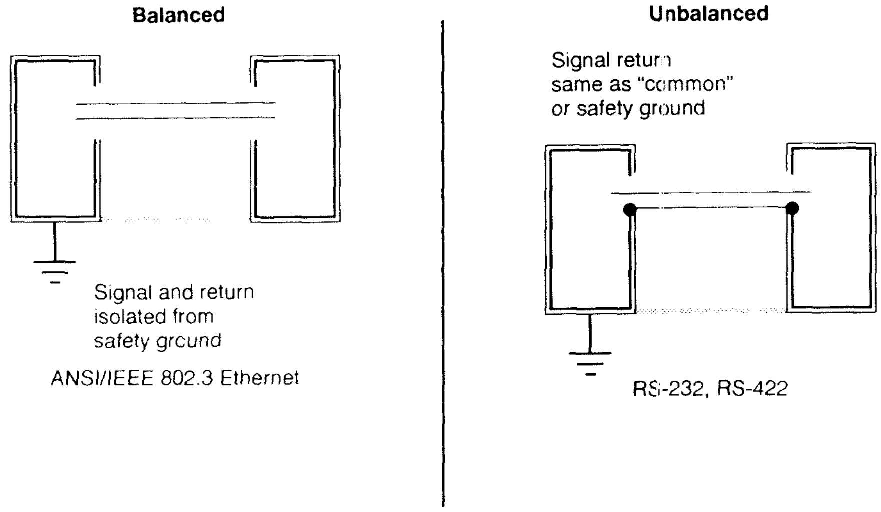

Exhibit IX-1-13. Balanced and Unbalanced Communications Topologies

isolation is inadequate relative to the magnitude of the surge voltages. For example, in installations that use unshielded twisted pair to connect buildings in a campus network, lightning induced transients can break down the isolation barrier and damage devices on either or both ends of the cable segment.

In balanced communications topologies, ground skew concerns are limited to control of transients. However, for shielded coaxial cables typically used for Ethernet backbones care must be taken to ground only one end of the shield, as called for in the guidelines for this type of cable. If both ends of the shield are tied to grounds of different potentials, the shield completes the ground loop resulting in susceptibility to ground offset and ground skew, which can overheat the cable segment or interfere with communications traffic.

For LANs using RS-232 or other unbalanced communications links (or for certain exceptions to an otherwise balanced system), an awareness of ground problems is invaluable. In unbalanced networks, differences in electrical potential between multiple ground connections are almost always a problem. In these environments, specific installation techniques recommended by the system vendor must be carefully followed.

Simple tests for the diagnosis of electrical power for network installations are described in Exhibit IX-1-14.

IMPROVING BUILDING WIRING SYSTEMS

When it comes to designing for maximum reliability of network devices, about all that can be properly asked of building wiring systems is that they be installed

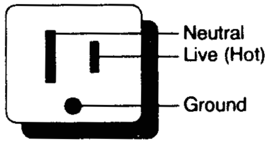

Test	Problem Sought	Test Positive Consequence		Recommended Action	Tools	Method
		Safety	System Performance			
Circuit wiring check	• Reversed Polarity • Open Ground • Open Neutral	X X X	 X	Code Violation Safety Issue Alert Facilities Manager	Inexpensive wiring checker Voltmeter and Probes	Test each receptacle. Note abnormal readings on LAN planning sheets.
Circuit voltage check	• Line-neutral nominal voltage out of bounds? (less than 108 V or more than 127 V rms)	X	X	Alert Facilities Manager: Load balance problem, utility feed problem or high impedance neutral	Voltmeter and Probes Oscilloscope and Line Viewer	With circuit under normal load, test each receptacle. Note abnormal readings on LAN planning sheets
	• Neutral-ground voltage more than 3 V?	X	X	Alert Facilities Manager abnormal, wire gauge, distance or bad connection.		Use wiring diagnostic chart for interpretation of readings.
	• Neutral-ground voltage more than 0.8 V –1.5 V?		X	Use low impedance, power conditioning transformer UL neutral-ground bond.		
	• Open Neutral	X	X	Alert Facilities Manager – safety issue		
	• Out of bounds and outage incidence over time	X	X	Assess need for UPS system or power distribution upgrade	Recording Power Quality Monitor	Leave instrument online at least 1 day, best 1 week. Note results.

