THE MAGAZINE OF 7x24 EXCHANGE INTERNATIONAL

NewsLink

FALL 07

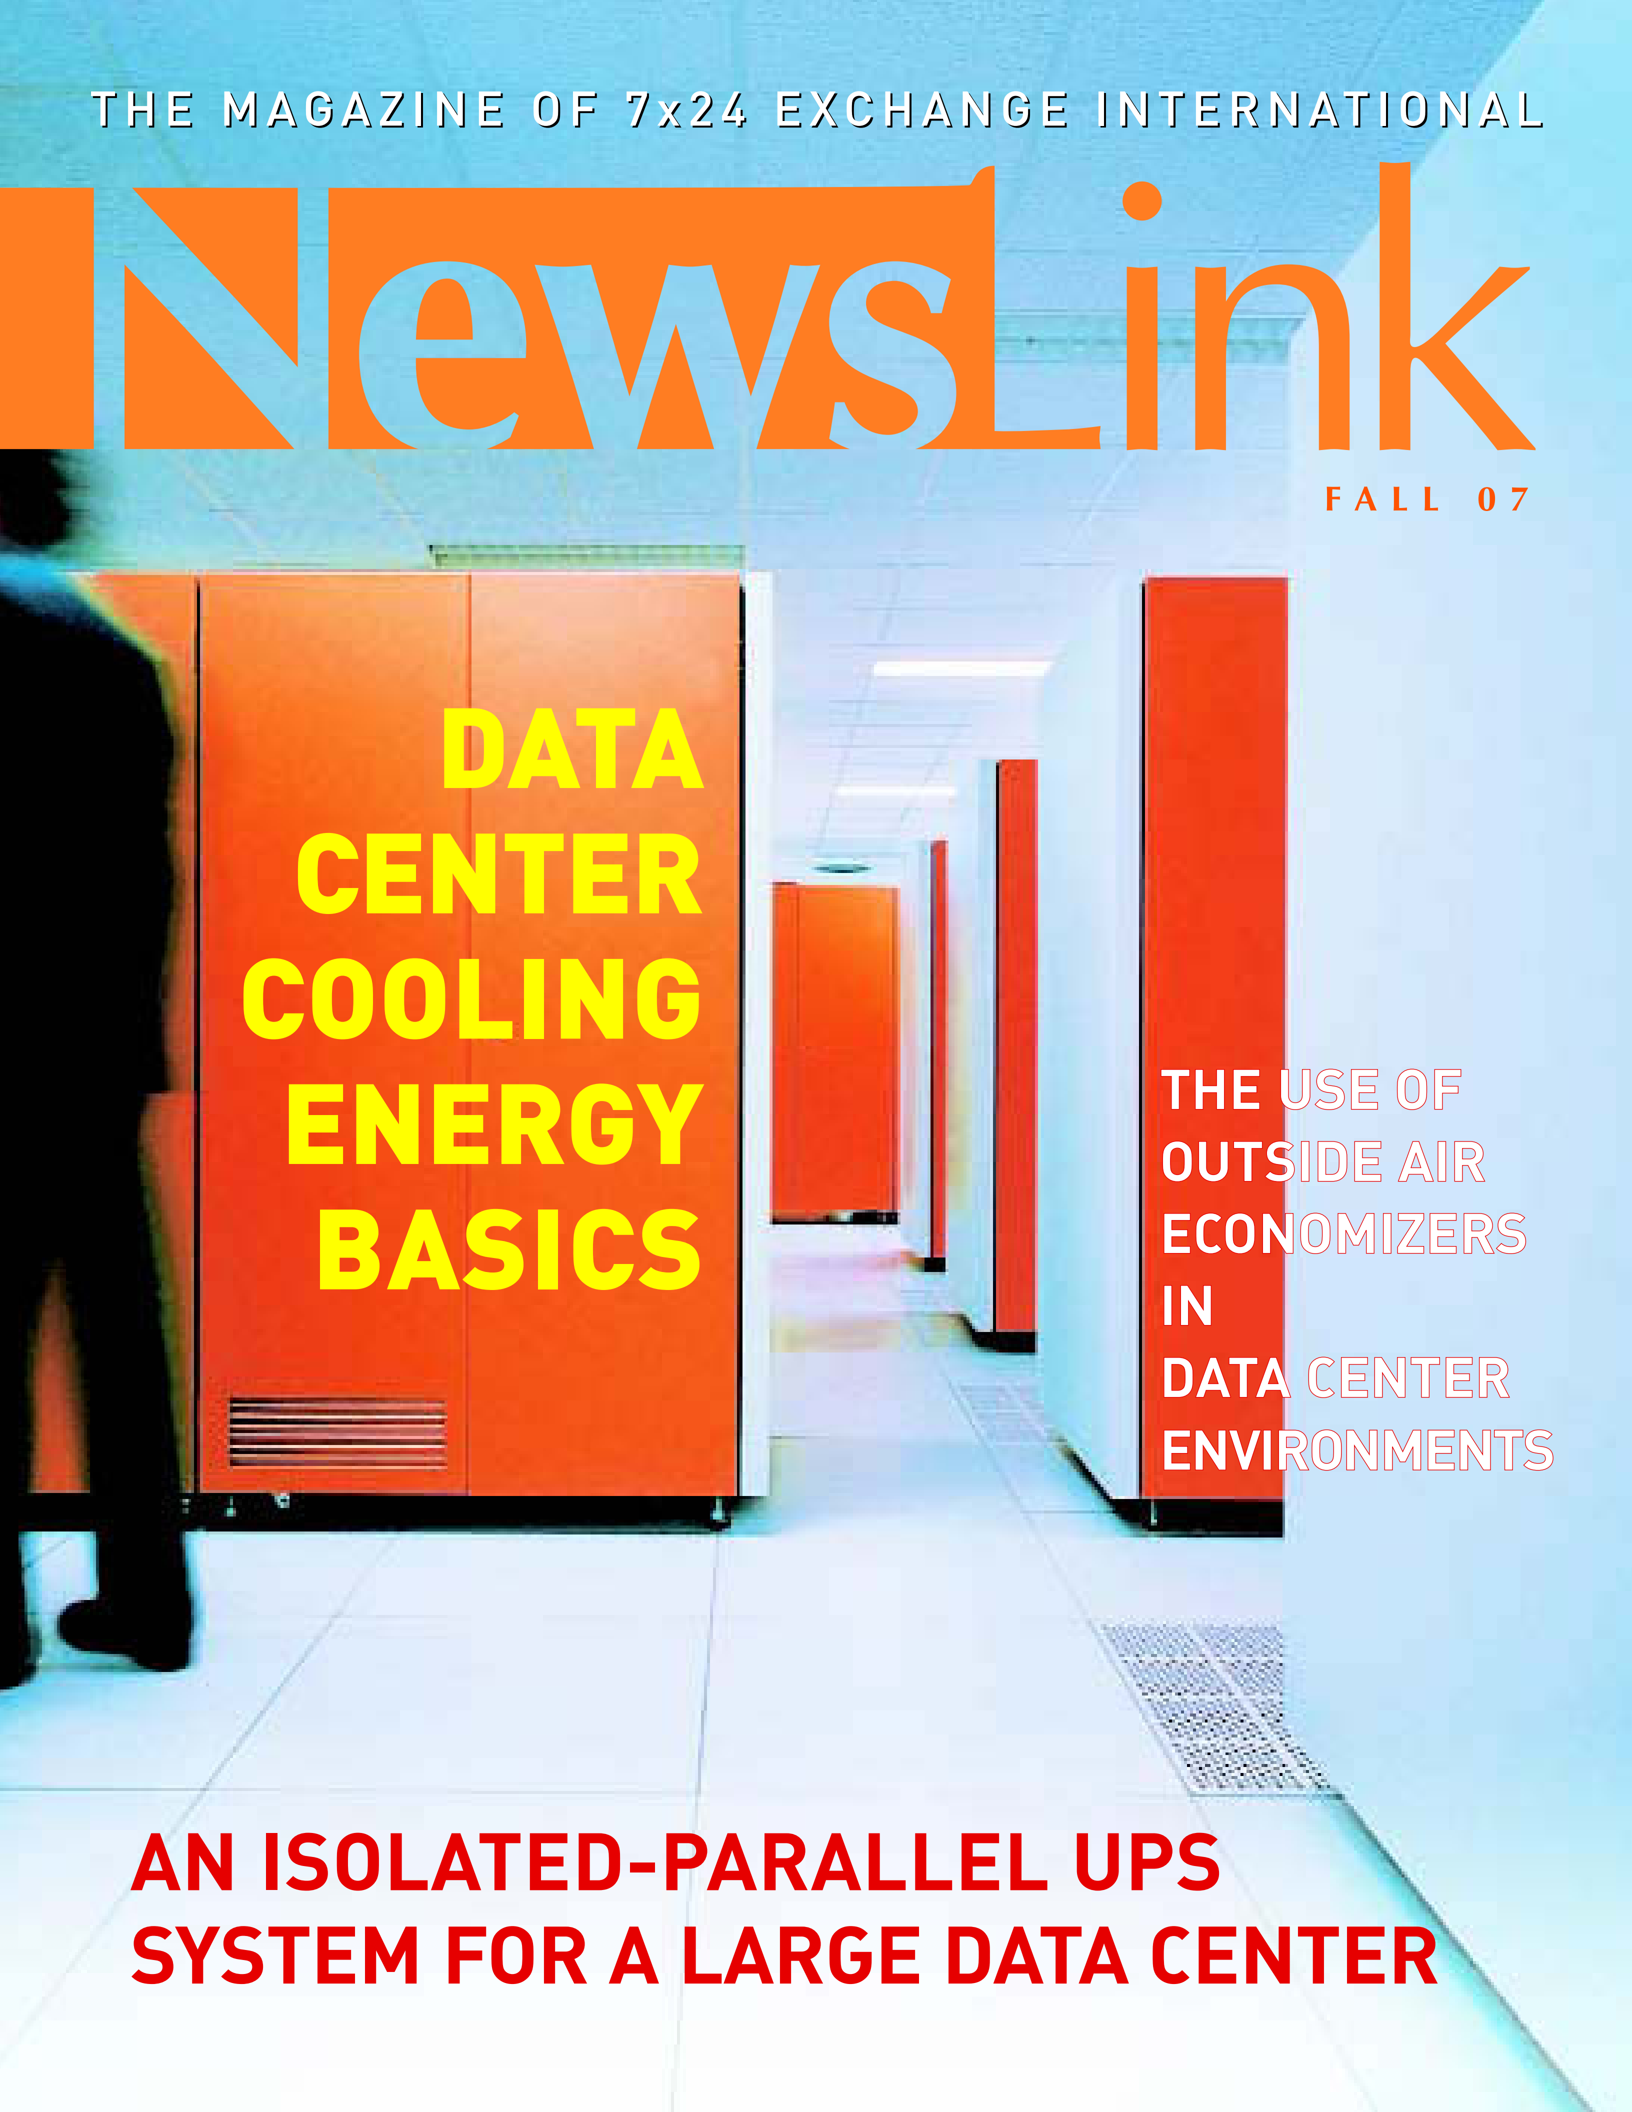

DATA CENTER COOLING ENERGY BASICS

THE USE OF OUTSIDE AIR ECONOMIZERS IN DATA CENTER ENVIRONMENTS

AN ISOLATED-PARALLEL UPS SYSTEM FOR A LARGE DATA CENTER

SOMEONE HAS TO LEAD

ADVANCED

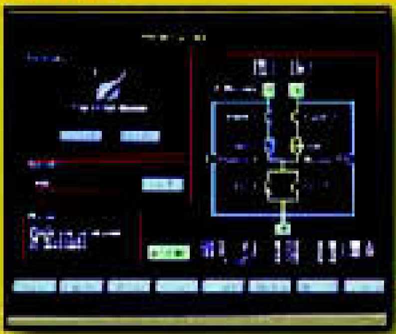

STS Communication

- New Wavestar communications package
- Triple redundant logic and power supplies
- Redundant LCD screen with Waveform capture
- Industry's first biometric security system

NEW KITS & VOLTAGE VERSION

Branch Circuit Monitoring System

- Kits for any existing distribution cabinets
- Track, report and manage current, voltage and power measures to server level
- Manage current voltage and power in server cabinets
- Communicate through modbus or SNMP protocol and/or local monitor

IMPROVED

PDU Performance

- Re-engineered rigidized doors
- Vision panels and improved locks
- Hinged interior panels
- Fits BCMS with no alterations
- Flexible custom designs

Power Distribution, Inc.
4200 Oakleys Court
Richmond, VA 23223
800.225.4838
804.737.1703 fax
e-mail: sales@pdicorp.com
web site: www.pdicorp.com

PDI
creating the perfect wave

Contents

7x24 Exchange International

The end-to-end reliability forum.

www.7x24exchange.org

DIRECTORS & OFFICERS

Chairman of the Board
Robert J. Cassiliano
Business Information Services, Inc.
(201) 672-0630

President & Chapter Representative
William Leedecke
Vanguard
(610) 669-5707

Vice President
Roy L. Chapman
American Express
(602) 766-6502

Vendor Representative
Douglas H. Sandberg
ASCO Power Technologies
(973) 966-2079

Director
David Schirmacher
Goldman Sachs & Co.
(212) 357-7482

Administrative Director
Kathleen A. Dolci
(646) 486-3818 x103
kathleen@7x24exchange.org

Membership & Education
Tara Oehlmann, Ed.M.
(646) 486-3818 x104
tara@7x24exchange.org

