ZONIT®
Structured Solutions

Solution Guide
99%+ Efficient Data Center Power Distribution

Zonit® Structured Solutions is redefining data center power delivery practices by bringing to market the world's smallest and most efficient zero-U automatic transfer switch based on unique patent(s) pending technology. The Zonit Micro Automatic Transfer Switch (µATS™) enables more efficient patent pending methods of data center power distribution, based on its unique technologies & features/benefits cost point. The Zonit µATS™ has five primary usage scenarios:

1. **Redundancy**
   Enables redundant power feeds for single power cord devices
2. **99.9% Efficient Power Distribution**
   Saves 3-15% of annual data center power costs
3. **Increase UPS Efficiency by 3-5%**
   Saves 3-5% of UPS power conversion losses
4. **Cost Effective N+1 Reliability**
   Eliminates the need for 2ndpower supplies in N+1 horizontal server tiers

Figure 1 - The Zonit µATS™

5. **Re-Location Flexibility**
   Allows easy "hot moves" to relocate powered-up equipment. Ideal for doing deployments that are temporary or will be moved at a later date.

Zonit µATS™ Detailed Usage Scenarios:

Redundancy: The µATS™ allows single power supply devices in the data center to be connected to redundant A-B power sources. This eliminates one of the primary causes of downtime, human error. Much downtime is caused by people plugging in one too many devices to a branch circuit, causing its circuit breaker to trip. The µATS™ is also very useful when used with single-power supply network distribution switches located in telecommunication closets. It allows them to be fed by filtered utility line power by default, with a UPS as the backup power source. This allows the UPS to be taken out of service for maintenance or testing, 9-5, M-F without network downtime.

Zonit Structured Solutions
1790 30th Street, Suite 140, Boulder, CO 80301 □ Phone: 720-266-0050 □ www.zonit.com □ info@zonit.com

99%+ Efficient Power Distribution: The µATS™ enables 99%+ efficient power distribution methods in the data center. It allows data center devices to be preferentially connected to filtered utility power on the A side and an Uninterruptible Power Supply (UPS) on the B side, which is only use when A is down. This saves 3-15% of data center power usage which can save millions of dollars in energy costs. The µATS™ is 99.75% efficient, using only 125mA in normal operation. This maximizes data center PUE.

Typical data centers use double conversion Uninterruptible Power Supply (UPS) units or much more recently, flywheel UPS devices. The best double conversion UPS units used in data centers have power efficiencies that vary as their load changes as shown in Fig. 2. They typically average 85-90% efficiency, with a maximum of ~97%. Flywheel UPS units average ~94% efficiency at typical load levels. This level of efficiency was acceptable when power costs were stable and relatively low

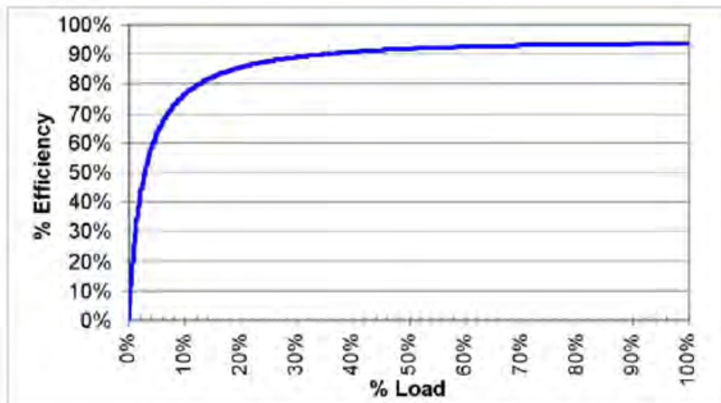

Figure 2 – UPS Efficiency

Power is now quickly changing from an inexpensive commodity to an expensive buy that has substantial economic & environmental costs.

Modern computer power supplies are almost all auto-ranging (accept 110-240V input) and all switched (Draw on the Alternating Current [AC] input power for just a short period of time and then convert this energy to Direct Current [DC], then repeat). Power supplies of this type are more resistant to power quality problems, because they only need to "drink" one gulp at a time, not continuously. If the input AC power voltage range is controlled within a known range, they will function very reliably. They do not require perfect input AC waveforms to work well. All that is necessary is that they receive sufficient energy in each "gulp" and that the input power is within the limits of their voltage range tolerance. This makes it possible to use a data center power distribution system that is much more efficient than a fully UPS supplied power system at a very reasonable capital expense.