Active Ground Loop	• Current flow in communications cables	X	X	Alert Facilities Manager May indicate illegal neutral-ground bond in branch circuit. First choice, correct wiring; next choice; break the loop. Correct network performance problems with baluns or fiber-optic splice in communications line.	Current Clamp (Amp probe) and digital multimeter	Clamp amp probe around communications cable or system power cord. Reading of 1 A or more means trouble. Reading of 0.1 A or less is probably ok. (With a break-out cable, clamp probe around ground wire only.)
High Frequency Noise Incidence	• Preinstallation site characterization. Predict future system problems and service requirements.		X	Use low-impedance, power conditioning transformer with neutral-ground bond. Or power filter with series inductor with appropriate let-through performance characteristics in normal and common modes. Device should limit peak voltage and limit impulse edge-speed to appropriate values. Device should not shunt surge currents to ground. Use distributed, point-of-application methodology for best results.	Oscilloscope and Line Viewer Recording Power Quality Monitor	With circuit under normal load, test key or representative receptacles. Compare results vs. system error log or experience-based norms.
	• Problem site: correlate system fault with power events.		X			

Exhibit IX-1-14. Power Quality Tests for Network Installations

and maintained in accordance with the requirements of the National Electrical Code. The managers of well-managed facilities regularly inspect the building branch circuit connections, including those at wall receptacles, junction boxes, and distribution panels. Unfortunately such well-managed facilities are rare and inspections of electrical systems are infrequent.

Because the integrity of an existing building wiring system may be uncertain, a special branch circuit methodology known as dedicated isolated ground (I/G) is often recommended or required by better system integrators as part of a standard site preparation procedure. Most often, this method entails remaking critical branch circuits prior to system installation.

Isolated Ground Circuits

To more properly address the problem of interference from conducted electrical noise on building branch circuits, safe exceptions were added to the National Electrical Code that are intended to provide a ground-noise reduction method for computer system installations. These I/G circuits are usually identified by special orange colored receptacles.

Unfortunately, contractors and electricians often misunderstand the special installation and connection techniques for I/G circuits. And even when they are installed correctly, usually at a premium in cost, I/G circuits are rarely any more free of electrical noise than a properly executed standard branch circuit.

The IEEE Green Book (IEEE Standard 142-1991) is a thorough, up-to-date reference resource on accepted and recommended grounding practices for industrial and commercial power systems.

IMPROVING POWER QUALITY

Surge Suppression

Transient voltage surge suppressors (TVSSs) were originally developed to protect lights and motors from lightning-related voltage surges. They may limit immediate damage to a system from large lightning-like transients, but although it is widely believed that surge suppressors also block electrical noise, they do not. Exhibit IX-1-15 shows a portion of an alternating current sine wave. Fast transient events that occur at points B and C are not blocked by the surge suppressor.

Beyond their failure to block electrical noise, these devices introduce circuit problems of their own. The primary shortcoming of TVSSs is the tendency of these devices to convert normal mode events to common mode events. Most commercial surge suppression devices direct unwanted surge currents to ground, and the ground can conduct those currents into the computer system.

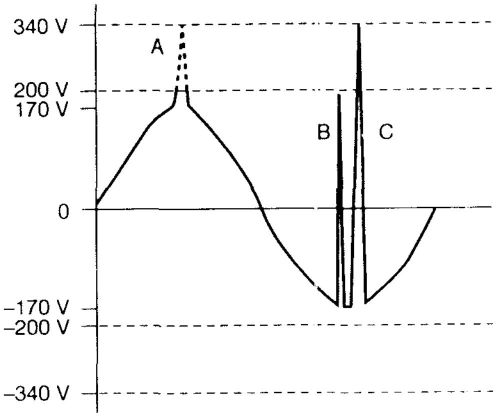

Exhibit IX-1-15. Transient Pulses in an AC Sine Wave

For these reasons, TVSS technology is not recommended for standalone mission-critical electronic systems and especially not for networked systems with multiple ground connections.

Filters

Power Filters. A new category of power treatment device offers a kind of technological middle ground between surge suppressors and transformer-based power conditioning systems. Varying performance and specifications that confuse even experts make filter selection a difficult task, however. At the low-cost end are various grades of surge suppressors with added electromagnetic interference (EMI) or radiofrequency interference (RFI) filters. At the high end are series-connected power inductors combined with surge suppression and output filtering modules.

Well-designed high-end filters offer peak voltage and edge-speed control that is significantly better than what is achievable through simpler surge suppression technology. A high-end power filter's common mode performance can be better than a surge suppressor but will never be as good as isolating transformers that are part of a complete power conditioning system. Some systems or sites may require the level of common mode performance offered only by transformer-based systems.

Voltage Regulators

Earlier sections of this chapter explored why computer power supplies do not need supplemental voltage regulation. External voltage regulators are redundant at best, and at worst they introduce problems of their own.

There are generally two types of regulators, tap changing transformers and constant voltage transformers. Both devices tend to be unstable in certain situations. This characteristic can extend the duration of an otherwise harmless power quality defect, which makes it harder for a computer power supply to recover after a momentary "blink" in AC power.