Conferences
Brandon A. Dolci, CMP
(646) 486-3818 x108
brandon@7x24exchange.org

322 Eighth Avenue, Suite 501
New York, NY 10001
(646) 486-3818

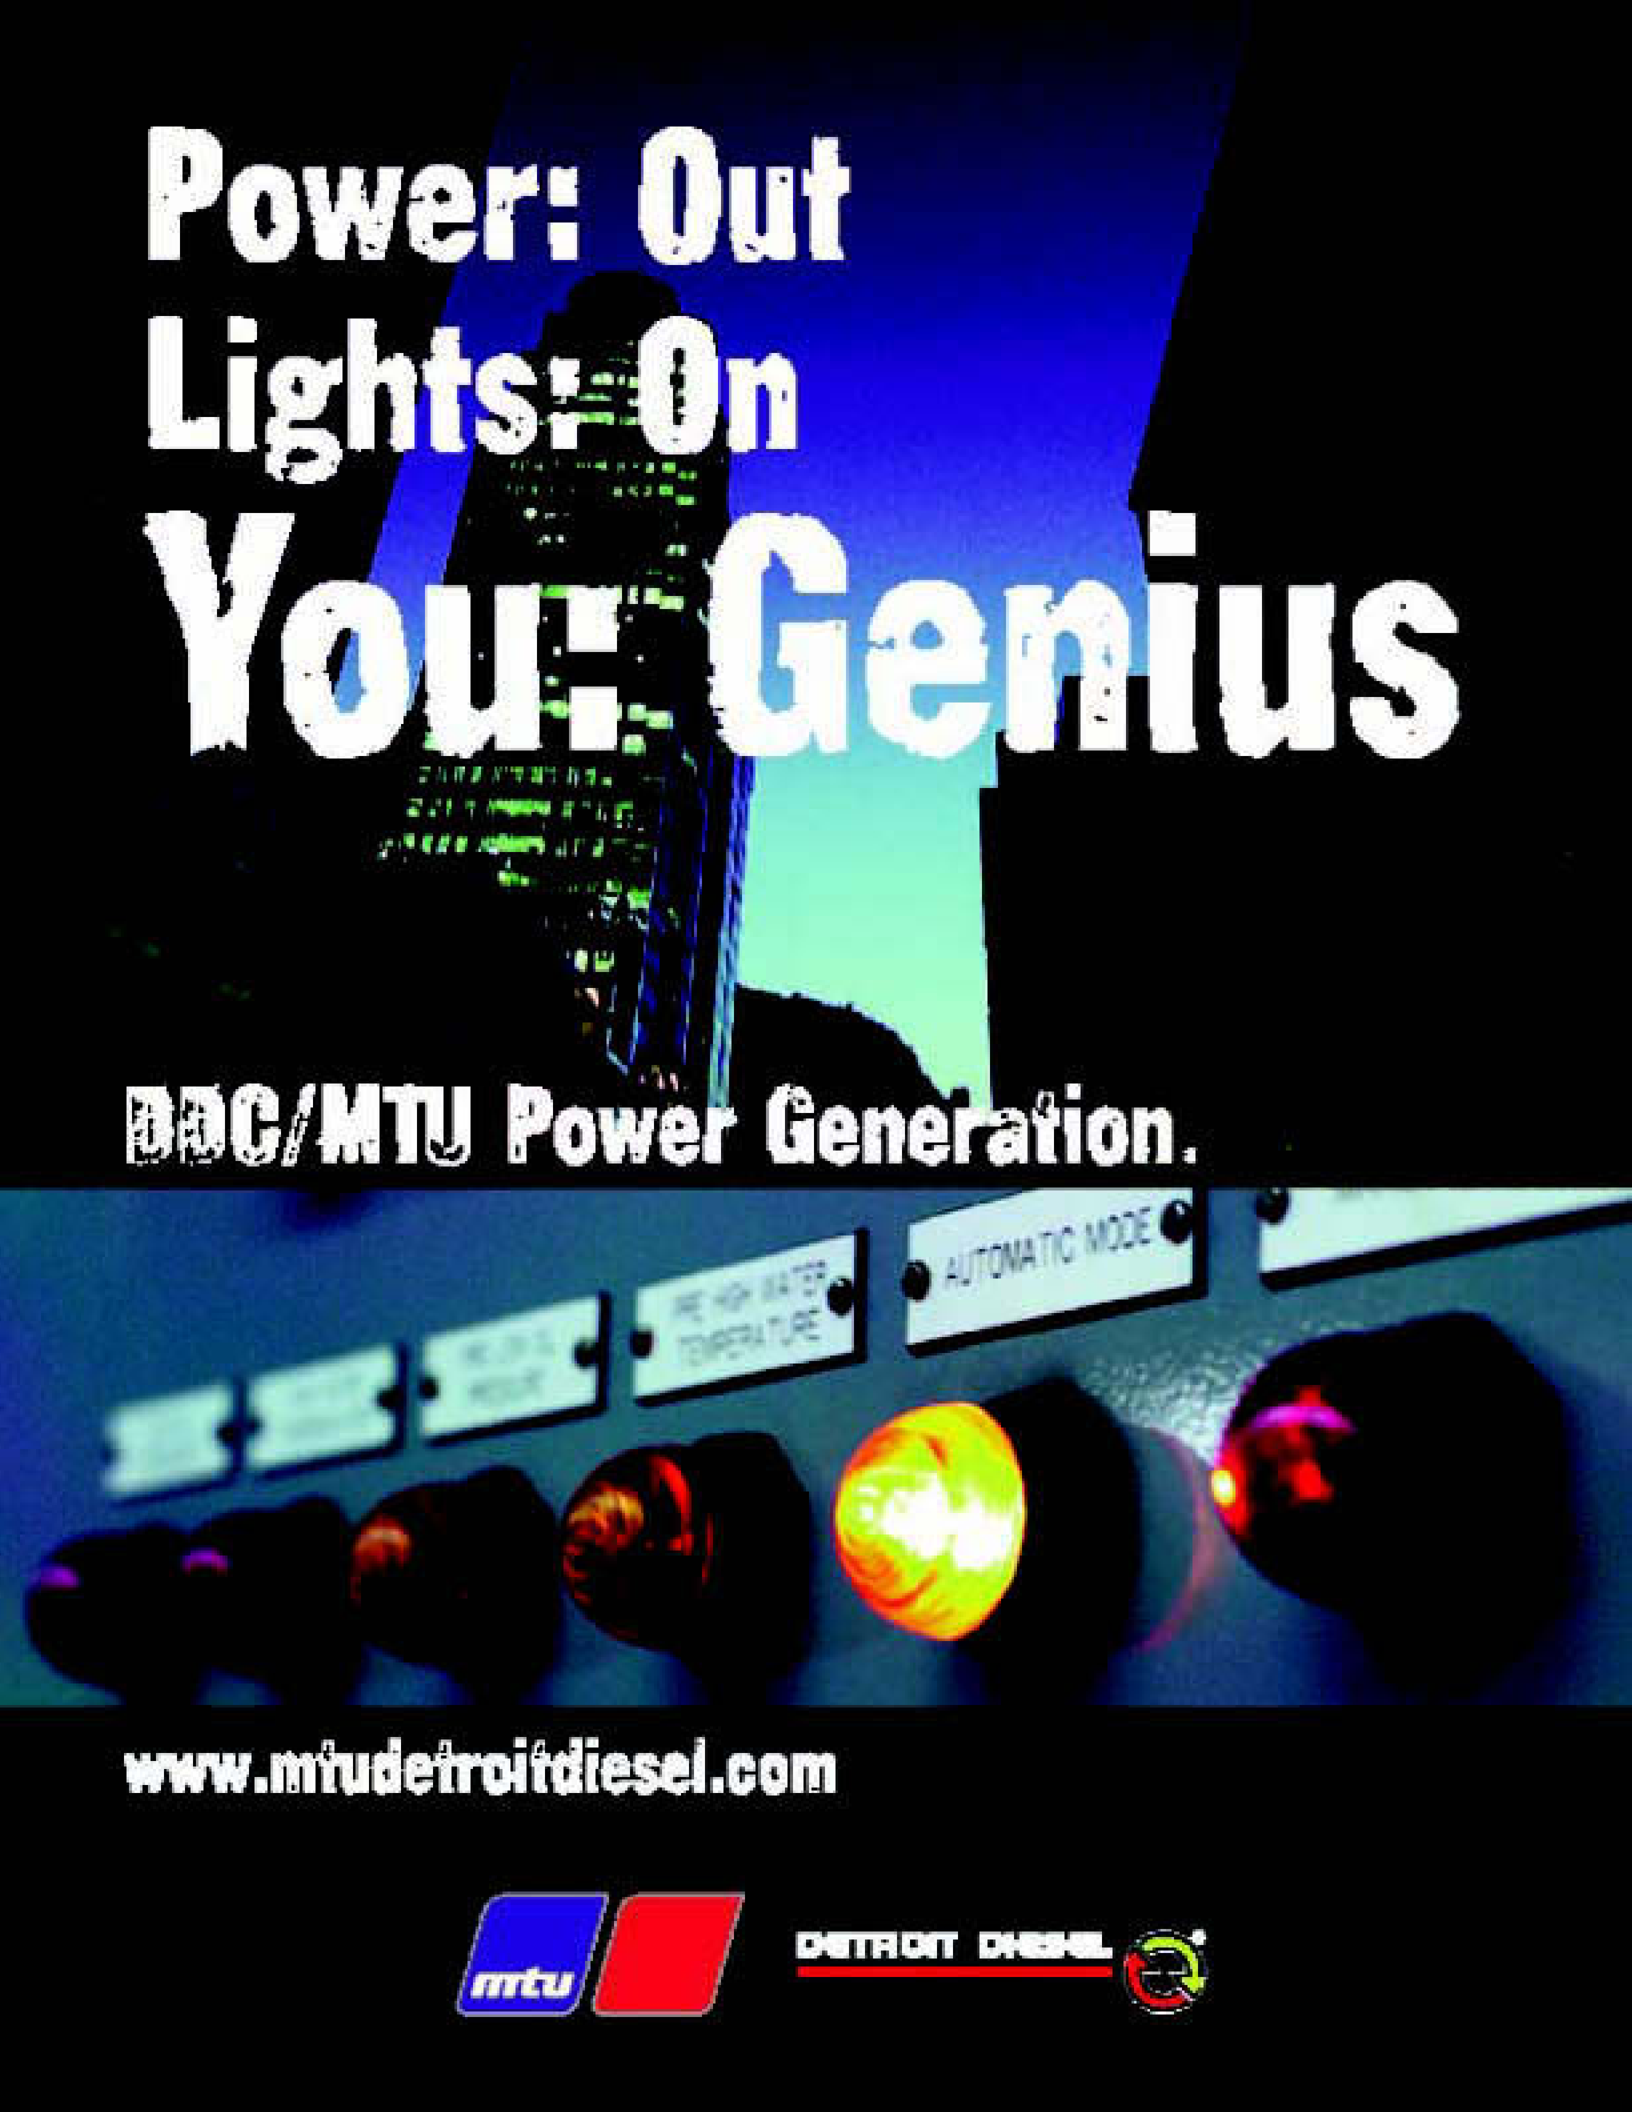
Power: Out
Lights: On
You: Genius
DDC/MTU Power Generation.
PRE HIGH WATER TEMPERATURE
AUTOMATIC MODE
www.mtudetroitdiesel.com
mtu
DETROIT DIESEL

Battery Diagnostic System
BDS-40

Trust Your **Batteries.**

- Preventing battery failure
- Avoiding costly downtime
- Reducing maintenance cost
- Optimizing useful battery life
- Increasing safety

Utilities

Data Centers

Telecom

Mission Critical

TOP PRODUCTS · 2006 · Electrical

Alber
Battery Test Experts

AN ISOLATED-PARALLEL UPS SYSTEM FOR A LARGE DATA CENTER

by Mike Mosman, PE

DuPont Fabros, a leading owner, developer, operator and manager of wholesale data centers, maintains several large data centers which it provides to tenants under long-term leases. In 2006 they were planning a new facility at their data center campus in Ashburn, Virginia. DuPont Fabros' data center design standards require the critical power delivery systems in their facilities be flexible enough to match individual tenant load requirements, while providing fault isolation that would limit internal disturbances to a single tenant. In addition to the technical requirements for the critical power systems, the overall facility infrastructure must be designed to be economical for both capital and recurring costs.

Existing DuPont Fabros data centers in their campus have rotary flywheel type Diesel-UPS modules arranged in isolated-redundant configurations. These UPS modules are applied in groups of 8 to 12 Diesel-UPS modules with two modules designated as redundant units. These systems are designed with static transfer switches to transfer critical load from a failed primary module to a redundant module. In order to achieve the fault tolerance requirements of the DuPont Fabros design standards, UPS modules are dedicated to specific tenant areas for fault isolation. Rotary Diesel-UPS technology as implemented in these facilities provides an efficient and resilient system that is economical and compact compared to an electrical system incorporating large system-redundant static UPS systems with batteries. For DuPont Fabros the iso-redundant systems have provided a high level of resilience while reducing capital and maintenance costs.

LESSONS LEARNED

In operating their existing facilities, DuPont Fabros identified certain aspects of the isolated-redundant UPS systems they wished to improve upon. For instance, the reliance on static transfer switches to provide the system redundancy makes isolated-redundant systems susceptible to the problems inherent in the coordination of UPS output characteristics and STS input tolerances. Another issue identified is that excess capacity in the isolated modules can not be easily used for delivery of additional power to other tenants. Likewise, the total redundancy in the system is always limited to the number of redundant modules that are dedicated to the alternate source of the static transfer switches.

In order to serve their growing customer base, DuPont Fabros made the commitment to build their new facility to an unprecedented scale. Their engineer, CCG Facilities Integration Incorporated, programmed a 300,000 square foot building with 170,300 square feet of raised floor divided into 16 computer rooms, each of which can be further partitioned into one, two or four spaces. Allowing for redundant capacities, a 40 MW UPS system, a 17,000 ton chiller plant, and an 80 MVA electrical service with 72 MW of backup diesel engine-generators was programmed to meet the maximum requirements of the tenants that would occupy these spaces.

The programming and design phases of this new facility gave its owners an opportunity to review the their previous data centers, and consider ways to improve future facilities. Instead of an iso-redundant UPS configuration, DuPont Fabros agreed to implement an innovative integration of parallel-redundant and isolated-redundant UPS configurations developed by CCG called an Isolated-Parallel (IP) UPS configuration. An IP system had never been used in the US, but it held the promise of providing some additional benefits over their prior critical power delivery systems.

IP BASICS

An Isolated-Parallel UPS system is made up of a group of Diesel-UPS modules which share a common paralleling bus, or IP bus. Each module within an IP system has a utility input backed up by a single diesel engine-generator, and a reactive choke and synchronous alternator combination for power conditioning. The alternator has a coupled flywheel for energy storage to be used as the ride-through power source between utility outage and diesel engine-generator takeover. Each module has an output bus for its dedicated critical loads. Each critical output bus is also connected to the IP bus through another reactive choke (IP choke) which allows real power to flow, but restricts high fault kVA flow due to the reactive nature of fault currents. The utility input and diesel engine-generator associated with each module can also be used supply essential mechanical equipment that does not need ride-through power.

Two IP systems, each with 16 modules were designed by CCG for DuPont Fabros' new facility. The Diesel-UPS modules are each designed to serve a combination of critical and essential loads, while providing continuous backup to the utility power. Electrical service is brought into the building via 16 pad-mounted transformers rated 5 MVA each. A service transformer secondary voltage of 600 volts is used to reduce the size of feeders and breakers in the facility. The 2250 kW engine-generators, which are part of the Diesel-UPS system and installed within the building, also operate at 600 volts. As an added reliability enhancement the entire 600 volt system is high-resistance grounded to allow the system to operate unaffected by a line-to-ground fault anywhere on the system.

The critical power distribution panels are capable of handling up to 200% of the their normal load, enabling DuPont Fabros to provide "A" and "B" distribution points within their tenant's computer rooms and not be concerned about an uneven A/B split of critical power demand. While Static Transfer Switches are not required to produce the redundancy between "A" and "B" UPS modules, they may be used if required for a specific tenant application. An even greater advantage to the owner and landlord is the ability to handle an "overload" by one of his tenants. As long as the total load in the building remains below the IP system's redundant load point, any computer room may exceed its allotted power level up to the air-conditioning capacity provided within the space. Since load is leveled across all connected modules in the IP system, the excess requirement of the overloaded computer room is automatically made up by those computer rooms which are yet below their load allotment.

IMPLEMENTING A NEW CONCEPT

Piller Power Systems was chosen as the supplier for the Diesel-UPS flywheel modules, engine-generators and associated switchboards. As part of their scope of work Piller commissioned a computer model of a 16-module IP Diesel-UPS system from DBR Consult, an independent consulting firm in Germany. The computer model was used to predict the load flow and transient stability of 16 generators and 16 flywheel/alternators operating on a common parallel bus through 16 chokes, and serving 16 independent steady state and transient loads of varying amount ranging from 0 to 200%. The results of Piller's modeling indicated very stable operation and controlled load flow within the normal ratings of system components.

To implement CCG's IP design concept, Piller's systems engineers and project management specified the required parameters of system components, and coordinated their delivery and testing. They designed the control systems necessary for the integrated operation of the UPS modules, engine-generators, and switchboards. Unique and innovative schemes were developed to provide independence between Diesel-UPS modules while controlling load flow especially when modules are in different modes of operation.

In order to verify the IP system, CCG and DuPont Fabros required Piller to assemble four Diesel-UPS modules, four engine-generators and four sets of module switchboards into a 4-module IP system within their New York testing facility. A comprehensive factory test was conducted at rated load and witnessed by DuPont Fabros and CCG. The test results agreed very closely with the computer model. A further test was conducted on a prototype of the IP choke at a high-current test lab that verified the reactive chokes on which the IP system depends for its isolation would not saturate and lose impedance under a fault condition. The first Isolated-Parallel UPS system had a promising start.

A DESIGN REALIZED

Construction of the new facility, designated Ashburn Corporate Center #4, or ACC4, began in June, 2006. DuPont Fabros engaged Holder Construction as the general contractor with Dynalectric as the electrical subcontractor. By the end of June, 2007 the first 16-module system was fully commissioned. The commissioning effort was specifically designed to demonstrate the special features of an IP configuration that, combined, differentiate it from iso-redundant or standard parallel-redundant UPS configurations. These are:

- **Load sharing and equalization among all modules.** When briefly on flywheel as the critical power source, or while on diesel engine-generators as the backup power source, the critical loads are evenly distributed across all units connected to the IP bus, even when critical loads vary from 0 to 200% on any unit's output bus.

- **Fault isolation to a single output bus.** A fault anywhere on the system will at most affect only one Diesel-UPS module's critical output bus. The voltage on all other critical busses remain within the ITI (CBEMA) curve. Fault currents everywhere in the system are limited by the IP chokes to less than 100,000 amps rms, even with 16 diesel generators connected.

- **Bypass to the common IP bus.** If a UPS module fails the output is not dropped, but rather is sustained by power flow through its IP choke from the IP bus. An automatic or manual initiation of the bypass circuit for that unit will transfer its critical load to the IP bus directly where it will be shared equally by all other modules connected to the IP bus.

- **Distributed redundancy.** There is no designated redundant module. All modules share the total critical load, and the amount of redundancy is merely the difference between the total load and the total capacity.

- **Independent modules.** While the IP system has a central Master Controller, it is not required to control load flow, or respond to outages and failures. All modules operate with completely independent individual controls, and no single point of failure was discovered anywhere in the system.

The results of site testing done during the commissioning process once again agreed with the predictions of the computer model. The load flow analysis was exactly comparable to real measured values. Furthermore, the entire system was very stable under transient perturbations. Even so, the full system commissioning tests revealed the need for some minor control component adjustments. After these changes the IP system's specified robustness and fault tolerance was fully realized. The new IP system has reliably and repeatedly demonstrated its ability to start and parallel 16 engine generators on the common IP bus well within the spindown time of the UPS flywheels at full building load. It is also able to seamlessly handle engine start and synchronizing failures. By the end of the commissioning effort the Isolated-Parallel UPS system configuration was a proven concept.

The inclusion of the IP design in its portfolio supports DuPont Fabros' commitment to build their data center infrastructures to the highest specification level commercially available, while providing battery-free maintenance, distributed redundancy, load sharing flexibility and fault tolerance.

Mike Mosman, PE is Vice President of CCG Facilities Integration Incorporated. He can be reached at mmosman@ccgfacilities.com.