The power system design shown in the left side of Figure 3 below controls voltage ranges very well. It has been in use for many years. The Zonit µATS™ enables new power distribution methods coupled with this power system design that are very efficient.

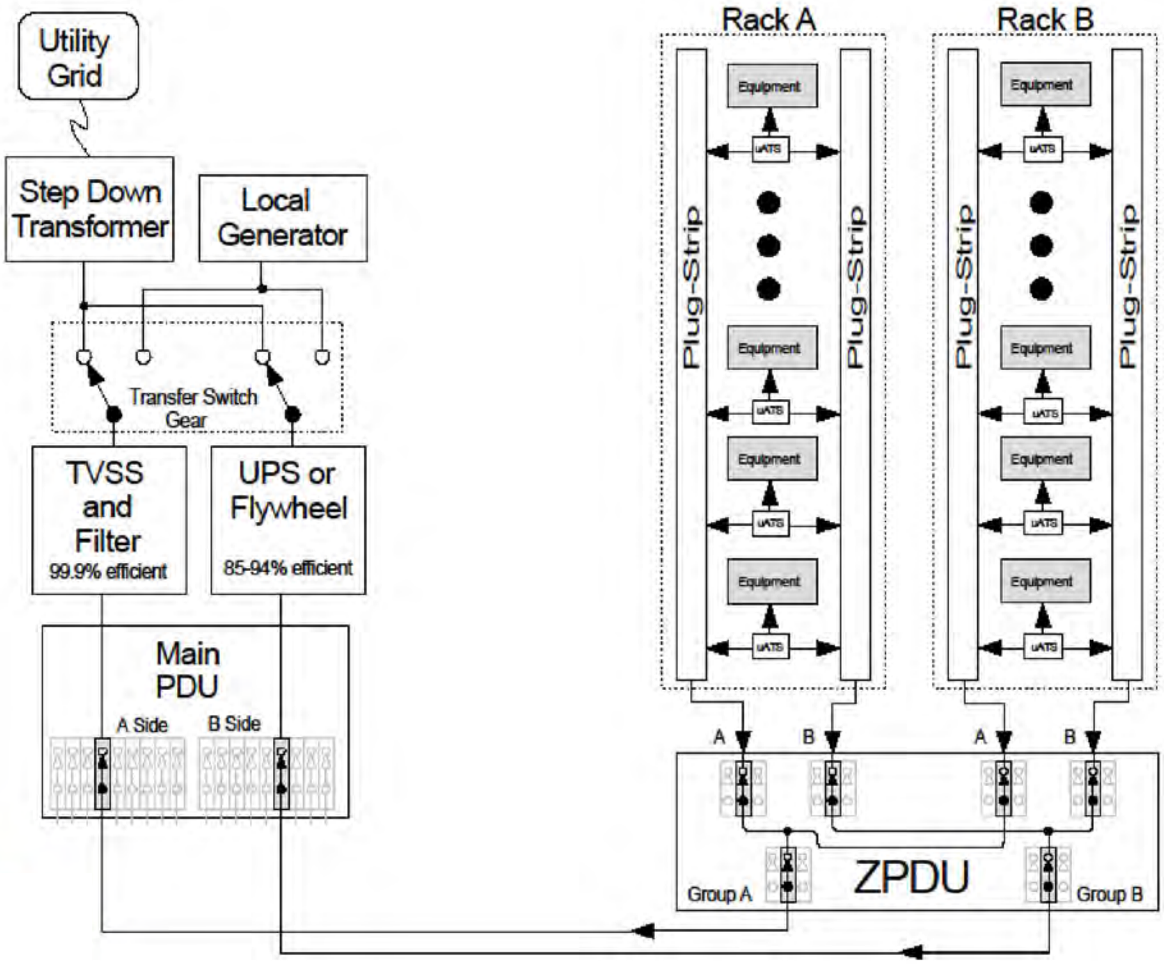

Figure 3 – 99%+ Efficient Data Center Power Distribution System

- ***Transient Voltage Surge Suppression (TVSS)***
  TVSS units are a mature technology that have been used in the power grid, data centers, hospitals, and factories since the 1960's. High-quality TVSS units are made by many of the same manufacturers (General Electric, Harris, MCG, etc.) as UPS units and many offer warranties against equipment damage, like UPS units. They are at least 99.9% efficient for the power they filter and are much simpler and lower cost than UPS units. TVSS products are fast acting filtering units that catch and damp transient voltage and current surges, through the use of passive filtering technology and large-amperage disconnect or shorting methods. They are very reliable and were originally developed for use in the power grid and mission critical facilities such as hospitals, etc. that were using electronic equipment that was not UPS connected. Properly specified units will reliably stop transient power events from going into the data center. They can self-monitor their health and report via SNMP.

- ***Zonit µATS™ units***

The Zonit µATS™ was designed to work properly with modern switched computer power supplies and it detects power loss and switches between the A-B power sources within the timing guidelines in the CBEMA Voltage vs. Time curve graph, shown right. The µATS™ was also designed to detect critical power quality problems, (over-voltage, under-voltage sag, and AC frequency shifts) that can affect computer power supplies and transfer to the B source before these problems can affect the operation of the power supply. This makes it practical to use filtered utility line + UPS power designs without risking downtime. The Zonit µATS™ will always use the A power source if it is available and of acceptable quality. This allows data center managers to always know and plan on what the load on the A and B sides of the power system will be, a requirement for power capacity management.