Tap Changers. Tap-changing, voltage regulating transformers operate by switching between multiple transformer output taps, boosting or dropping output voltage in increments. These devices run cooler than older ferroresonant regulators and operate relatively quietly. Still, they represent a complicated solution to a non-problem. The technology does not guarantee effective noise isolation (performance varies by design), and with poor designs voltage transients are created when the device switches between taps.

Ferroresonant/Constant Voltage Transformers. Ferroresonant transformers were invented in the 1930s to provide constant voltage to neon lights. They provide consistent output voltage over a range of input voltages and also offer good noise isolation. Noise isolation is desirable but voltage regulation is not very useful for microcomputers. Worse, a tank circuit effect actually causes competition with the computer's power supply for energy to recover from a "notch" in the incoming power. This can extend the duration of the notch beyond the ride-through capabilities of the computer power supply, causing the computer to re-boot.

Isolating Transformers and Power Conditioning Systems

High-Impedance Types. Ferroresonant transformers and other high isolation transformers provide very good protection from conducted line noise—much more protection from transients than can be provided by surge suppressors or power filters. However, they typically have high output impedance which can reflect or amplify electrical noise generated by computers and related hardware. High transfer impedance makes them incompatible with the high current crest factors that are typical of switching power supplies. Problems of heat and audible noise make them unpopular in environments shared with people.

Low-Impedance Types. Power conditioning systems built using transformers with low transfer impedance were developed in the early 1980s. The design is much more compatible with the power consumption characteristics of switching power supplies. Low transfer impedance eliminates the efficiency and instability problems associated with earlier high-impedance isolating transformers. Low output impedance is made possible by an output filter circuit that eliminates all high-frequency noise. These systems typically provide complete protection from conducted transients without any bad side effects. They also provide a safe and absolutely clean reference ground in close proximity to digital electronic devices.

UPSs

An uninterruptible power supply (UPS) is a battery-based power backup device that typically incorporates one of the previously mentioned surge and noise suppression technologies. The type of power conditioning technology employed determines the UPS's power conditioning capability when the AC line is functioning properly. Other issues in choosing a UPS are transfer time, inverter wave shape, battery maintenance, and the UPS architecture.

Transfer Time. It should be faster than 16 to 20 milliseconds (ms) to transfer from line to battery power without affecting the computer system's power supply. Most commercial UPSs are fast enough: transfer times of 2 ms to 5 ms are typical.

Inverter Waveshape. The inverters in most standby UPSs sold for microcomputer applications produce output current whose waveshape is described either as pseudosine, modified sine, or a stepped approximation of a sine wave. This waveform has the same peak and root mean square (rms) voltage content as a sine wave and is ideal for driving the input bridge diodes of switching power supplies. This form of inverter output has certain advantages over inverter designs that deliver a pure sine wave, in particular, efficiency, reliability, and lower cost.

However, sinewave output is more universally applicable. Some of the more sophisticated computer power supply designs include autoranging detection and switching circuits, or power fail detection circuits, and these can sometimes be confused by non-sinusoidal voltage waveforms.

The only unacceptable inverter output waveform is a true square wave. A square-wave output may have the proper rms voltage, but its peak voltage level will be much lower than the switching supply expects. Low peak voltage causes the duty cycle of the computer power supply diodes to be extended—this can stress the components to the point of failure.

Battery Maintenance. The most critical and most often overlooked component in any UPS is the battery. UPS batteries are made from the same materials and are generally of the same design as car batteries. As with a car battery, they can wear out just when they are most needed.

UPS batteries in older or in contemporary low-cost designs can be tested by pulling the plug on the system when the computer is up and running and timing its operation until the UPS battery is drained. Safer UPS designs include a test routine that simulates a power failure and runs the computer on the battery for a period of time. The test is triggered by pushing a test button or through UPS monitoring software that resides on the computer system and communicates with the UPS. More sophisticated designs use subtle measurement circuits to regularly test battery condition in the background and alert the network manager through front panel messages or through system-resident monitoring software.

Replacing batteries in most small UPS (systems under 2 kilovolt-amperes) requires off-site service by trained technicians. Larger UPSs require on-site maintenance by trained service technicians. New designs allow the user to replace batteries on site. This feature can simplify maintenance logistics, reduce costs, and facilitate proper recycling of batteries, which pose a potential hazardous waste problem. Some designs also offer "hot swap" battery packs that are ideal for networks with high availability, total up-time requirements. Hot swap batteries can be replaced without bringing down the UPS or the system it supports.

Power Conditioning Capability. The power conditioning capability of any UPS, online or standby, is only as good as the filtering circuitry built into the UPS. There are a limited number of UPS designs that provide superior filtering capabilities. These designs can be identified by the presence of a full output isolation transformer. These UPSs are slightly larger and heavier than those that use small autotransformers to step inverter output to normal line levels or boost low input lines without draining battery reserves. Autotransformers do not provide full output isolation.