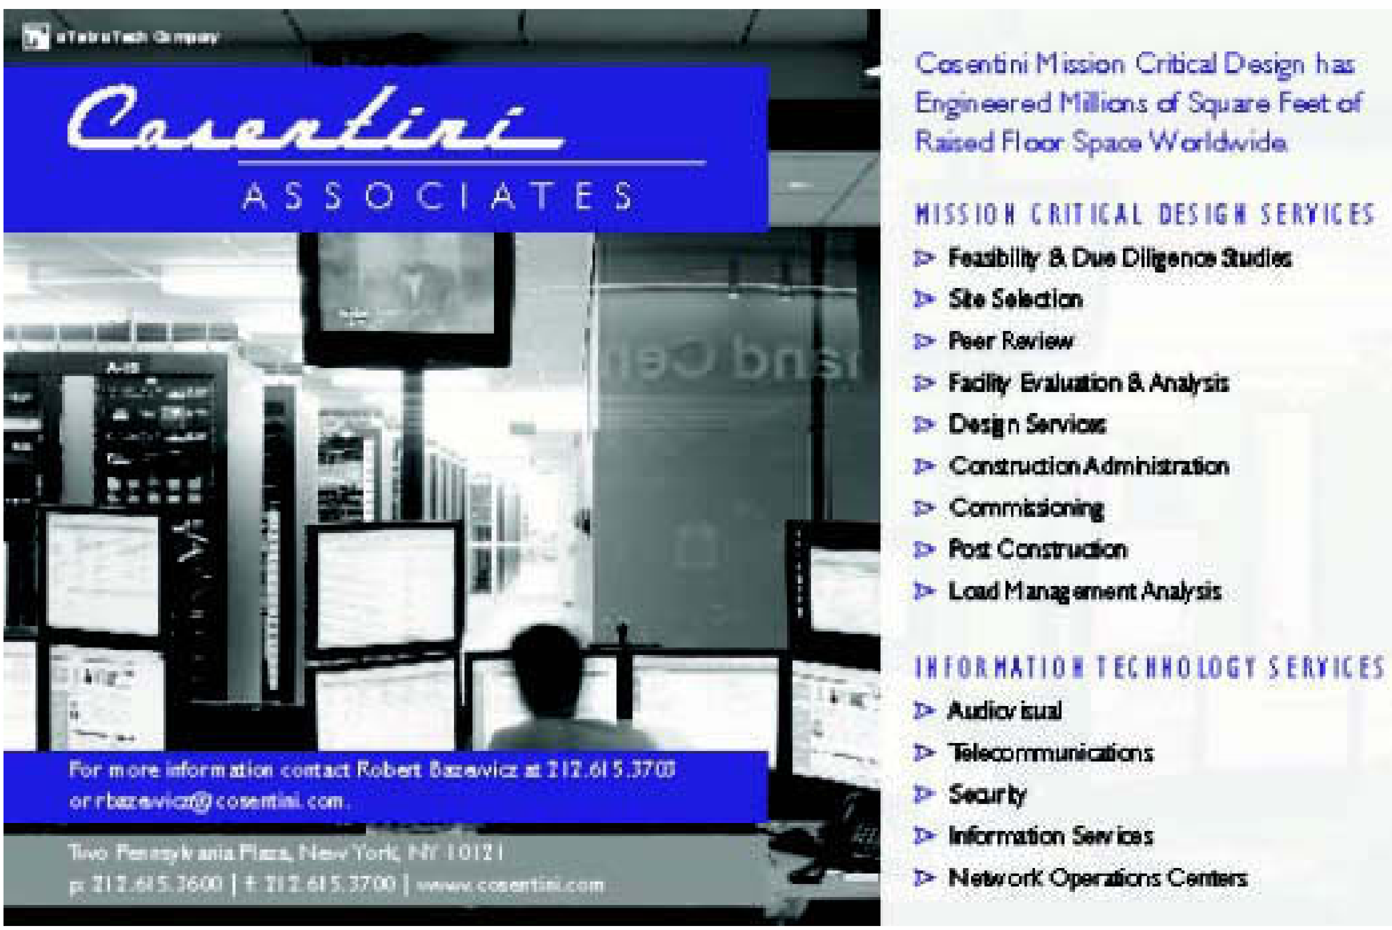

Russelectric custom engineers every power control system to specific customer requirements

"We wanted triple redundancy in the controls for our critical power system.

Russelectric provided redundant 17" touchscreen controls plus full manual operation."

At Russelectric, special requirements or unusual designs are not interruptions of our business...

they are our business.

Russelectric specializes in using proven technologies and components in innovative ways to satisfy even the most exotic customer requirements. We'll never compromise on quality, safety, or performance. And, above all, we'll never ask our customers to settle for less than exactly what they want.

Independently owned and operated, Russelectric has only one goal... to take care of our customers.

1-800-225-5250
www.russelectric.com

Russelectric
POWER CONTROL PEOPLE YOU CAN RELY ON

DATA CENTER FACILITY CONTROL SYSTEMS: COMMERCIAL VS INDUSTRIAL

by Steven Blaine, PE

System uptime is the crucial objective of data center operations yet, unlike other mission-critical facilities, data centers are typically built with commercial DDC (Direct Digital Control) control systems. Consider this: a study at Lawrence Berkeley National Laboratories identified controls as the single largest source of HVAC system problems (see below). The biggest threat to data center availability may be the control system.

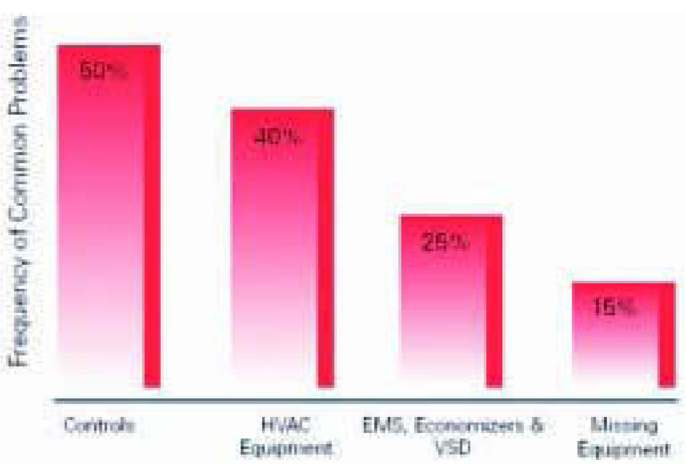

Figure 1: Frequency of common problems encountered in a 60 building study performed by Lawrence Berkeley National Laboratories (2)

It's time for owners and system designers to consider the benefits of industrial PLC (Programmable Logic Controllers) / SCADA (Supervisory Control and Data Acquisition) control systems.

Building automation controls are generally regarded as a simpler subset of process control. In most buildings, thermal mass is large, response times are slow and the consequences of system failure are usually not severe. But in data centers, this is not true. If the control system does not respond quickly and appropriately, a data center may experience a destructive and rapid failure – even if redundant chillers, air handlers and power sources have been installed.

Data centers have unique and demanding HVAC requirements. One study by ComputerSite Engineering showed that during a cooling failure, the air temperature in a modestly loaded data center could see a 25°F temperature rise in only 10 minutes (see below). As heat densities increase, the response time will decrease to just a few minutes with an even higher temperature rise. This is enough to seriously disrupt or damage sensitive computer equipment.

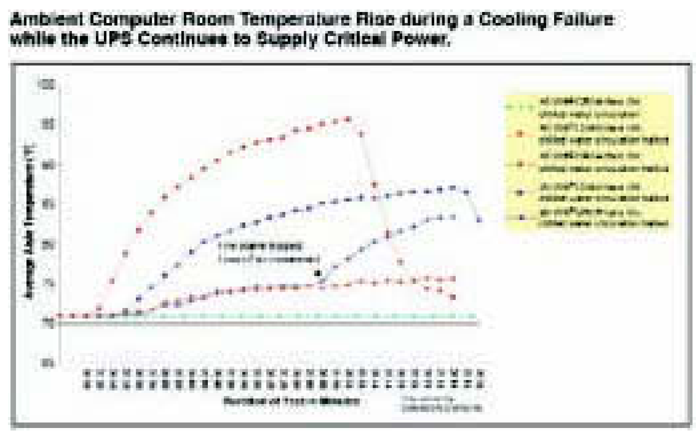

Figure 2: Continuous Cooling is Required for Continuous Availability – From the Uptime Institute (3)

In spite of these stringent requirements and the serious consequences of failure, most data centers are built with the same commercial DDC style control systems used in office buildings. This is in contrast to other mission-critical environments such as semiconductor cleanrooms, or pharmaceutical laboratories, where industrial controls such as a combination of PLCs with SCADA computers or even DCS (Distributed Control Systems) systems perform many of the same functions.

Cost Differences

A rule of thumb for control systems is this: industrial controls total installed cost is approximately $2000/point. Commercial systems cost approximately $1000/point. For reference, a recent data center project was completed with 500 I/O points. This represents a difference of $1M vs. $500K. This does not consider the difference in maintenance and service contract costs (which is typically higher for commercial controls) but is a reasonable idea of the difference in up-front costs. So, besides price, what differentiates industrial from commercial style controls? Following is an overview of the five main areas where commercial and industrial systems differ.

Critical Power Solutions

Who's backing up your backup power systems?

Ensuring reliable power takes more than the right equipment. You also need the right partner. At Schneider Electric, we've been delivering proven power solutions for more than 100 years. We bring together the products, resources and application expertise it takes to develop critical power systems that meet your exact needs. And we back it up with comprehensive service offerings and lifetime support so you can improve performance and uptime throughout your entire operation. Talk about staying power. To learn more, visit www.criticalpowernow.com or call 1-888-SQUARED.

SQUARE D

Quality of Devices

Automated control starts with the measurement of ambient and system parameters. The measurement process is a chain of sensors, transducers, analog to digital conversion and software processing. Errors and uncertainties at any point in the chain affect the accuracy of measurement and ultimately, the capability of the control system. For both DDC and industrial control systems, the largest source of inaccuracy is typically the sensor itself.

Sensors for temperature, humidity, flow, pressure, voltage and current are all used in data centers. Commercial sensors have a minimal accuracy requirement but are normally chosen for their low cost and sometimes, architectural attractiveness. Industrial controls generally use more accurate and robustly packaged devices.

DDC systems can use directly connected resistance measurements for temperature, and 0-10 VDC or 4-20 ma for other input signals. Industrial systems nearly always specify 4-20 ma current loops which are more impervious to electrical noise and wiring degradation. In commercial installations, sensor wiring is not always placed in conduit. Industrial sensor wiring is typically in conduit where it is protected from physical damage and electrical interference. At the controller, input signals are converted from analog to digital with different levels of precision. Commercial controllers have 10 or 12 bit resolution. Industrial controllers have 14 or 16 bit resolution. While not always significant for environmental parameters, higher resolution coupled with more accurate sensors and lower noise signals means better measurement.

Controllers and Software

All input and output signals eventually connect to some sort of controller – the central element of any control system. Commercial systems use a mix of "unitary" controllers to control a single piece of equipment and larger building controllers for facility wide programming tasks. Industrial systems use PLCs which also come in a range of sizes and intended applications. The differences between these controllers can be discussed in terms of form factor and physical robustness, I/O type and capacity, and processor programming capability and flexibility. These differences are discussed further in the long version of this paper.

Programming Capability

DDC programming languages have evolved from text based versions of high level computer languages like BASIC and PASCAL into graphical drawing versions. A DDC programmer creates a program or control strategy by placing a box representing the function block on the screen and connecting the inputs and outputs appropriately.

Once these graphical representations are complete, they are translated or compiled into machine readable code and downloaded to the controller. Each vendor has their own programming languages that are specific to their equipment – sometimes different software products for different controllers. DDC vendors often supply control program templates optimized for specific HVAC functions. The templates can match standard HVAC applications quite perfectly.

Programming a PLC is very different from programming a DDC. Like DDC manufacturers, PLC vendors each have their own programming software. In contrast, programming templates are not normally provided for PLCs. PLC manufacturers do offer a common software product that typically programs all of the PLCs they sell. There has also been a significant effort to standardize programming languages used by all PLCs. IEC 1131-3 is the international standard for programmable controller programming languages. This specifies the syntax, semantics and display for a suite of PLC languages. The result is that today, most PLC programs share a common look and feel regardless of the manufacturer. In the USA, PLCs are usually programmed in ladder logic. This visual language is quite familiar to electricians. In fact, its name comes from the hardwired relay control diagrams used to run machinery that look like ladders.

System Performance

The two types of systems conceptually can look very similar. The difference is performance. Industrial systems are designed for "real-time" control. Like a DDC, a PLC program looks at sensor input, performs logic or calculations and writes outputs. The speed of processing and communication in PLC systems is faster than DDC systems. This allows inputs to be read from anywhere in the system, logic solved, and outputs written to anywhere else in the system – in real-time.

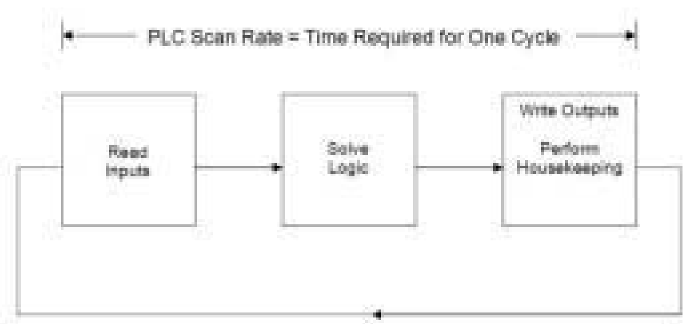

Figure 3: Explanation of PLC Scan Rate

The time it takes for a PLC to read inputs, solve logic, write outputs and perform overhead functions is called "scan rate". Scan rates for PLCs, even in large programs with distributed I/O, are generally measured in milliseconds. DDCs have program execution frequencies measured in seconds.

PLC and DDC programs differ fundamentally in flexibility. The programming functions in a PLC are more numerous and powerful. There is a richer instruction set for math, logic and bit manipulation. Many PLCs allow encapsulation of instructions to create user defined function blocks. This is a powerful tool that sophisticated users leverage to create simple, re-usable code. Finally, modification of PLC programs can be done "on-line" which means the controllers do not need to be stopped if the program needs to be changed.

System Architecture and Redundancy

Reliability should consider the dependability of individual items but also a system in which a failure in one part doesn't affect others. With careful engineering, control systems can be designed for fault tolerance.

One method of achieving fault tolerance is to provide distributed control. Valid for either commercial or industrial systems, small inexpensive controllers can be dedicated to individual machines or processes. In this case, the loss of a single controller cannot shutdown the entire facility. This type of design is called "passive automation". It implies that the system will operate properly even if the automation system is not performing its function.

Not all situations allow for passive automation. There are cases where a single controller must make decisions that require inputs from or outputs to multiple systems: large area temperature control, power restart, load shedding, humidity control, chiller sequencing and pressurization. These should always function continuously.

Instead of distributed control, another method of achieving high reliability is to build a fault-tolerant, redundant control system. With this approach, just a few or even one single controller pair can run an entire facility and no single failure can prevent continuous operation. A good design of this type requires careful consideration of each system component. These may include redundant controllers, network media, power supplies and SCADA servers. It may also include dual power feeds, redundant and separate supervisory networks. PLCs have evolved sophisticated capabilities in this area but DDCs have not.

Other Control System Considerations

The following is a list of features that should be compared when considering a particular architecture or controller for data center applications. DDC functionality has increased tremendously in the last few years and the latest systems from the leading suppliers can provide most or all of these features. It still must be said that some of these items are more difficult or impossible to accomplish with DDC systems:

- **Ability to "hold last state"** – during communication loss or programming downloads, this ability can prevent loss of control or a lengthy recovery period.
- **Master/backup loops** – critical control loops are sometimes programmed so that one controller is the master but a second is the backup. In case of a controller failure, the loop continues to operate.
- **Hot swap of modules** – PLC modules are often designed to be replaced under power. This feature prevents the necessity of powering down a controller to perform a repair.

Conclusion

We have seen how industrial control systems differ from commercial systems in their use of more accurate and rugged sensors and devices, signal types and wiring methods. Industrial controllers are more robust, have higher performance, faster networks and more flexible programming capability. Redundancy options with industrial controls can address the most difficult control issues without relying on "passive automation".

While both DDC and PLC/SCADA systems are capable of controlling the facility equipment found in a typical data center, owners and system designers should be aware of the differences in performance, flexibility and reliability, and not expect to achieve industrial control system performance on a commercial control system budget.

Steve Blaine, PE is an Instrumentation and Controls Specialist with IDC Architects, a CH2M HILL Company. He can be reached at steve.blaine@ch2m.com.