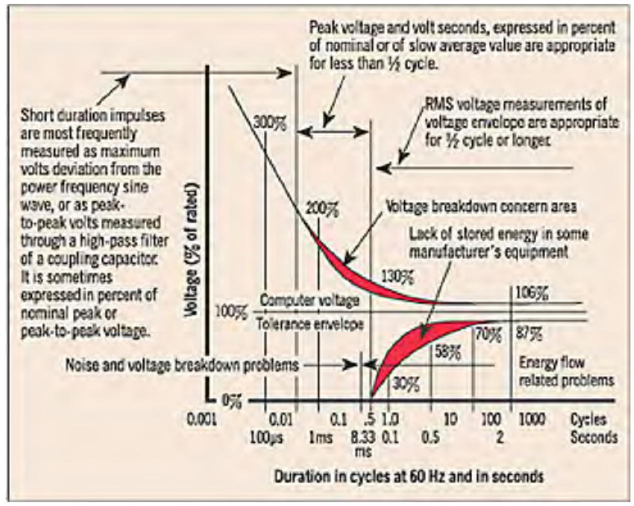

Figure 4 – CBEMA curve

The Zonit µATS™ is engineered to be deployed in parallel in large numbers, unlike many other automatic transfer switches. This is done by controlling the timing of the switches between the A and B power sources, which is done so that the generators throttle response under load has sufficient time to stabilize. If not, power source switches between the A side and the B side (generator side) could introduce an increasing amplitude resonance on the load of the generator, potentially affecting it.

The Zonit µATS™ also comes equipped with a Zonit patent-pending *Virtual Circuit Breaker*. The unit will sound an audible alarm if current levels are exceeding the 15A limit and if the over-current draw continues, the µATS™ will disconnect from the power source and illuminate a red LED. It can then be reset via a reset button on the unit. The internal interrupter fuse is only blown in the event of a catastrophic failure, such as a direct short-circuit of the device(s) plugged into the output of the µATS™.

The Zonit µATS™ is compatible with all standard server power supplies and eliminates the need to buy servers with two power supplies when using filtered utility line power, providing a huge cost savings. It can be reused with many generations of servers, which makes it a very long-lived, cost-effective, and green solution.

- ***Utility Company Step-Down Auto-Ranging Transformer***
  This unit is the on-site unit (typically sited outdoors next to the building) that reduces the grid feed voltage to a desired level, most commonly 480V. It can be specified to be auto-ranging, which means that it will limit the range of the voltage fed to the customer premises. This is done automatically, via a series of windings on the transformer core "taps" that are switched between automatically as the input voltage varies to insure that the output voltage stays in the desired range. This change can be done relatively inexpensively and provides redundant protection against utility line over-voltage problems in addition to the Zonit µATS™.