UPS Architecture Review

There are two basic architectures for UPS systems: online and standby. Hybrid designs are a subset of the standby technology family but blend in some of the performance characteristics of online technology.

Online UPSs. The online UPS architecture supplies continuous power through the UPS's inverter. This architecture is illustrated in Exhibit IX-1-16. The inverter is typically fed either by the battery (during outages) or by the battery charger (during normal conditions). Online designs provide truly uninterrupted power with no transfer time. This was an important performance characteristic, especially for midrange and mainframe systems. However, in an

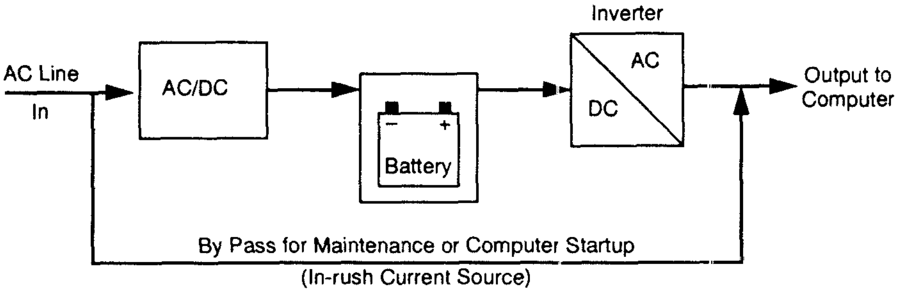

Exhibit IX-1-16. Block Diagram of an Online UPS

online UPS, the power electronics function continuously. Online designs typically have half the MTBF performance of standby units, whose power electronics operate only when needed (i.e., when commercial power fails). The inherently higher cost and generally lower reliability of online units have led to a relative decline in their popularity.

Offline/standby UPSs. These supply a computer with commercial AC through passive circuits, switching to their power electronics and batteries only when an AC line fault is detected. This architecture is illustrated in Exhibit IX-1-17.

Solid state power switches allow for instantaneous switching between line and battery. Most standby UPS designs include a line failure detection circuit that inhibits switching for 2 ms to 4 ms to confirm the outage is real and avoid activation of the UPS unnecessarily.

Hybrid Designs. These are a newer subset of the offline/standby class. They are labeled hybrid because they incorporate additional circuitry that provides useful performance advantages that blend the best characteristics of online and standby designs.

Line Interactive. The classic or purist definition of a line interactive UPS architecture is one that runs the computer from both AC and DC sources simultaneously. The advantages include significant overload tolerance and fine regulation of output power in response to changing input conditions and dynamic output requirements. Such power supplies are challenging to design, expensive to build, and consequently are not very common.

The term *line interactive* is now used by some UPS manufacturers to mean a modification of a standby design. In these, gross output voltage regulation is provided by engaging an autotransformer to step output voltage up to

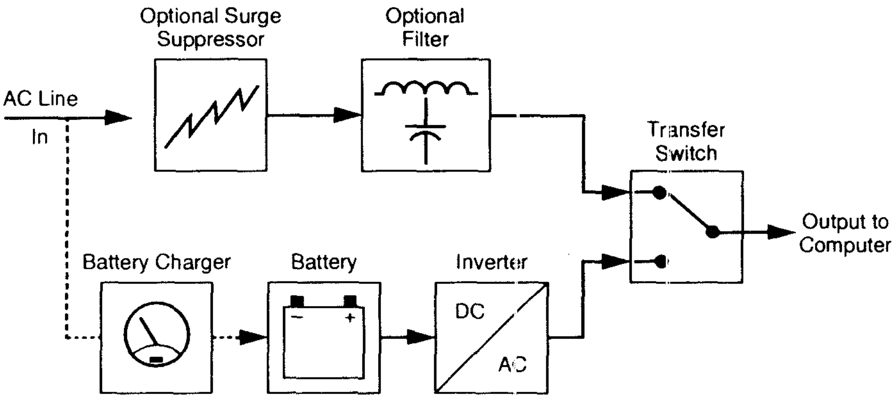

Exhibit IX-1-17. Block Diagram of Offline/Standby UPS

nominal line levels. The primary advantage to this technique is that it prevents the UPS batteries from becoming drained during extended brownout (low-voltage) conditions.

TriPort/Isolated UPS. A recently developed hybrid UPS design is called TriPort architecture. It is sometimes also referred to as a line interactive design. This type of hybrid UPS is shown in Exhibit IX-1-18. This design incorporates a transformer with three sets of windings. One winding receives power from the AC line. This winding energizes two secondary windings. One provides isolated and conditioned power to the computer; the other supplies the battery charger and inverter circuit path. When AC power fails, the charger/inverter winding becomes a primary winding and continues to energize the output winding that serves as the power source for the computer. The computer is always supplied with clean, conditioned AC, fully isolated from conducted line transients as well as from inverter output transients.