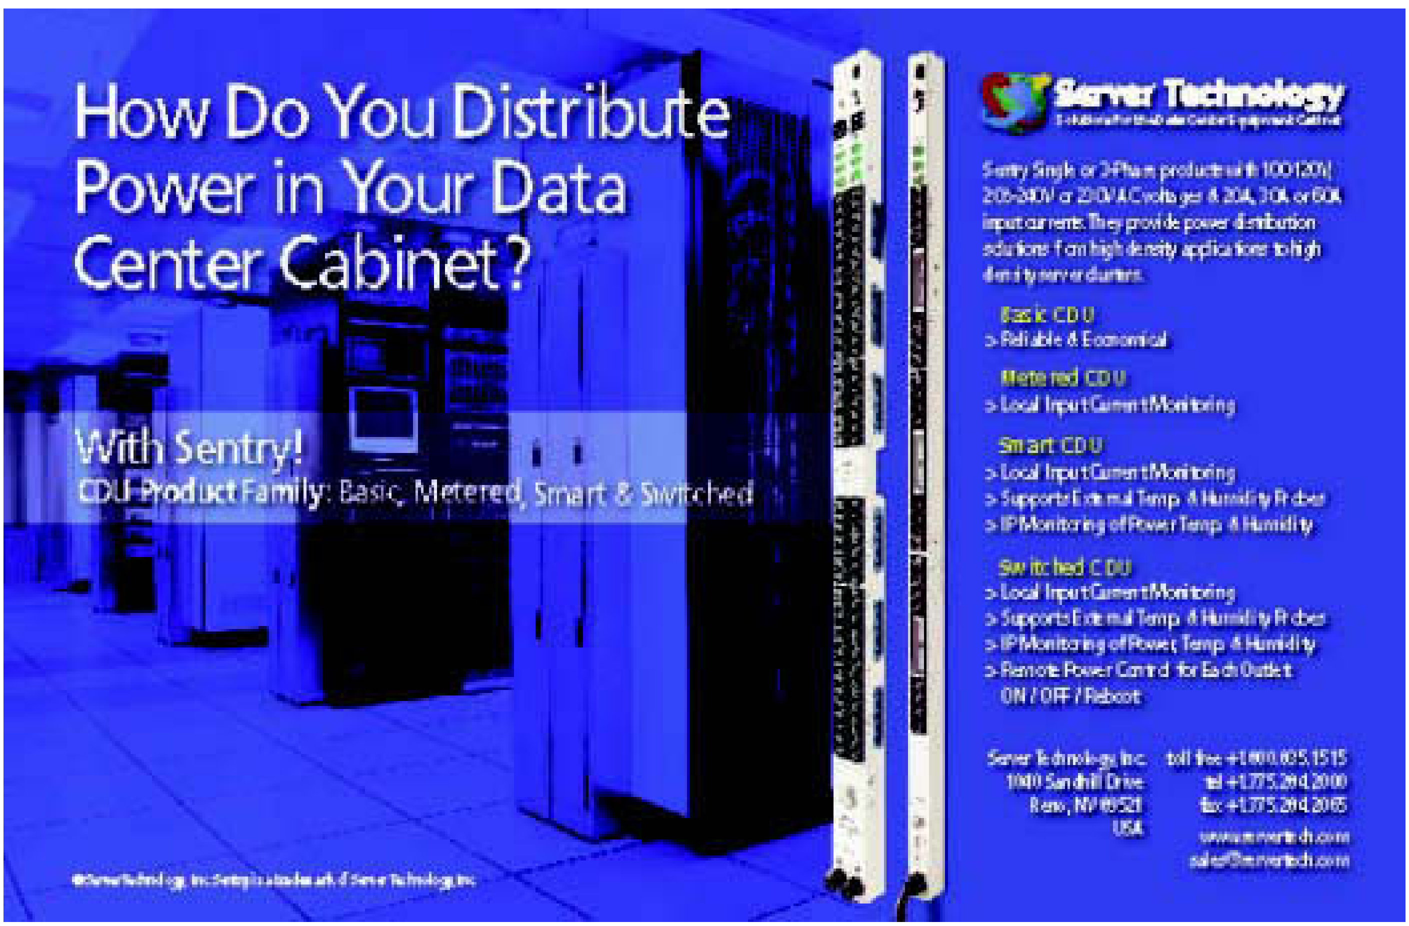

DATA CENTER COOLING ENERGY BASICS

by Donald Beaty

Introduction

The focus on data center energy is rapidly increasing. Energy Star, EPA, Congress Public Law 109 – 431, ASHRAE, Green Grid, LBNL are all involved in data center energy one form or another.

The are many ways to approach energy usage and conservation. Innovation can certainly reduce energy but there is plenty of energy saving opportunity that is relatively easy at the basic level. The basic level starts with "right sizing" which applies to both cooling equipment and cooling systems. This article starts with a practical example of "right sizing" and then describes the following 5 steps to achieving "right sizing" and the associated energy savings.

- **Step 1** – Determine the Actual Maximum Load
- **Step 2** – Determine the Load Profile
- **Step 3** – Determine the Future Load
- **Step 4** – Determine the Operating Conditions
- **Step 5** – Equipment and System Selection

Practical Example of Right Sizing

How efficient is it to size something 100 times larger than necessary? Will it use the same amount of energy as something that is sized precisely for the load it serves? Common sense correctly tells us that grossly oversizing is not efficient. A practical example of "right sizing" is to consider various ways to fill a glass of water.

OVERSIZING... How much water will actually enter the glass if it is filled with a fire hose? Depending on how long you try to fill the glass, the entire area around the glass will be covered with the overspray which is essentially wasted water. The wasted water is a direct parallel to the wasted energy associated with mismatching the load and the cooling system capacity.

RIGHT SIZING... In the same glass example, what if the glass is filled with water from a sink faucet that is easily controlled; all the water enters and stays in the glass with no waste. This is an example of matching the supply with the need; "right sizing".

WASTE DUE TO POOR CONTROL... Take the glass example still one more time. What if the glass is filled with a faucet that has an automatic shutoff? The glass will be filled but you have to move the glass around to re-engage the sensor so that the faucet does not prematurely shutoff. The net result is the glass will be filled but the automatic shutoff will keep the faucet running for a period of time after the glass is filled. This is an example of inadequate control and not synchronizing the cooling supply with the load.

Right sizing is a good energy saver and is easily achievable. One way to accomplish this is to use the following 5 steps.

Step 1 – Determine the Actual Maximum Load

What is the maximum load for each piece of equipment? Especially in a multi-vendor environment, this can be quite an undertaking. Typically the terms used vary greatly resulting in inherent inaccuracies due to not using a common baseline and unit (e.g. load for 1 piece of equipment stated in watts while another piece of equipment stated in amps). For example, some terms used to express power or load include:

- Watts or KW
- Amps
- Circuit Amps
- Wiring Amps
- Nameplate Amps
- Full Load Amps
- Peak Amps
- Measured Load
- General Allowance in KW/Rack

Frequently the nameplate amps are used because it is a number that is readily available (attached to the equipment). Unfortunately the nameplate information is a regulatory requirement and is focused on safety; not on accurately characterizing the load.

Typically the nameplate amps grossly overstate the load. To compensate for this, many apply an adjustment factor (e.g. 40 to 60%). There is great inconsistency in the measured load versus the nameplate value and therefore a standardized adjustment factor can be inaccurate (oversize or undersize).

The most effective approach to identifying equipment load is to request an ASHRAE Thermal Report for each piece of equipment. This report identifies the measured load (heat release) in watts and provides this information for multiple configurations of the equipment such as number of processors, amount of memory, or the amount of hard disks.

Step 2 – Determine the Load Profile

Imagine looking at the compute load or workload for an entire data center in 5 minute increments across an entire year. It is hard to imagine that there would be no fluctuation in load and easy to imagine significant fluctuation.

“Six-nines” reliability

To operate with the utmost reliability, your data center must achieve 99.9999 percent uptime. PowerCommand® pre-integrated power generation systems from Cummins Power Generation Inc. help give you end-to-end reliability. They're specifically designed to fit and work together—providing 24/7 assurance of uninterrupted power, exceptional data reporting and the capacity to match your expanding infrastructure.

Our network of local distributors is at your service to deliver everything from design assistance and installation to 24/7 maintenance and support.

Contact your local distributor today or visit www.cumminspower.com/local

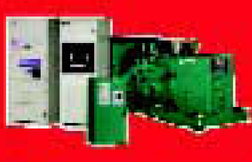

Our energy working for you.™

Power Generation

Total power solutions

PowerCommand power systems are fully integrated for seamless transfer of power to your loads, delivering better performance with increased functionality. They combine digital paralleling systems, automatic transfer switches, remote monitoring and microprocessor-controlled generator sets.

For more information, contact your local Cummins Power Generation distributor or visit www.cumminspower.com/solutions

© 2007 Cummins Power Generation Inc. All rights reserved. Cummins Power Generation and Cummins are registered trademarks of Cummins Inc. PowerCommand is a registered trademark of Cummins Power Generation. "Our energy working for you." is a trademark of Cummins Power Generation. CPG-417 (9/07)

Typically within a data center, there is more than one software application running. Further the activities and load on each application are often varying. This compute load profile is also complicated by software upgrades and applications that are installed but not serving load (e.g. cannot run until all sites have the software installed).

The cooling system is not only impacted by the compute load but by climate as well. It is very beneficial to produce annual load profile curves for the compute load, climate, and a composite of compute load and climate. The same profile curves should be provided for an hour period and day in addition to the annual.

These profiles should be produced even if they are not much more than educated guesses using minimal input from IT operations, etc.

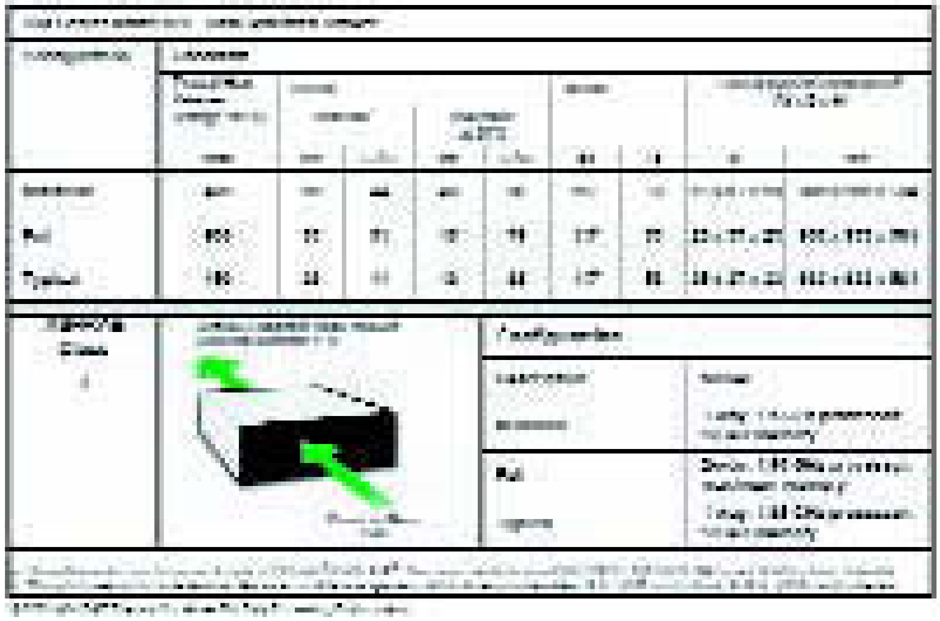

Figure 1 shows an ASHRAE Thermal Report for the IBM model 520. Where a Thermal Report is not available, simple measurements can be made to provide a general characterization of the piece of equipment and informal produce a Thermal Report.

Step 3 – Determine the Future Load

ASHRAE has tables that identify the projected life of cooling equipment (10 to 25 years). IT equipment has a much shorter life. For ease of comparison, if the IT equipment life is 2 to 5 years, there will be 5 generations of IT equipment in one lifetime of cooling equipment.

As a result cooling system design and sizing should consider both the current load and the future load. Often, the data center must remain in operation during a partial or full upgrade/refresh of IT equipment. This means that upgrading the cooling system to add capacity or cooling interface points must occur with minimal to no disruption; this is very costly and difficult to accomplish.

Step 4 – Determine the Operating Conditions

The operating conditions (temperature and humidity) have a significant impact on:

- Amount of cooling being accomplished by economizer (free cooling)
- Pumping and fan horsepower since the larger the Delta T (temperature differential), the less the required flow.

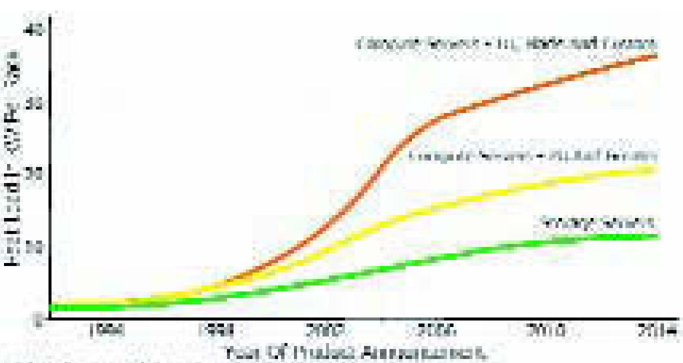

Figure 2 is a modified chart from the ASHRAE book, Datacom Equipment Power Trends and Cooling Applications. This chart provides projections through to 2014. These loads are the highest loads for any single manufacturer for any given class of equipment. The intent is to apply an adjustment or derating factor to these loads to determine the future load for a specific facility.

- Amount of humidification required and its associated energy

ASHRAE Book, Thermal Guidelines for Data Processing Environments, provides both recommended and allowable ranges for temperature (20 to 25˚C or 68 to 77˚F) and humidity (40 to 55% RH) as well as temperature rate of rise.

Step 5 – Equipment and System Selection

Develop several cooling alternatives based on the current and projected capacity needs as well as the operational load profiles. Optimize equipment selections based on operating ranges. Alternatives should include various degrees of scalability. For example, the current load may be significantly less than the ultimate load. A single piece of equipment or system may not be efficient at both the current or day 1 load and the ultimate or final load.

Comparison of the equipment choices and systems should be based on Total Cost of Ownership (TCO). The TCO cost should include the cost to remain in operation during a IT and cooling upgrades.

Donald Beaty is President of DLB Associates Consulting Engineers. He can be reached at dbeaty@dlbassociates.com

Summary

Data center energy usage is significant. Right sizing through the 5 step process or any other structured approach has the potential to yield significant energy savings.

- **Step 1** – Determine the Actual Maximum Load
- **Step 2** – Determine the Load Profile
- **Step 3** – Determine the Future Load
- **Step 4** – Determine the Operating Conditions
- **Step 5** – Equipment and System Selection

"STARLINE Track Busway is leading the industry in changing the way Data Center Design Professionals distribute power for Mission Critical applications. Eliminating all whips, cable tray and distribution panels, STARLINE full access plug-in power modules offer improved reliability, scalability and the ability to easily change and relocate in a matter of seconds."

STARLINE
TRACK BUSWAY

Universal Electric Corporation • 168 Georgetown Rd • Canonsburg, PA 15317 • 1 800 245 6378 • www.uecorp.com

The Use of Outside Air Economizers

Introduction

Compared to most types of commercial buildings, data centers are energy hogs. A large data center, for example, can easily consume as much electrical power as a small city. Consider an average size data center – say a facility with an average demand of 1 megawatt (MW) over the entire 8,760 hours of the year. The cost of driving that 1 MW computer load is $700,000 per year (assuming a cost of electricity of $0.08/KWH). The cooling load associated with the 1 MW load is 285 tons. At an average chiller efficiency of 0.5 KW/ton, the cost of running that chiller is approximately $100,000 per year. Of all the facility's HVAC equipment, the chiller is easily the largest energy consumer. As such, significant HVAC energy savings can be realized by reducing chiller energy. The fact that the most energy efficient chiller should be selected for the facility goes without saying. However, an even more important issue is that reducing the number of hours of chiller operation can have a larger impact on the reduction of that piece of the building energy pie than simply selecting a more energy efficient chiller.

Syska Hennessy Group has formed a Green Critical Facilities Committee to address these very issues relating to sustainable design as they relate to critical facilities. Of necessity, the members of this committee represent the various industry specialties (engineering design, information technology, commissioning, etc.) in order to assure that a wide range of sustainable design elements are considered for every design. Our need (and especially our clients' need) to incorporate more sustainable and environmentally responsible design elements – especially in a facility that consumes staggering amounts of energy every single hour of the year – must accept a double-pronged approach toward a cooling system design. On the one hand, more energy efficient equipment must be selected. On the other hand, a method of reducing the hours of operation of the equipment must be incorporated into the cooling system design. Bracketing this double-pronged approach is the absolute necessity to assure that the overall system reliability is never compromised.

There are two types of economizers which can accomplish a reduction in hours of chiller operation – waterside economizers, and airside economizers. Waterside economizers are explained briefly below; the remainder of this paper will concentrate on issues relating to airside economizers.

A waterside economizer uses the building's cooling towers to cool the chilled water by taking advantage of the hours of the year during which the outdoor wet bulb temperature is sufficiently lower than the chilled water supply set point. In essence, rather than running the chiller during those hours, the cooling tower water is bypassed around the chiller and diverted to a heat exchanger so that it can cool the chilled water directly. This type of economizer has certain advantages and disadvantages, none of which will be addressed in this paper.