A good opportunity to upgrade the efficiency of the data center power system is when a power capacity upgrade is needed or new construction is done. If UPS power is used for both the A and B sides, the A UPS can be moved to the B side and an upgraded utility step-down transformer and a TVSS unit can be installed. This has relatively low impact on the power distribution system in the data center from the PDU to the rack. If uniform A-B power distribution to all racks is not in place, the Zonit Power Distribution System (zPDS) should be considered. It is designed for uniform and flexible A-B power distribution that supports all of the changes in power distribution needs that occur at the rack during the lifecycle of the data center. See www.zonit.biz to learn about the zPDS.

Increase UPS Efficiency by 3-5%: Normal practice in data centers with matched A-B UPS units is to load them approximately equally, neither more than 50% of capacity, so that either can fail over to the other. Unfortunately, this means that both UPS units will be run in a less efficient configuration than is possible. The UPS load efficiency curve shown in Figure 2 demonstrates this. A UPS is most efficient when it is run as close to 100% load as possible. When the electrical load of the data center in shared between two UPS units, they will each typically run in the 30-40% load range. This is inefficient.

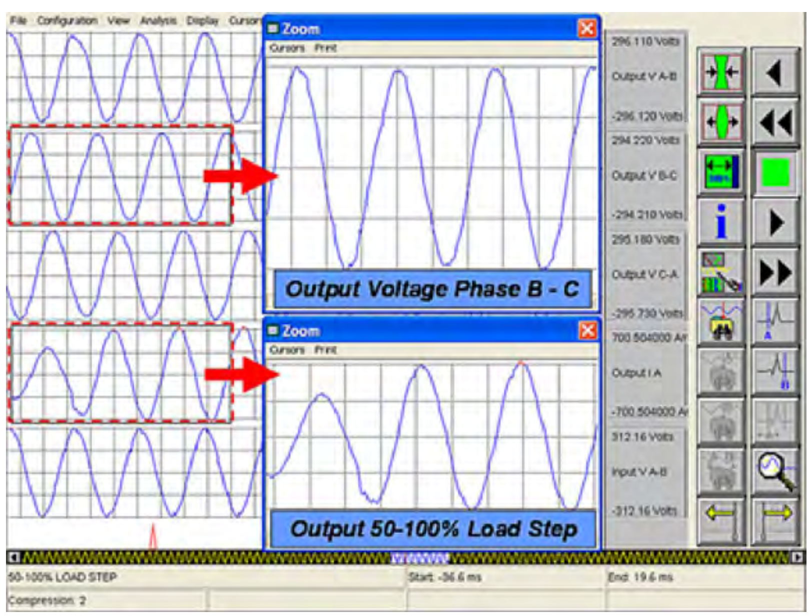

Figure 5 – UPS Load Step Test

The Zonit µATS™ is designed to always select the A power source if it is available and of sufficient quality. This feature allows all of the load in the data center to be placed on the A side UPS in normal operation. This is done by connecting devices via µATS™ units to the A and B power sources. A µATS™ can be used for every single power supply device and the B or N+1 power supply of devices that have dual or multiple power supplies. This "***Load Shift***" Method puts the entire electrical load on the A UPS which then runs more efficiently. The B UPS runs with little or no load in normal operation. It must support a "50-100% load step function" which means that it can go from 0% to full load without any problems, but most modern UPS units sold during the last 10 years can do this. An example of a UPS load step test is shown in Figure 5. Note that the output voltage remains normal.

Cost Effective N+1 Environments: Data Centers are the fastest growing electrical energy consumers in the USA today. Large server farms are the drivers in this growth. The huge numbers of servers deployed in these data centers currently are almost all "commodity" Intel X86 architecture compatible CPU's. This is what powers most of the big server farms running large Web sites, cloud computing running VMWare™ or other virtualized solutions, and high performance computing (HPC) environments. It is the most competitive and commoditized server market segment and offers the best server "bang-for-the-buck". This is why it is chosen for these roles.