PUTTING SOLUTIONS IN PLACE

There are a variety of products and implementation options for providing a conditioned electrical environment for LANs. The network designer or manager

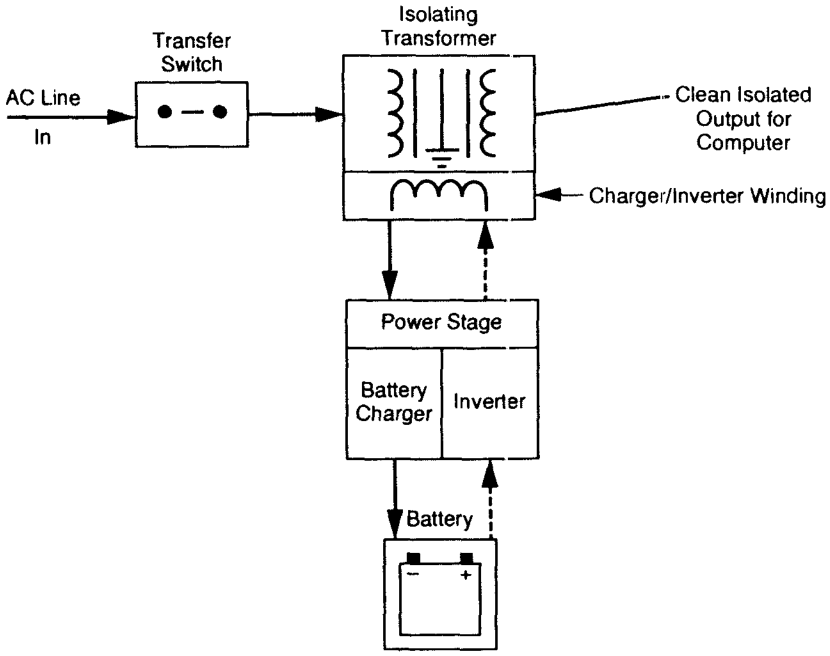

Exhibit IX-1-18. Block Diagram of TriPort/Isolated UPS

must consider the problem of adequately controlling transient energy without corrupting the grounding system (i.e., making it unsafe or inappropriate as a signal reference).

Distributed Computer Room Power Topology

Central versus Distributed. FIPS 94 recommends that the computer system be placed close to its power source to eliminate load-induced common mode noise. Noise is a high-frequency phenomenon, and the impedance of a circuit increases with the length of the circuit and with the frequency or edge-speed of the transient voltage induced on the circuit. The likelihood and magnitude of noise increases with distance from the circuit's source. Clean, dedicated, computer grade power from a central source is often corrupted by the time it reaches the distributed elements of a network.

The distributed nature of a network system favors a distributed approach to power conditioning—it simply works better to solve the problem locally. Other advantages include lower total cost, greater flexibility to reconfigure the system, and practicality. Even so, there are central options for UPS and power conditioning deployment. Central and distributed schemes are shown schematically in Exhibit IX-1-19.

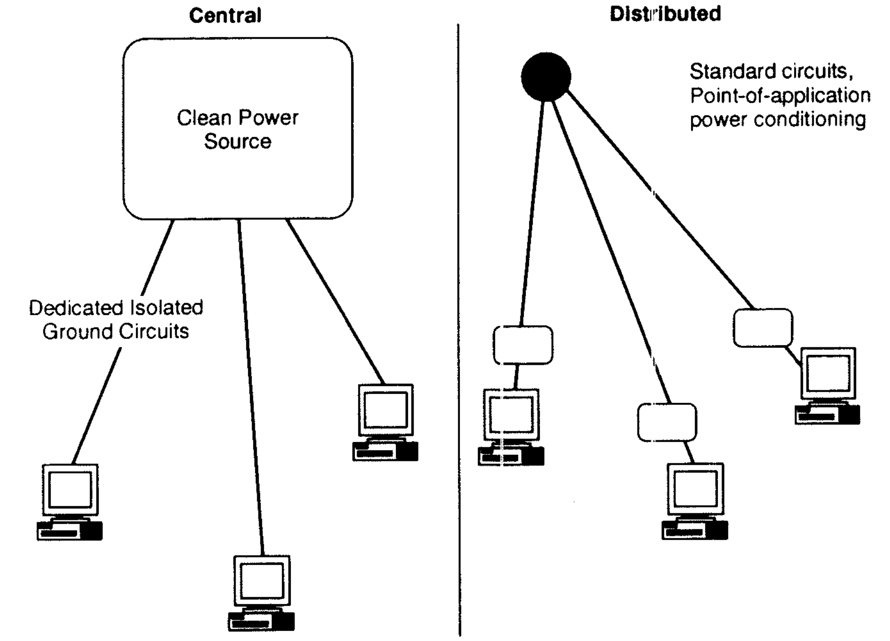
Central
Clean Power Source
Dedicated Isolated Ground Circuits
Distributed
Standard circuits, Point-of-application power conditioning