Airside economizer – brief overview

An airside economizer takes advantage of the hours of the year during which the outdoor enthalpy (energy content) is less than the return air enthalpy. Under such outdoor conditions, using outdoor air reduces the load that would be experienced at the air handling unit when compared to using the return air flow.

When the outdoor enthalpy is less than the return air enthalpy AND higher than the supply air enthalpy, some mechanical cooling is still required in order to meet the requirements at the supply air setpoint. Under these circumstances, the chiller will be required to operate, though not at as high a load as would be required for a 100% return air system. This is a partial (or integrated) economizer.

When the outdoor enthalpy is less than the required supply air setpoint, no chiller operation is required, and the actual supply conditions can be met by either mixing outdoor air with return air (if the outdoor air is below the supply air setpoint), or using 100% outside air (if the outdoor air is at the supply air setpoint). This is considered a full economizer. Significant energy savings can be realized whether a partial or full economizer is utilized. Figure 1 shows the basic flows involved with a 100% recirculating system and an economizer system.

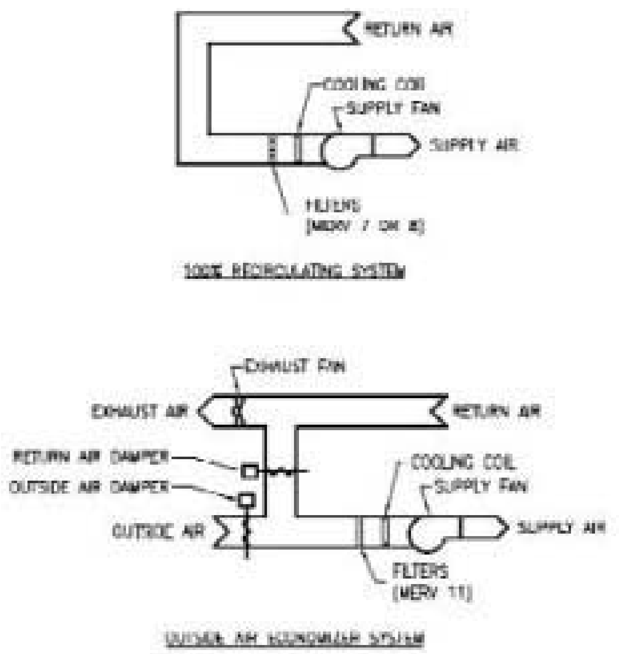

Figure 1: Comparison of 100% Recirculating System and Outside Air Economizer System

How many hours of economizer use in various cities

Exactly how many hours of the year are available for economizer use? The weather data for several representative cities – Dallas, TX; New York, NY; San Francisco, CA; and London, UK were evaluated. For each of these cities, dry bulb and dew point conditions were examined and compared with an ideal facility using a supply air setpoint of 68˚F dry bulb (DB) and a dew point of 50˚F. (For more discussion of why these conditions were selected, see discussion under

Temperature/Humidity Control, below.) Using these weather data, three "bins" were established to collect and classify all the yearly data:

- the number of hours during which an airside economizer is available to provide 100% of the facility's cooling needs
- the number of hours during which an airside economizer is available but cannot meet all the facility's cooling needs (partial economizer)
- the number of hours during which an airside economizer should not be used (i.e. the return air conditions are more favorable than the outdoor air conditions)

The results are summarized in Table 1, below:

YEARLY UTILIZATION OF AIRSIDE ECONOMIZER

Representative Cities	Available Hours of Full Economizer	Available Hours of Partial Economizer	No Economizer Availability
San Francisco, CA	8,563	197	–
New York, NY	6,634	500	1,626
Dallas, TX	4,470	500	3,790
London, UK	8,120	300	340

***Calculation of available hours based on a 68°F DB / 50°F Dew Point supply air.*

Table 1: Hours of Economizer Use by City

Common sense would normally dictate that an outside air economizer in the hottest climates would not have a good payback. That logic may be applicable to a typical office building where there are approximately 2,500 hours of use in a year. However, a data center must run continuously, 24 hours a day, for a total of 8,760 hours per year. The number of hours of availability is, therefore, greatly increased. In the case of Dallas – the warmest climate considered – the hours of availability amount to more than half of all the hours of the year. For cities like San Francisco and London, where the annual hours of full economizer availability are higher than 8,000, the use of an economizer requires almost no complicated analysis and should most definitely be considered. (There may be other issues that come into play that could restrict the use of an airside economizer, such as the lack of availability of sufficient building openings to the outdoors for air intake and exhaust. Under those circumstances, a waterside economizer should be considered.)

Flip side – two schools of thought on use of economizers for data centers

Historically, the industry has in general avoided the use of outside air economizers when dealing with data centers. Even ASHRAE's TC 9.9, the technical committee which deals with issues of data center design and operation, has avoided making any recommendations about the application of outside air economizers until more research can be provided to either support or reject its use for data centers.

The main points to be considered when such large quantities of outside air are introduced into a data center are as follows:

- Introduction of outside air into a data center can be a source of dirt, dust, and other airborne contaminants which can place the investment of computer equipment in that center at risk.
- Introduction of outside air into a data center can play havoc on the ability to control the space's relative humidity.

Both of these issues should be addressed during the design phase for any data center facility. Both issues have been addressed successfully in the past in other types of facilities that have both critical needs AND a requirement for large quantities of outside air, such as hospitals and laboratory facilities.

Yet, the data center design community will not accept the use of outside air economizers if doing so will result in lower facility reliability or reduced uptime. Examining the related design issues and addressing each item one by one, there is a case to be made that the issues of air contaminants and humidity control can both be resolved as evidenced by recent case-studies published by LBNL (Tschudi, 2007). With those issues resolved, there should be a strong impetus to use outside air economizers based on improved energy efficiency. If a large data center that has been designed according to today's "best practices" uses as much energy as a small city, the savings that can be realized implementing an economizer system can be equivalent to the energy used by a large community.

Filtration of outside air

The first and easiest item to address is the issue of air contaminants. When introducing a large amount of outside air, it is necessary to increase the filtration at the air handlers. With 100% recirculating systems, filters with a Merv rating of 8 or 9 are used. These filters are intended to remove only the particulates that are generated by the activity within the space. When outside air is introduced, it is necessary to increase the Merv rating to 10 or 11 so that the filters can extract the increased loading of particulates (i.e. the smaller particulates) associated with construction, traffic, industrial processes, and other outdoor pollutants. The resultant "cleanliness" of the outdoor air in terms of particulate content will be as good as the resultant "cleanliness" of a recirculating system with the lower Merv rating. The higher Merv rating filter will create a higher operating pressure at the fan, and this is associated with an increase in energy usage. However, this extra energy usage is small in comparison to the savings associated with reduced operation of the chiller plant.

Temperature/Humidity control

With the publication of ASHRAE TC9.9's document "Thermal Guidelines for Data Processing Environments" (2004), the industry has come to a consensus about the optimal thermal environment for data processing equipment. The recommended environment is 68˚F to 77˚F DB and 40% to 55% relative humidity at the INLET of the