There are several reasons to put multiple (dual or N+1 are the most common configurations) power supplies into these servers. The first is to eliminate a single point of failure through redundancy. However, modern power supplies are very reliable with Mean Time Between Failure (MTBF) values of about 100,000 hours = 11.2 years, well beyond the typical service life of the equipment. The second reason that multiple power supplies are used is to allow connection to more than one branch circuit. This is the most common point of failure for power distribution, as discussed earlier. Also, having dual power connections makes power system maintenance much easier, by allowing one power source to be shut down without affecting end user EDP equipment.

However, putting multiple power supplies in Electronic Data Processing (EDP) equipment has costs. The additional power supply(s) are an added expense. They are almost always specific to each generation of equipment, and therefore must be replaced in each new generation of equipment, which for servers can be as short as three years in some organizations.

Power supplies also have a loss factor, they are not 100% efficient and the least expensive way to make a power supply is to design it to run most efficiently at a given load range typically +- 20% of the optimum expected load. Power supplies have an efficiency curve that is similar to UPS units, such as was shown in Figure 2.

This presents another issue. If two power supplies share the server load, they can be less efficient than optimum. Depending on the specification of the power supply, was it optimized to run best in a single or dual configuration? The product manager for the server manufacturer may sell that server in two configurations, with one or two power supplies. In that case, he may chose to specify only one power supply model, since to stock, sell and service two models is more expensive than one model of power supply. The power supply model he chooses usually will be the one that matches the most common configuration sold. This trades capital expense (the server manufacturer can sell the server at a lower initial price point) vs. operational expense (it uses more power to run) . This is because with two AC to DC power supplies, the DC output bus will almost always be a common shared passive bus in the class of commodity server that is most often used in large scale deployments. Adding power source switching to this class of server to gain back efficiency (only one power supply at a time takes the load) but is generally too expensive for the market being served. It also adds another potential point of failure that costs to make redundant if needed for greater reliability. To use single power supply servers, their input power must either be protected via a UPS (which results in power efficiency losses as discussed earlier) or their input power must be auto-switched from a primary to a backup source. We can now discuss the advantages of auto-switching at the device level using the Zonit µATS™ vs. auto-switching at other points in the power distribution topology.

The µATS™ makes switching at the device or near device level both possible and desirable. The advantages are easy to understand. A population of highly reliable ATS units at the device level produces much higher per device power reliability levels than an ATS that switches a branch circuit or an entire panelboard can, due to the statistics involved. The chances of all of the µATS™ units in a given population failing at the same time and therefore affecting all of the auto-switched servers is infinitesimal vs. the chance of an ATS failing that is closer to the root of the power distribution topology. Consider the following example.

> 1 Panelboard ATS with an MTBF of 200,000 hours
> 1/200000 = 5.0e-06 chance of failure in any given hour.
> Note: This would be a very expensive unit w/ this MTBF # and while reliable, it **is** a single point of failure for **all** of the devices that it powers.
>
> 200 µATS™ units each unit with an MTBF of 200,000 hours
> 1/200000 = 0.005% chance of failure per unit in any given hour and the chance of 200 units failing simultaneously
> = 200,000 raised to the 200th power divided by 1
> =6.2230152778611417071e-1061.

This is essentially zero chance of simultaneous failure and is over 1000 orders of magnitude better than a single ATS. This is a key advantage of the patent pending data center power distribution methods we are describing that the Zonit µATS™ makes possible. Reliability is so very important to data center operators, especially for companies that measure their downtime in hundreds or thousands or millions of dollars per hour. It is hard to over-emphasize this point.

Re-Location Flexibility: Data Center moves and deployments are not always planned or go to plan. Sometimes you have to improvise. The µATS™ allows data center operators to easily deploy equipment and then relocate it later without powering it down. The fact that the µATS™ moves *with* the device allows it to be unplugged on one source, plugged into an extension cord, then unplugged on the other source, then moved. We call this a "hot walk". The fact that Ethernet based TCP/IP connections will tolerate being unplugged and then re-connected (all that is needed is to reconnect to the same VLAN, even if on a different switch port) makes "hot walks" practical.

More Information

Please contact Zonit at info@zonit.biz for more information on implementing these more efficient power distribution system designs and realizing the economic and technical advantages they deliver.