Exhibit IX-1-19. Central and Distributed Power Conditioning Schemes

UPS and Battery Backup. Some network administrators choose a large centrally located UPS system to power a network installation. The central approach offers more control, ease of management, and (seemingly) lower cost on a dollars-per-watt basis.

But a central UPS has to be large enough to accommodate future expansion and is often hard-wired into a dedicated electrical distribution system. Future moves and changes become an expensive proposition. A central UPS also presents a single failure point that will affect entire network segments when a fault occurs.

Expensive to install, maintain, and upgrade and with the performance limitations of central power conditioning, a central UPS is a poor choice for networks. An increasingly popular approach is to provide distributed battery backup capabilities specifically targeted to a network's critical resources. The distributed approach typically offers lower cost, greater flexibility, and better overall reliability.

Other UPS Considerations. A UPS should be rated to deliver twice the necessary backup time, for reasons described earlier in this chapter.

The level of interface sophistication that is required has to be evaluated. Interface capabilities range from simple signaling for automatic server shutdown to sophisticated SNMP interfaces with remote device control that can send service-required alerts or enable output power control to facilitate remote device rebooting (i.e., reinitiate a power-up sequence) or remote disabling of supported network devices.

Power Conditioning. Power conditioning systems offer the same deployment options as UPSs, that is, they can be central or distributed. A central approach requires installation of a system of special, dedicated power distribution circuits with insulated and isolated ground connections to deliver clean power throughout the network.

The emerging standard is point-of-application power conditioning. It offers economic advantages and better performance than central approaches.

Recommended Implementations

Exhibits IX-1-20 and IX-1-21 illustrate the best practical approaches for meeting the full range of power and grounding requirements of a distributed system.

A Departmental LAN. Exhibit IX-1-20 shows a LAN or LAN segment that is supplied entirely by a single low-voltage power distribution system within a single building. In an ideal installation small power conditioning systems (which use isolation through dedicated transformers) provide clean power and a solid reference ground for each device on the network. The otherwise pristine installation shown in Exhibit IX-1-20 can be corrupted.

Improper installation of data communications cables (i.e., running data cables in parallel with power circuits or near other sources of electromagnetic

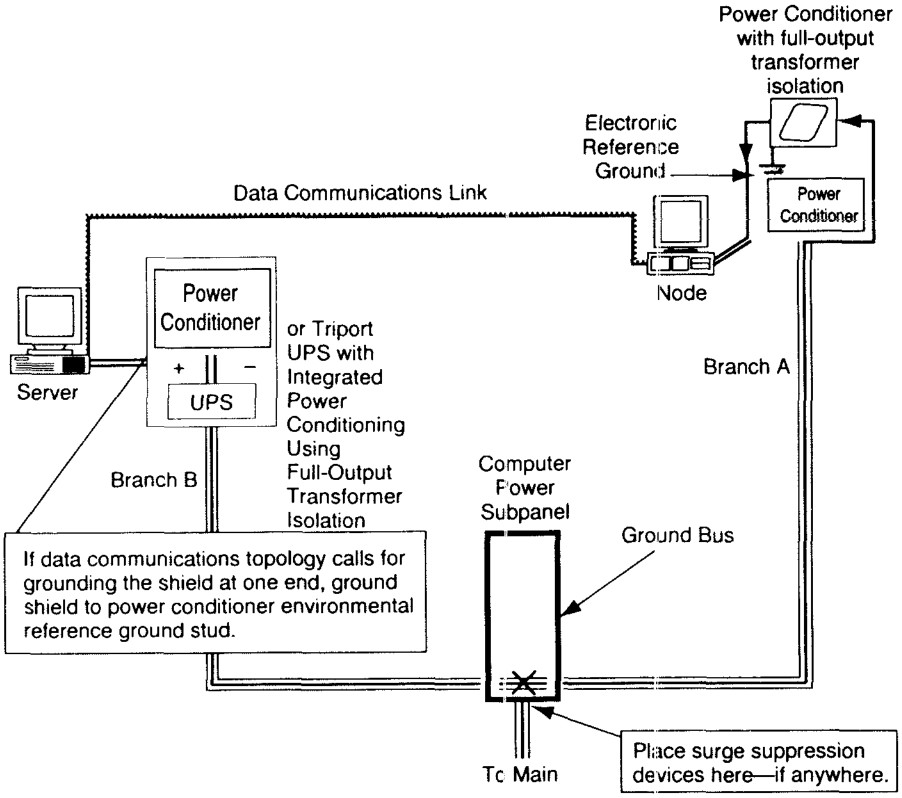
Power Conditioner with full-output transformer isolation
Electronic Reference Ground
Data Communications Link
Power Conditioner
Node
Power Conditioner
+
−
UPS
or Triport UPS with Integrated Power Conditioning Using Full-Output Transformer Isolation
Server
Branch A
Branch B
Computer Power Subpanel
Ground Bus
If data communications topology calls for grounding the shield at one end, ground shield to power conditioner environmental reference ground stud.
To Main
Place surge suppression devices here—if anywhere.