Simulate Airflow and Temperature to Meet Your Cooling Challenges

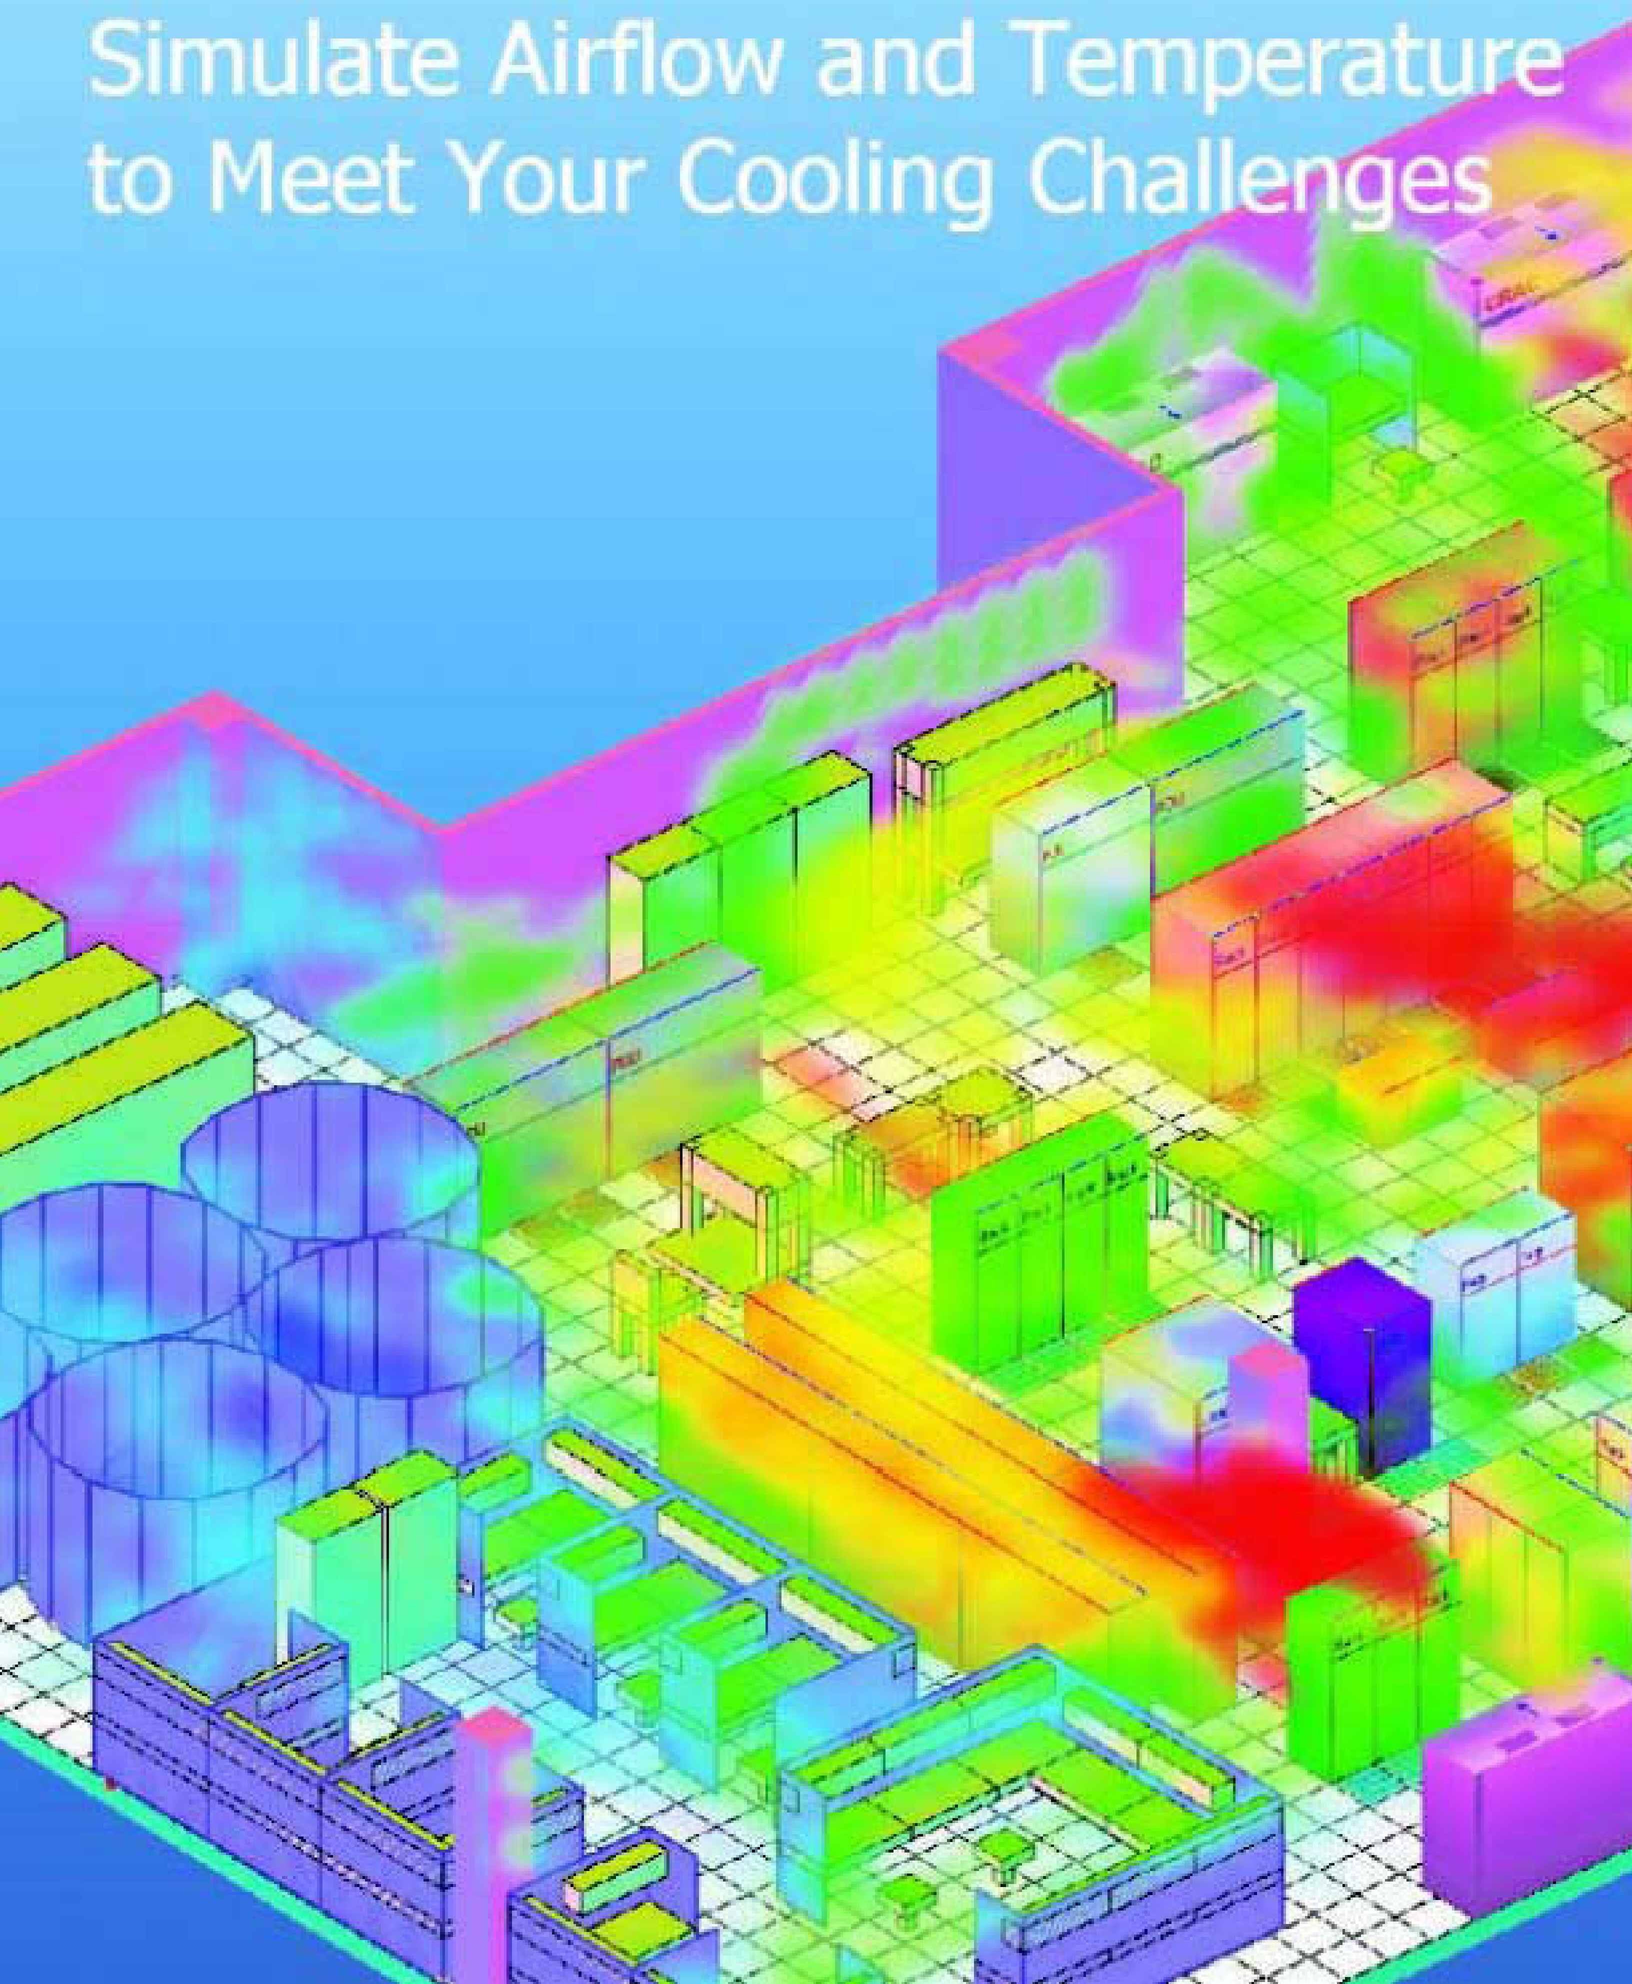

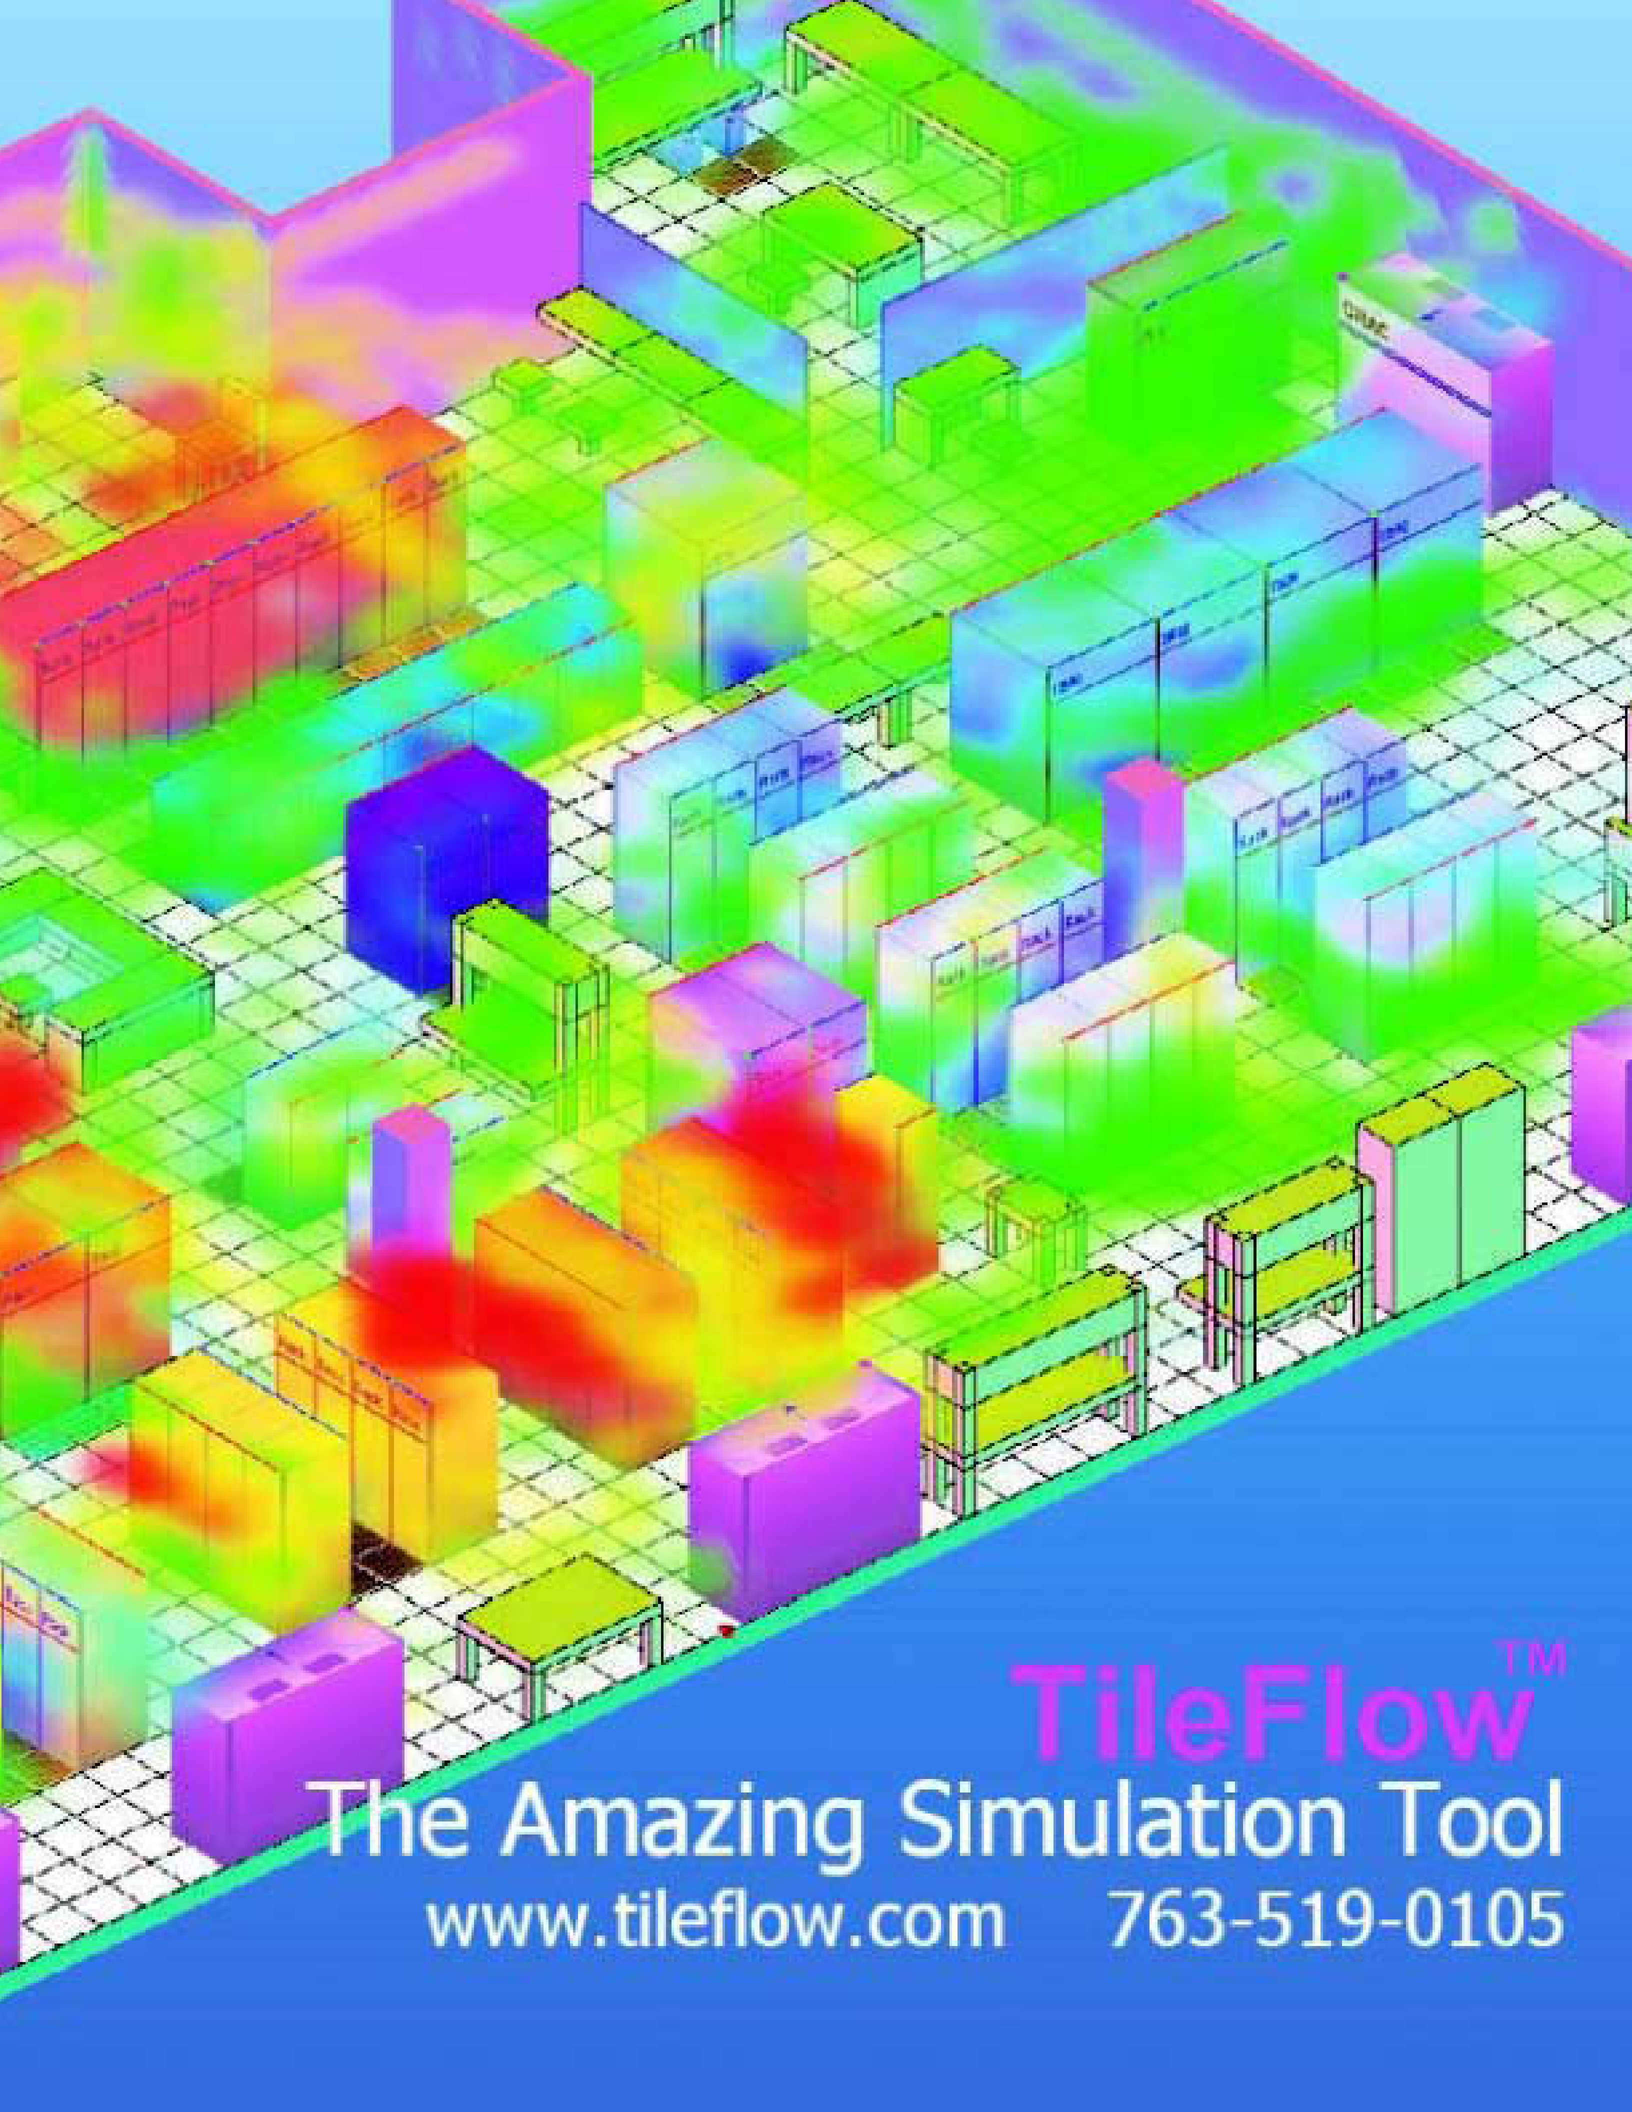
TileFlow™
The Amazing Simulation Tool
www.tileflow.com 763-519-0105

equipment. (See Figure 2.) The temperature at any other location within the data space is irrelevant. This is a critical point to understand, and it highlights the basic differences between cooling for comfort and cooling for equipment.

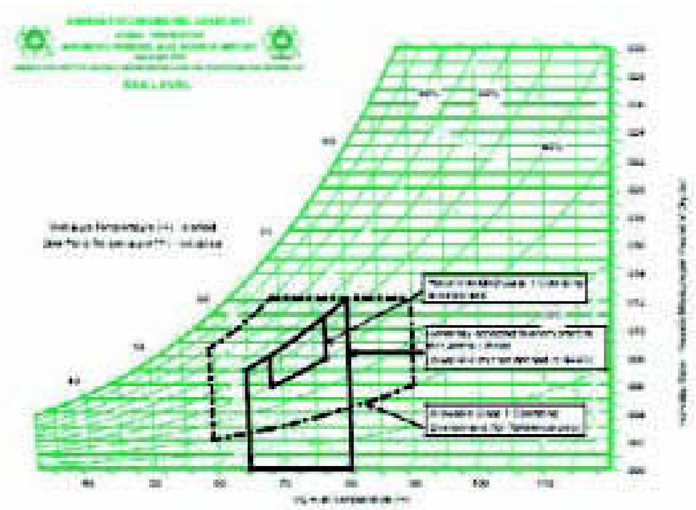

Figure 2: ASHRAE-Recommended "Class 1" Operating Environment

In a comfort cooling environment, the supply air is usually provided by an overhead system within a range of 55˚F to 60˚F dry-bulb (DB). This air is then thoroughly mixed into the occupied space, and the temperature perceived by the occupants is represented by this mixed condition. The thermostat that controls the occupied space must be somewhere in the occupied space to insure optimal comfort.

In a data processing space with cooling provided from an underfloor plenum, the cold supply air is separated from the hot discharge air by arranging equipment cabinets in alternating hot and cold aisles. This separation of air streams protects the equipment from overheating by preventing the hot discharge air from recirculating back to the equipment inlets. Therefore, when put to practice properly, the inlets to the equipment will be close to the supply air temperature. The thermostat should be placed in the supply air stream to be as close to the controlled environment as possible. The biggest mistake that designers make (and the computer room air conditioning equipment manufacturers propagate this mistake) is to put the thermostat in the return air stream. The facility manager cannot properly control his equipment's thermal environment (i.e. at the inlet to the equipment) by sensing the return air temperature that is physically distant from the controlled environment.

Since the Thermal Guidelines promote the warmer design temperatures, designing to a fixed supply air temperature opens up new possibilities. If the hot aisle/cold aisle concept is properly implemented, and sufficient air is supplied into the cold aisles to preclude hot air recirculation from the hot aisle back to the cold aisle, the cold aisle will remain uniformly at one temperature – the supply air temperature. There is no reason to cool supply air to 55˚F or 60˚F DB when 68˚F DB and higher is recommended for the equipment inlet temperature.

There are many benefits of designing around 68˚F DB as a supply air temperature:

- The thermal environment for the equipment will fall within the recommended envelope, in compliance with the Thermal Guidelines. The result will be optimized equipment reliability.
- There are many hours of the year during which the outdoor temperature falls within the range of 60˚F and 68˚F. By designing around the higher supply air temperature, these hours become available for full outside air economizer use. The onset of chiller operation is delayed, and the total hours of chiller operation are reduced. (The analysis presented above for various cities assumes a 68˚F supply air setpoint and the availability of these extra hours of economizer utilization.)
- Selecting a cooling coil for the higher supply air temperature also allows more sensible heat to be removed from the space for a given air flow. One reason for this is that the amount of latent heat removed unnecessarily by the coil is reduced or eliminated, as well as the unnecessary humidifier operation needed to return the lost humidity back to the space. As such, the air handler coil and the heat transfer process will operate more efficiently.
- With the higher supply air temperature, the chilled water temperature can be raised. The result is that the chiller will operate more efficiently for the fewer hours that it does operate.

This discussion has not yet addressed the issue of humidity control. The Thermal Guidelines define the top and bottom of the thermal environment envelope (as viewed on a psychrometric table) in units of relative humidity. This poses a design challenge since relative humidity varies depending on where in the space it's measured. The higher the local temperature, the lower is the relative humidity. The result is that one can properly control the dehumidification or the humidification processes only if the relative humidity is measured where supply air temperature is measured (assuming that we are controlling on supply air temperature, as noted above). This is not always practical.

There is a simple solution using absolute humidity and/or dew point sensors. Since there are no latent loads in data spaces, the measured absolute humidity or dew point will be uniform throughout the space – from the supply air, to the cold aisles, to the hot aisles, to the return air. This is not the case for relative humidity. If these sensors are placed in the return air stream, they are in the perfect location to measure the return air enthalpy. (Temperature sensors must also be included in the return air stream to help determine whether to use the air side economizer or mechanical cooling. However, to avoid unstable temperature control and "fighting" between adjacent air handling units, these temperature sensors must not be used to control the cooling coil. Supply air temperature sensors must be used for that purpose.) The dew point sensors in the return air stream then serve a dual purpose – they're used as the controlled point for space humidity control, and they're used as part of an enthalpy economizer.

Using an enthalpy economizer is the last component to consider in the use of outdoor air economizers. The issue of enthalpy economizer was introduced earlier in this paper, and was used only in the context that enthalpy relates to the energy content of an air stream. In essence, an enthalpy economizer looks at both the temperature and humidity of both the outdoor air and the return air, compares the condition of each, and determines which air stream (or what combination of both air streams) will utilize the least amount of mechanical cooling. With a full economizer, the chiller and its associated equipment is turned off. There is no mechanical cooling. With a partial economizer, some mechanical cooling is needed, but cooling the outdoor air rather than the return air to the supply air setpoint leads to a net lower usage of energy. Humidity IS considered under these conditions. In most major cities in

EXER THERM™

SPOTS POWER FAILURES BEFORE THEY HAPPEN

The Next Technology Step for Mission Critical electrical equipment

24/7 Continuous Thermal Monitoring INSIDE the enclosure

Modbus output for easy integration

Why take a snapshot when you can now have the WHOLE PICTURE

ExerTherm USA / North America
599 Albany Ave, North Amityville, NY 11701
www.ExerThermUSA.com
sales@ExerThermUSA.com
Tel: 631-841-2300 Fax: 631-841-2305

this country and Europe, an enthalpy economizer makes economic sense. For some cities in the southwestern parts of the US, where the climate is dry when the outdoor temperatures are high, a simple dry bulb economizer can work just as well and would cost less to install and implement.