Exhibit IX-1-20. In-Building LAN with Single Power-Distribution System

interference such as fluorescent light fixtures, or large motors) can let electrical transients into the network through the data cables.

If surge suppression devices, which divert surge currents to ground, are placed out on the branch circuits, the integrity of the system reference grounds can be compromised.

Illegal neutral to ground connections at the panel (X) or somewhere within a branch circuit would divert to ground some portion of the normal return currents on the neutral conductor. Sometimes these illegal connections are made intentionally to solve computer performance problems caused by common mode noise. By mistake or otherwise, these connections can induce a repeating, low-frequency waveform on unbalanced communications cables, potentially corrupting the data stream. It is not uncommon to find illegal neutral-ground connections in newer as well as old buildings.

Extended LAN. Exhibit IX-1-21 shows a LAN that overlaps separate power distribution systems or extends beyond a single building.

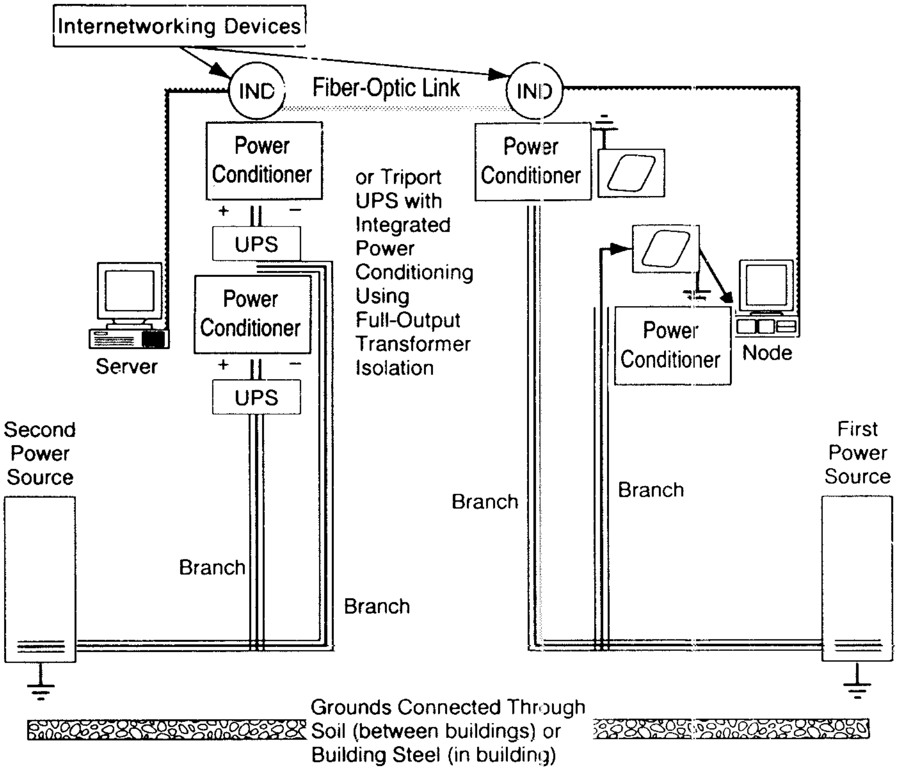

Exhibit IX-1-21. Multiple-Building or Multiple-Power Distribution System LAN

Whether the situation is two sources in one building or two separate buildings, the electrical grounds can be at different electrical potentials. The best solution to the possible ground skew and ground offset problem in these environments is to use a fiber-optic communications backbone between LAN segments. Fiber links do not make an electrical connection between different power distribution and grounding systems. The circuits supporting each local segment should have a common ground source if the building has been wired according to the National Electrical Code. If a fiber connection cannot be used and LAN segments in different buildings are being connected, then supplemental data line transient protection is recommended.

The advantage of the fiber-optic backbone is that each LAN segment can be treated as a departmental environment. Power conditioning systems should be installed on internetworking devices as well as on the individual network components in each LAN segment. For mission-critical networks, backup power is recommended not only for file and communications servers but also for critical path bridges, routers, and concentrators.

CONCLUSION

There are a lot of things to consider when trying to design a trusted, fault tolerant network. First, the value of the work that depends on the network environment has to be determined. A model that can be used to place a dollar value on network reliability has been provided. If a high value is placed on system uptime, then it pays to do things right.

Many network owners look to redundancy to provide added system fault tolerance. Redundancy is expensive—and for some systems it does make sense. But for all systems, attention to providing the correct electrical environment will provide a solid foundation on which to build reliable network systems.

As networks become increasingly complex and critical to an enterprise's business processes, they require managing. This is not a trivial task. Installing a network management system is an expensive and complex undertaking. Some networks warrant this level of investment in reliability. But for all systems, a network that operates reliably is inherently more easily managed than one that does not. Again, it pays to install a solid power foundation to improve network reliability and simplify network management.

This means assuring correct building electric wiring, selecting appropriate communications topologies, and careful placement of appropriate power conditioning devices throughout the network. The solution most like that used in data centers calls for the local placement of small, dedicated power conditioning systems that use low-impedance transformers as a key component. This provides a clean reference ground for each network device. It also protects the device and the network from the effects of transient electrical energy. Critical network devices such as servers, hubs, and routers should in most cases also be supported by battery backup when continued uptime or controlled shutdown is required.

In those networks that do warrant a comprehensive network management system, UPSs used in that network need to be included in the management scheme. A remotely manageable UPS has the capability to send alerts over the network when its batteries need replacing or when other conditions arise that require preventative or immediate maintenance. These advanced UPS management agents can also be used to restart locked devices or to gather information about power failure incidence in that segment of the network.

In the end, the decision as to how much time and money should be spent on preparing the LAN operating environment should be based on the costs of downtime for a particular network or network segment, the estimated mean time to "fault" occurrence, and site specific risk factors related to power quality that can adversely impact overall system fault incidence.

Total Protection Solutions for network computing and communications equipment

Founded in 1979, ONEAC Corporation is a company whose power conditioning and uninterruptible power systems are characterized by an exceptional ability to enhance fault tolerance and reliability for computing and communications equipment. At the heart of these systems is a technologically advanced design that makes them the ideal choice to protect electronic systems in sophisticated applications anywhere in the world.

UPSs with power conditioning

ONEAC design incorporates a low impedance, full output isolation transformer while other UPSs typically use transient suppression devices and filters. ONEAC's premium grade power conditioning effectively eliminates all fast edged transients, conducted noise and other contaminants that can damage or interfere with sensitive electronic systems. UPSs with filters let many of these contaminants through.

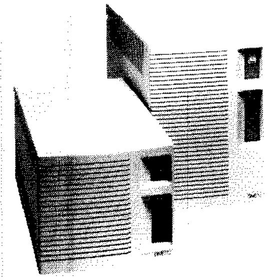

Power Conditioners

ONEAC power conditioners offer the same premium grade power as ONEAC UPSs for applications that do not require battery backup. Effectively eliminating all harmful power contaminants, ONEAC power conditioners deliver measurably better protection than surge suppressors and dramatically improve network reliability.

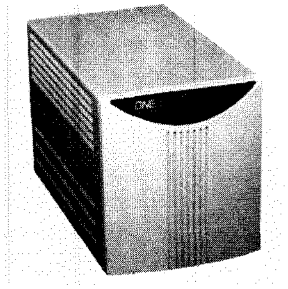

Communication Line Protectors

ONEAC telephone and data line protectors shield voice and data switching equipment from the devastating effects of overvoltage and over current on telecommunications lines. ONEAC's OnLine products incorporate three stages of protection: solid state, PTC's and ONEAC's patented Transient Filtering that eliminates damaging fast-edged events that pass through conventional protectors. A broad line of circuits and packages for customer premise, central and remote office applications are available.

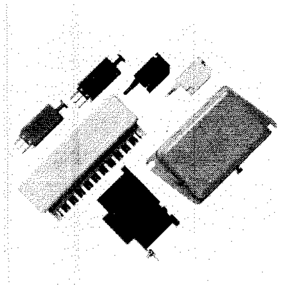

ONEAC®
A CHLORIDE POWER PROTECTION COMPANY

ONEAC is a UL/BSI
registered corporation —
Certification No. A2900

ISO 9001

ONEAC USA • 27944 N. Bradley Road, Libertyville, IL 60048-9700 • Tel 847-816-6000
ONEAC UK • 18 & 20 Blacklands Way, Abingdon Business Park, Abingdon, Oxfordshire 14 1DY, UK • Tel +44 (0) 1235 534721