With a full economizer in operation and the control components in place as described above, humidity control becomes a straightforward process. If the resultant mixed air condition is too dry (i.e. the dew point is too low), a humidifier adds moisture into the air stream to increase the dew point. To reiterate what was noted above, this process must be controlled by dew point or absolute humidity sensors to assure that cooling coils don't dehumidify while the humidifiers simultaneously add moisture to the space.

If the resultant mixed air condition indicates too high a dew point, the chilled water temperature and supply air temperature should be bumped downward in one degree increments over a several hour period to wring moisture out of the space through the cooling coil. This would normally occur whenever the outdoor conditions are hot and humid. However, during such conditions the economizer would be off since the return air would provide the more favorable enthalpy. This condition would be no different than any other system without an economizer.

Conclusions

With the large energy usage and costs associated with data centers, the incentive to find ways to reduce the staggering costs of operating these facilities is huge.

Outside air economizers provide one of the best ways of reducing these costs. Yet, there has been considerable resistance to these economizers in data centers. The main reasons for this resistance have been the fear of outdoor contaminants entering the facility's critical equipment, the perception that humidity control can become complicated and unstable, and the difficulty of keeping humidity or enthalpy sensors calibrated. With the application of appropriate design principles and control strategies, the first two issues can be addressed in a manner that assures that the reliability of the facility is not compromised. The third item can be addressed simply with the statement that there have been significant improvements in the quality of sensors in the last several years. The newer sensors are more reliable and are able to hold calibration longer.

Finally, the need for a thorough commissioning program at the inception of the project coupled with a program of continuous retro-commissioning throughout the life of the facility cannot be stressed enough. The commissioning effort is especially critical in guaranteeing that the outside air economizer does, in fact, perform according to the design intent.

Vali Sorell, PE is a Senior Associate at Syska Hennessy Group. He can be reached at vsorell@syska.com

References — ASHRAE. 2004. Thermal guidelines for data processing environments. Atlanta: American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

Tschudi, William, 2007. Data Center Research Roadmap. LBNL (http://hightech.lbl.gov/datacenters-bpg.html)

COOLER

SMARTER

GREENER

Optimize what you've got with Upsite's products and services

Cut the energy and operating costs of installing additional cooling equipment and reduce your carbon footprint with Upsite Technologies' state-of-the-art solutions, including KoldLok®, for optimizing your data center and ensuring site uptime.

Upsite's suite of products and services currently optimize more than 6 million ft² of data center space.

upsite technologies™

To learn more about data center cooling challenges and easy-to-implement solutions, please download our Executive Summary, *Reducing Bypass Airflow Is Essential for Eliminating Hotspots* at upsitetechnologies.com/[illegible]

Fan Energy Reduction in an Existing Data Center – A CASE STUDY

by Michael H. Schwarz, PE

INTRODUCTION

Many modern-day data center facilities contain constant air volume air-conditioning systems for electronic equipment cooling, which typically consist of packaged Computer Room Air Conditioning (CRAC) units serving a common underfloor air distribution plenum. Data centers are also usually provided with multiple redundant CRAC units with no automatic controls to adapt the air-side system to real-time load conditions. The paradigm is to operate CRACs simultaneously and continuously, ensuring adequate cooling and multiple back-up units, but possibly causing short-cycled cooling airflow and excessive plenum static pressures – resulting in wasted energy. Energy consumption related to CRAC unit fans can range from 5-10% of the total energy expended in a data center at a full load, including computer loads and electrical distribution losses, and in a typical cooling infrastructure is second only to cooling compressor energy consumption. But unlike compressor energy, CRAC unit fan energy consumption typically remains constant and does not fluctuate according to the data center server load they serve, which is usually variable. According to the U.S. Environmental Protection Agency's recent Report to Congress on Server and Data Center Energy Efficiency, in 2006 data centers consumed a total of 61 billion kilowatt hours, or 1.5% of the total electricity used in the United States. CRAC fan energy is a key target for energy reduction in existing and new data centers. However, before airflow is reduced to below design values, the airflow patterns in an existing data center must be understood so the impacts may be tested and quantified.

This was the case at a data center site of a leading global financial services firm located in the mid-Atlantic United States. As part of risk assessment and energy reduction studies, KlingStubbins created Computational Fluid Dynamics (CFD) airflow models for each data center room to understand the site-specific airflow patterns. The main objective of the modeling effort was to develop a quantitative tool for analyzing air management best practices, and to determine which CRAC units are required to operate given the existing data center load conditions.

CASE STUDY DATA CENTER ROOM

Each room in the data center is 45,000 ft^2 and 25 feet high, with a design cooling capacity of 90 watts/ ft^2, with 46 CRAC units that are spread along the perimeter and in a row across the center of the space. Physical obstructions below the 36 inch raised floor include chilled water piping, cable conveyance systems, electrical conduits, structural steel supports for electrical equipment, and partition walls. The majority of the server racks are a type that receives cooling airflow through its bottom and side access doors. There is also an assortment of other rack styles, including front-to-back and bottom-to-top airflow configurations, open frames, and tape storage units. Figure 1 is a complete room CFD model, with simulation results displayed for maximum rack inlet air temperature and some objects hidden for clarity.

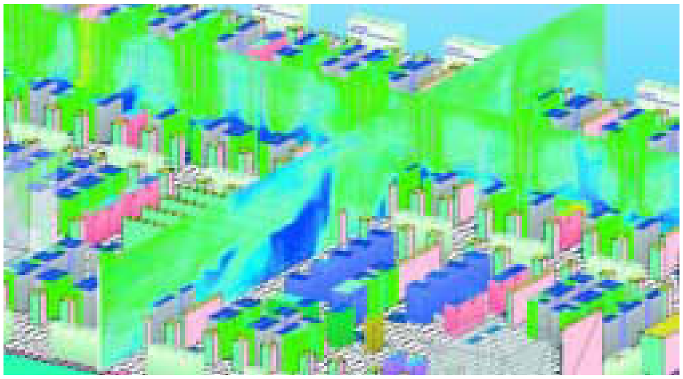

Figure 1

A survey of the data center was conducted, and raised floor properties such as floor opening measurements and locations of perforated floor tiles were determined. Above-floor features, such as server rack locations and airflow characteristics, wall transfer air openings, and overhead cable trays were also surveyed. A backpressure-compensating flow hood was used to measure supply airflow rates at a sampling of perforated tile locations across the raised floor, and measurement of return airflow was conducted using a velocity grid and manometer. Supply air temperature was also measured at these locations. These measurements provide firm data for calibrating each model to actual conditions.

Allowances were made in the models for how air is received by the server racks from the room and directly from the raised floor, and leakage air associated with cable management "sidecars". When they could not be surveyed directly, separate airflow measurements were performed for floor openings directly underneath racks, so their airflow characteristics as a function of static pressure are known. Because the floor tile cutouts at the cable management sidecars are the most common cutout style, the data center airflows are more sensitive to their characteristics than other opening types. An iterative process was therefore required to determine the approximate typical free area of the cable management cutout to attain the measured airflow rates through the perforated tiles. An allowance was also made for distributed leakage due to airflow through seams between raised floor panels.

EXISTING CONDITIONS

Once the model was calibrated for underfloor airflow, racks in the model were configured with actual load data from the data center's power management system. For the case study room, rack loads totaled 1,720 kW, or approximately 43% of the design cooling capacity. Maximum rack loads ranged up to 6.6 kW/rack.

Because all CRAC fans are running and the data center is operating at partial cooling load conditions, some CRACs are in heating mode to maintain their return air temperature set point or are merely re-circulating room air, therefore reducing the effectiveness of the cooling airflow. Many fans are currently operating with their chilled water valve positions at 0% capacity. This is causing a wide range of supply air temperatures. The lowest supply air temperature measured during the perforated tiles sampled was 57˚F and the highest 69. The supply air temperature in the existing conditions model was fixed at 62˚F, as this was the average underfloor temperature recorded in the data center at the time of the survey.

Figure 2
Measured vs. Calculated Data Center Airflow

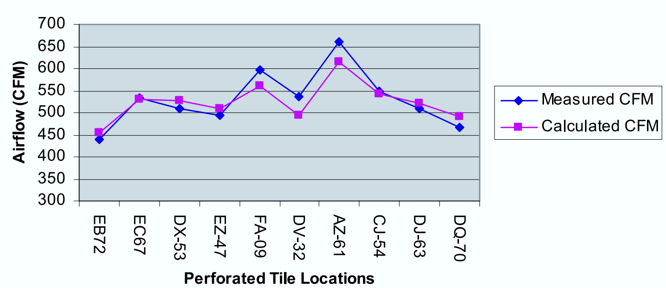

Comparison of raised floor airflows calculated in the base CFD model against actual measured values is a straightforward process, since inflow (CRACs that are operating) and outflow (floor openings and leakage) are physical attributes that can be measured. The accuracy of the model in terms of airflow is presented in Figure 1, and the results are within 10% of measured values for specific locations. Sources of complexity in validating temperatures across the CFD model are 1) the actual rack heat loads that were generated at the time temperatures were measured are less than the rack heat loads represented in the model and 2) the model is based on the maximum heat that theoretically can be generated by the computer equipment, even after it is de-rated from nameplate values. These full rack loads must be used for electrical and mechanical planning purposes, but since the operational diversity of the computer equipment is not included in the CFD model, calculated temperatures cannot be compared against measured values in detail. However, since the assumptions in the CFD models are conservative, the results may be dependably used in analysis efforts.

TESTING OF ENERGY REDUCTION EFFORTS

Since the data center is operating with significant spare cooling capacity as well as redundant CRAC units, the opportunity to immediately shut down selected fans was tested using the model. This scheme was called the "Airflow Reduction Scenario".

Due to the lower number of CRAC fans operating, air will be returned to the units at a higher temperature and closer to their adjustable set point, and thus supply air will automatically be delivered consistently at the design temperature of 55˚F. This temperature was assumed for all supply airflow in the model. Several iterations with different CRAC units operating and floor opening settings were simulated before a final run was created that maintains acceptable conditions. No server rack properties were adjusted in the model—it was simply an attempt to match cooling airflow to the existing data center load. The 17 CRAC units to be shut down in this scenario are indicated in Figure 3.

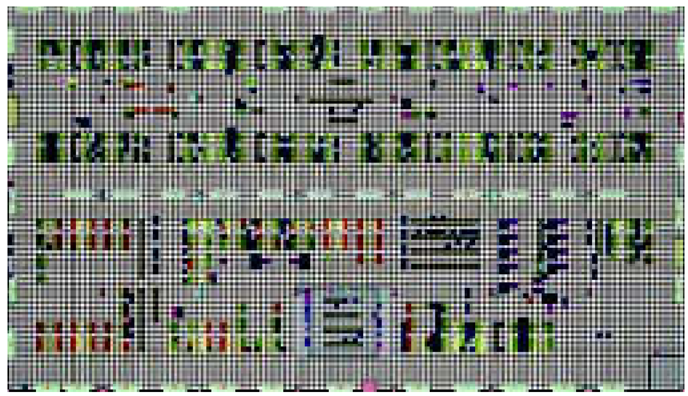

Figure 3

Next, a second scheme, called the "Airflow Management Scenario", was developed. This scenario includes recommendations for optimizing airflow through the server racks, including the installation of blanking panels at various locations and modifications to the front and back doors of selected cabinets. The best practices for blanking panel installation within the most typical cabinet type in the data center were developed using rack-level CFD models separate from the room models. The blanking panel best practices also allow for cooling

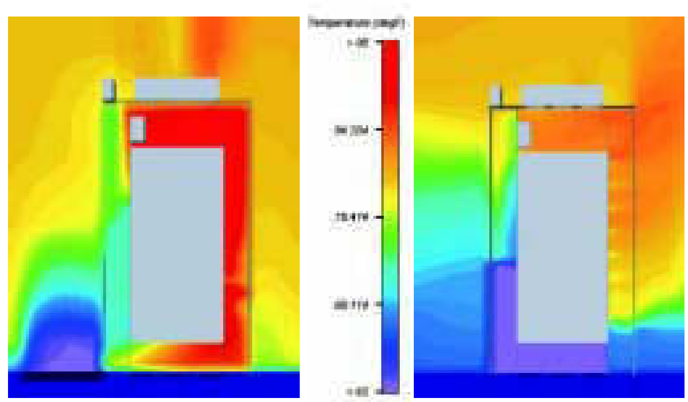

Typical server rack cross-section temperatures under existing conditions

Typical server rack cross-section temperatures after modifications

Figure 4

of a typical 5 kW-loaded cabinet directly from the raised floor plenum and without assistance from the server rack fan tray. Figure 4 indicates the temperatures inside the typical cabinet loaded at 5 kW before and after the best practice modifications. Most perforated tiles in the cold aisles near those racks were also eliminated, as that air will be short-cycled back to the CRAC units after the cabinet modifications are implemented. Further optimization was also pursued by reducing the number of perforated tiles at front-to-back airflow racks by matching the load with the appropriate quantity of tiles. Additional CRAC units to be shut down in this scenario are indicated in Figure 5.

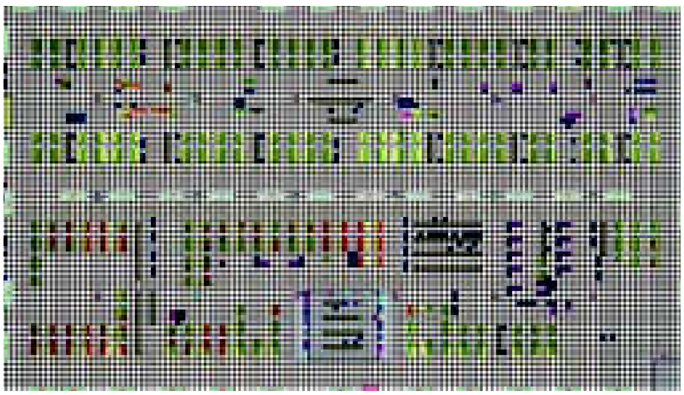

Figure 5

Note that all scenarios assume that the return air openings at the CRAC units are capped, so that no backflow occurs through the units. Also, at least one spare CRAC unit was included in all the scenarios and its performance in terms of raised floor location was confirmed.

CFD AIRFLOW MODEL RESULTS

In the Airflow Reduction Scenario, maximum rack inlet temperatures were held below 77˚F, with five racks exposed to inlet temperatures between 75 and 77˚F, 87 racks between 70 and 75˚F, and all other racks below 70˚F. This is within Environmental Specifications Class I as defined by ASHRAE in Thermal Guidelines for Data Processing Environments. The average temperature rise of cooling airflow is only 12.3˚F and total rack airflow is roughly 60% of total CRAC cooling airflow, indicating that some cooling air is still short-cycled. As a result of the airflow reduction, average underfloor plenum pressure will be lowered from 0.056″ w.g. to approximately 0.018″ w.g.

In the Airflow Management Scenario model, maximum rack inlet temperatures were held below 73˚F. Since it is assumed that blanking panels are installed in the most typical cabinet type, those racks have a bottom-to-top airflow configuration and therefore receive 55˚F cooling airflow at all intake points. 7 other racks are exposed to inlet temperatures between 70 and 73˚F, with all other racks below 70. The average temperature rise of cooling airflow is 18˚F and total rack airflow is roughly 90% of total CRAC cooling airflow, indicating that cooling air is used much more effectively than the airflow reduction scenario. As a result of the further reduction in airflow, average underfloor plenum pressure will be lowered to approximately 0.01″ w.g.

VARIABLE SPEED FAN CONTROL

A limitation of conventional CRAC unit cooling technology is that they are controlled according to return air temperature, which in a large data center is often the desired room temperature—not the hot aisle or rack exit air temperature on which the CRAC cooling capacity is often based on. But the most important factor in data center cooling is the rack inlet supply temperature, which is often independent of CRAC return temperatures, depending on the architecture of the space. A more appropriate method of control is to lock the supply air temperature and control the airflow so that it satisfies the cooling requirements of the racks.

An alternative method to realizing comparable fan energy savings is to implement a variable frequency drive (VFD) for each CRAC supply fan motor. For an existing CRAC unit, modifications are required to allow the VFD to be controlled based on return air temperature and the chilled water control valve to constantly provide a low supply air temperature independent from the cooling load. This allows for both a reduction in airflow and chilled water flow rate proportional to the data center load. VFD control will provide energy savings according to real-time data center load fluctuations, which depending on the load profile, may amount to much more savings than simply shutting units down based on the current allocated power in the data center.

OPERATING COST SAVINGS

For the case study room and the airflow reduction scenario, 17 selected CRAC units can be shut down, yielding an estimated annual energy cost savings of $154,000. With the airflow management scenario, 26 selected CRAC units may be shut down, yielding an estimated annual energy cost savings of $235,000 – a 57% reduction in data center fan energy cost.

The operating cost savings assume the sites' electric utility rate, with a consumption charge of $0.112 per kilowatt-hour and a demand charge of $17.33 per kilowatt. With these current rates and a power consumption of 7.6 kilowatts, each CRAC unit fan costs approximately $9,000 annually to operate constantly.

Depending on the first cost to implement the cabinet modifications and the removable CRAC unit return air inlet covers, simple payback for the energy reduction measures can be achieved within one year. The option of CRAC fan VFD control is a superior method of achieving the fan energy savings over simply shutting down and capping selected CRAC units, as this strategy will automatically compensate for fan failures and extreme changes in data center cooling load. Deployment of VFDs in a phased manner on only the CRAC units required for current loads conditions may also be considered.

CONCLUSION

This case study demonstrates the viability of fan energy reduction in an existing data center as well as the CFD modeling process. Because data center cooling systems operate 7x24x365, energy-saving initiatives usually result in significant operating cost savings, especially when maintenance costs are also considered. Cooling system performance with new IT loads may also be tested with the CFD models to confirm acceptable rack inlet temperatures, avoiding the trial-and-error involved with deploying equipment in the live data center. They are a valuable tool for assessing air management practices as well as creating a business case for various upgrades.

Michael Schwarz is an HVAC Project Engineer for KlingStubbins. He can be reached at mschwarz@klingstibbins.com.

MGE UPS SYSTEMS

Sure, you have the server capacity... But what about the power?

If you don't have adequate power availability and protection, your operations can come to a screeching halt!

Don't overlook the most important aspect of your data center infrastructure. Depend on **MGE UPS SYSTEMS** to help design and integrate power solutions that not only meet the current demand of power hungry servers, but are scaleable to accommodate future growth. MGE's experts in power infrastructure, UPS systems, sophisticated power management software and around the clock services make sure that you don't have unexpected capacity problems.

Need more info?

Visit **www.mgeups.com/dc** to register for one of our informative webcasts or for a free white paper on *Optimizing the Cost of Business Continuity.*

With over **40** years of global **UPS** experience, **MGE's** products and services provide the ultimate protection for enterprise-wide applications. Contact us today to learn how our power experts can safeguard your critical infrastructure.

800.523.0142

MGE
UPS SYSTEMS
www.mgeups.com

DESIGNING FAULT-DEFIANT ELECTRICAL POWER SYSTEMS FOR DATA CENTERS

by Jim Neumann

On October 4, 1957, 50 years before Fall 2007 7x24 Exchange conference, the American engineering community was jolted into the Modern Age when the former Soviet Union launched Sputnik-1, the first man-made object to orbit the Earth. Fearing that America was threatened by a shortage of engineering firepower, the U.S. Congress quickly passed the National Defense Education Act: in doing so, it also launched "The Space Race," the computer age, and millions of engineering careers.

For data center operators, a similar wake-up call came on August 14, 2003: in a mere eight minutes, the infamous Northeast Blackout moved the eastern seaboard like a tsunami, knocking out 250 power plants across a million square miles. Ten million people in the Canadian province of Ontario (about one-third of Canada's population) and 40 million people in eight U.S. states (about one-seventh of its population) were left without electrical power.

Geographically, these states are home to some of the world's most computing-intensive industries: financial services (New York), manufacturing (Ohio and Michigan), and telecommunications and healthcare (New Jersey). Of the entire FORTUNE 500, 319 companies are headquartered in states where the blackout took its heaviest toll.

For many of these companies, their data centers – even those protected by massive investments in on-site power generation, UPS, and fault-tolerant systems – were hit by surprise... and hit hard.

Figure 1: Hundreds of data center operators had a false sense of security about their resilience prior to the 2003 Northeast Blackout. (Photo: NOAA, the National Oceanic and Atmospheric Administration)

Data Centers: Corporations' Multimillion-Dollar Nest Egg

By the time power was restored, almost one-third of the nation's enterprise data centers had been affected by The Blackout. According to the data center research firm AFCOM, a survey of more than 500 data center executives found that nearly half suffered significant financial losses:

- Ten reported that they lost more than $10 million as a result of the outage
- Two reported losses of between $5 million-$10 million
- Fifteen reported losses between $1 million and $5 million
- Twenty five reported losses between $500,000-$1 million
- Fifty report losses of $100,000-$500,000
- One hundred twenty five reported loses of between $10,000 and $100,000

But for all of its notoriety, the Northeast Blackout was far from a one-time anomaly. More than 500,000 businesses and consumers experience electrical power problems every day, with the average power outage lasting two hours. The annual cost to the U.S. economy due to these power disruptions is estimated at $104–$164 billion... not including up to another $24 billion due to power quality phenomena.

According to the San Francisco Chronicle, the average public utility customer in San Francisco lost power for more than 4 1/2 hours in 2006, quoting statistics compiled by the utility and submitted to state energy regulators.

More recently, on July 25, 2007, a serious 11-hour power outage occurred in greater San Francisco. It affected major Websites –including Craigslist, Netflix, Technorati, Typepad and Gamespot – hosted at a 227,000 square foot commercial data center that was knocked offline, as much of the city blacked out.

According to Reuters, it took about five hours after the power came back on for Craigslist to be restored. And worse, according to the website Ecommerce Times, such major ecommerce sites can lose as much as $1 million per minute.

The lessons learned are simple but painful ones: interruptions in electrical power service can have costly, even catastrophic, consequences for data center operators, especially in facilities where they had invested heavily in disaster mitigation and recovery systems, and believed they were sufficiently prepared.

Figure 2: The cost of downtime varies by company, but the financial trade-offs to ensure maximum uptime is rarely understood until power problems take their toll.

The Weakest Links

Why do the most intricately-designed "failsafe" systems fail?

Often times the problem begins with the design itself. Rarely is it tested to ensure its operational resilience prior to construction; and following the completion of an installation, limited tools are employed to fully test how the "live" operation might deviate from "as-designed" performance.

This seems to be unique to the data center environment: in most other mission-critical applications – the space shuttle, aircraft, shipbuilding, oil platforms, etc. – an electrical design is fully modeled and a variety of simulations are performed using advanced modeling software to test the limits of the design prior to construction.

Further, as the installation matures, undocumented changes inevitably occur; e.g. equipment or circuitry is added to a critical bus. In a crises situation, these changes result in "surprises" that can disrupt the service or recovery process.

Another common problem is that many backup systems, once installed, are no longer maintained by the experts who designed them… so maintenance procedures may not always follow those prescribed by the designer. This can be exacerbated during service actions where vital electrical components are taken in-and-out of service by personnel who may not be intimately familiar with the overall design characteristics of the electrical infrastructure.

Failure to properly maintain backup systems – or more likely, to diagnose their preparedness for changing infrastructure and loads – can have serious ramifications: 40% of all business downtime is caused by electrical power problems… and 80% of those problems are rooted in internal causes (overloading, inadequate system updates or maintenance, failing to upgrade backup systems in response to equipment changes or human error)

For example, in typical installations, an acceptable startup time for an on-premise generator is several seconds from the time of failure. During those seconds, a series of UPSs and batteries – designed to carry the load for 15 to 60 minutes – will provide emergency power and "ride-through" until the generator starts. If the generator fails start in a given period of time, this back-up (in theory) provides for a controlled shutdown of the computer equipment.

But in order to provide the desired protection, UPS units must be properly maintained. One of the most commonly-neglected components is the UPS's battery: most batteries have a useful lifetime only a few years, and they lose their ability to hold a charge gradually over that time. So if a UPS started with one hour of runtime for the connected load, after a year, it may only provide 45 minutes, and so on.

The Skyrocketing Price of Prevention

Businesses dependent on their data center operations face skyrocketing costs to build and maintain the IT systems they need to meet demand. There are five primary factors driving the spike in the cost of operating a data center:

1. Demand, Real Estate Prices are up: According to the San Jose Mercury-News, data center vacancies in the area are "in the single digits." Prime office space in downtown San Jose costs about $2 a square foot per month… but data center rents are as high as $30 a square foot per month.

2. Construction prices are up: In 2006, the typical 50,000 square foot data center would cost approximately $20 million to construct, based upon a current rate of $400 per square foot. This cost is projected to at least triple – and some say, increase to $5,000 per square – foot by 2010, meaning that today's $20 million data center construction will cost at least $60 million three years from now.

3. Consolidated, higher-capacity infrastructure: According to Gartner Inc., 94% of IT departments are in the process of either consolidating their servers, or planning to do so. Randy Mott, CIO for Hewlett-Packard, has said publicly that he plans to reduce the number of HP data centers from 85 to six… a more than 90% reduction. Thus, organizations are putting all of their "eggs" into fewer and smaller baskets… raising the IT stakes of each newly-consolidated facility.

Sales of servers are estimated to increase 40% to 50% annually over the next four years, meaning that higher-density servers are being installed at a very rapid rate. Because server performance and density are more doubling every year, more and more IT resources are being squeezed into a smaller footprint… dramatically increasing the risk to reduced-site operations.

4. Utility consumption, prices are up: In the United States alone, data centers consume some where between 30 and 40 billion kilowatts of electricity annually. A typical data center consumes 15 times more energy per square foot than typical office space, and is estimated to be 100% more energy intensive, in terms of density. The more dense the server components, the more heat they generate, thus, the more air conditioning they require.

5. Obsolete infrastructure: It's not just older data centers that are being taxed by computing clusters. According to AMD, facilities built as recently as two years ago are ill-equipped to hold today's high-density server racks, due to inherent limitations in their power utilization and cooling output.

When data centers are constructed, they are subject to a physical limit as to how much power can be delivered to the site and stored. In 2006, the typical data center was designed to handle 40 watts of power per square foot; because of increasing server density, this is projected to handle up 500 watts of power per square foot by 2010.

Similarly, data centers' air conditioning systems are limited in how much cooling they can deliver in a given space. The problem is so pervasive that through year-end 2008, heat and cooling requirements for servers will prevent 90 percent of enterprise data centers from achieving anywheres close to their theoretical server density.

The Below-the-Floor Solution

As noted earlier, the most common cause of data center failures lies in the electrical power infrastructure supporting it. Most of these failures could be prevented in either of two ways:

1. Better up-front modeling prior to construction

2. Better post-construction diligence, diagnostics, and insight

Organizations in every industry tally the cost of downtime, and make calculated decisions about the cost-benefit trade-offs of reducing power-related disruptions by investing in power systems infrastructure.

But "throwing hardware at the problem" has not resulted in failsafe facilities; to the contrary, it simply introduced different points of failure, while creating a false sense of security that causes data center operators to stop probing for potential lapses in their infrastructure.

What is needed is a more ecosystemic approach… an approach many forward-thinking organization are calling "Power Analytics." First introduced at the **Spring 2006 7x24 Exchange Conference**, Power

Is Your Data Center Experiencing Growing Pains?

Unprecedented growth is causing big problems

The increasing power density of server racks has created a crisis at many data centers. They're running out of protected power from uninterruptible power supply systems. They're running out of computer room floor space too. Moreover, many IT operations have been given a mandate to cut energy costs and improve operating efficiencies, to help their facilities achieve LEED® (Leadership in Energy and Environmental Design) Certification.

S&C can help

S&C's PureWave UPS™ System provides total power protection . . . not just for your servers, but for critical air conditioning systems and mechanical as well. And the PureWave UPS System will help you attain your efficiency and energy-saving goals.

The PureWave UPS System is available in low and medium voltages and is expandable up to 20,000 kVA.

S&C PureWave UPS System rated 12,500 KW @ 12.47 KV protecting a large critical load

It's ideal for high-density data centers. It will let you meet Tier III or Tier IV electrical requirements as well as Class A or Class B cooling requirements. Only the PureWave UPS System can be installed outdoors without modification, saving valuable interior floor space and reducing cooling loads.

PureWave UPS System Ratings

System Rating		Voltage
kVA	kW	
313 to 2500	250 to 2000	208 to 600
2500 to 20 000	2000 to 16 000	4160 to 34 500

Let S&C show you how the PureWave UPS System can provide uninterruptible power for your data center at the lowest installed cost. Call us and we'll determine the system that's right for you.

The extraordinarily high efficiency of the S&C PureWave UPS System is in keeping with the goals of The Green Grid™, a consortium of information technology professionals and companies seeking to lower the overall power consumption of data centers and improve their efficiency. S&C supports these goals and is a member of this organization.

www.sandc.com/pwn

©2007 S&C Electric Company [illegible]

S&C ELECTRIC COMPANY
Specialists in Electric Power Switching and Protection
Offices Worldwide • 6601 N. Ridge Blvd., Chicago, IL 60626-3997 • Phone: (1) 773-338-1000 • Fax: (1) 773-338-3562

Since 1911

Analytics combines the best practices of "design" and "operations" into a single synergistic environment.

Using the original CAD model of the facility – which includes operating parameters and specifications down to the small component within the infrastructure – Power Analytics systems are able to diagnose and accurately predict the full capabilities of any given electrical power systems design.

Figure 3: The Power Analytics methodology is being adopted to protect private and public assets in a wide range of mission-critical facilities.

Further, once the system is installed and running, the "live" version of the Power Analytics system will detect the slightest variations within the infrastructure – say, a component that is slowing drifting from its normal operating parameter – and assess the downstream potential of that components deteriorating performance.

Since the performance of individual elements within power systems typically degrade with the passage time – due to aging, weathering, or stressing materials beyond their engineering limits (overloading, heat, etc.) – Power Analytics systems can identify and resolve potential problems with an extremely high degree of accuracy… typically, long before they reach the problem stage, and in sufficient time to effect repairs without business disruption.

A Crystal Ball for Power Systems Operations

While any number of "monitoring" systems have the ability to collect and manipulate data, none of these systems inherently have the capability to accurately predict and preempt disruptive electrical power problems; all require hundreds or even thousands of hours of programming following their installation in order to give the system usability that mimics "intelligence."

Power Analytics provides for an instant understanding of the systems intended design and limitations. This built-in knowledge is of extreme value by enabling the system to deliver a wide range of ancillary benefits that would standalone in their own right.

By making use of real-time operational data, this "intelligent" system makes accurate performance predictions and provides detailed operational insights into:

- Reliability – The trustworthiness of the system to perform as designed; the probability and frequency of failures… or more importantly, the lack of failures. Reliability metrics include probability of failure on demand (POFOD); rate of failure occurrence (ROCOF); mean time to failure (MTTF); and availability or uptime (AVAIL).

- Availability – The percentage of time that data can be instantly accessed, and that a system is available to deliver failure-free performance under stated conditions. The term is mostly associated with service levels that are set up either by the internal IT organization or that may be guaranteed by a third party datacenter or storage provider.

- Commissioning – The commissioning process is the first crucial step in new construction, consolidation and in some cases expansion. The model- based approach to Power Analytics can reduce the typical commissioning time from 25-50% not only providing huge cost and time reductions but for the first time guarantee that the design and the as-built drawings are in 100% agreement.

- Capacity – The storage and transaction processing capability of computer systems, the network and/or the datacenter. Capacity planning requires insightful forecasting, e.g. what if traffic triples overnight; what if a company merger occurs, etc. As a result of such the analyses and forecasts, systems can be upgraded to allow for the projected traffic or be enhanced so that they can be ready for a quick changeover when required.

- Configurability – The ease in which IT infrastructure and related systems can be maintained, upgraded, redeployed, and retired from mission-critical use. For example, understanding the effective life of specific hardware and software technology – while taking into account a systematic technology upgrade policy – allows companies to seamlessly 1) maximize the life of their technology investments, 2) incorporate new technologies, and 3) phase out older systems as their relative performance warrants.

- Energy Management – Energy management is a growing concern for mission-critical facilities, especially data centers, where the cost of powering IT equipment equals 20% of its cost. Because of the ever-increasing density of servers and switches, power requirements for the largest data centers are increasing 20% per year; some data center operators report that their monthly power costs have doubled, tripled, and even quadrupled in the past 36 months.

- What-if Simulations – Power Analytics systems create virtual environment that provides an off-line, mirror image of them. It enables users to make a "freeze frame" of their real-time, current environment in order to conduct detailed "what if" simulations reflecting the present configuration of your electrical infrastructure. Such simulations include testing of real-time configuration, maintenance, repair, and other procedures, before attempting them on live systems.

Summary & Conclusions

"Fault tolerance" is no longer an acceptable measure of success in mission-critical data centers.

The technology exists today – and is proven in world-class facilities – to preempt business-disruptive electrical power problems altogether… ensuring business continuity under the most extreme operating conditions.

That's not to say that companies should abandon their investments in UPSs, on-site power generation, etc., e.g. no driver would remove the airbags from his car, simply because it comes with new collision avoidance technology.

But it's important for data center operators to know that technologies like Power Analytics hold the promise to ensure that such emergency technologies are minimized as a means to ensure day-to-day operational resilience.

Jim Neumann is Vice President of EDSA Micro Corp. He can be reached at jneumann@edsa.com.

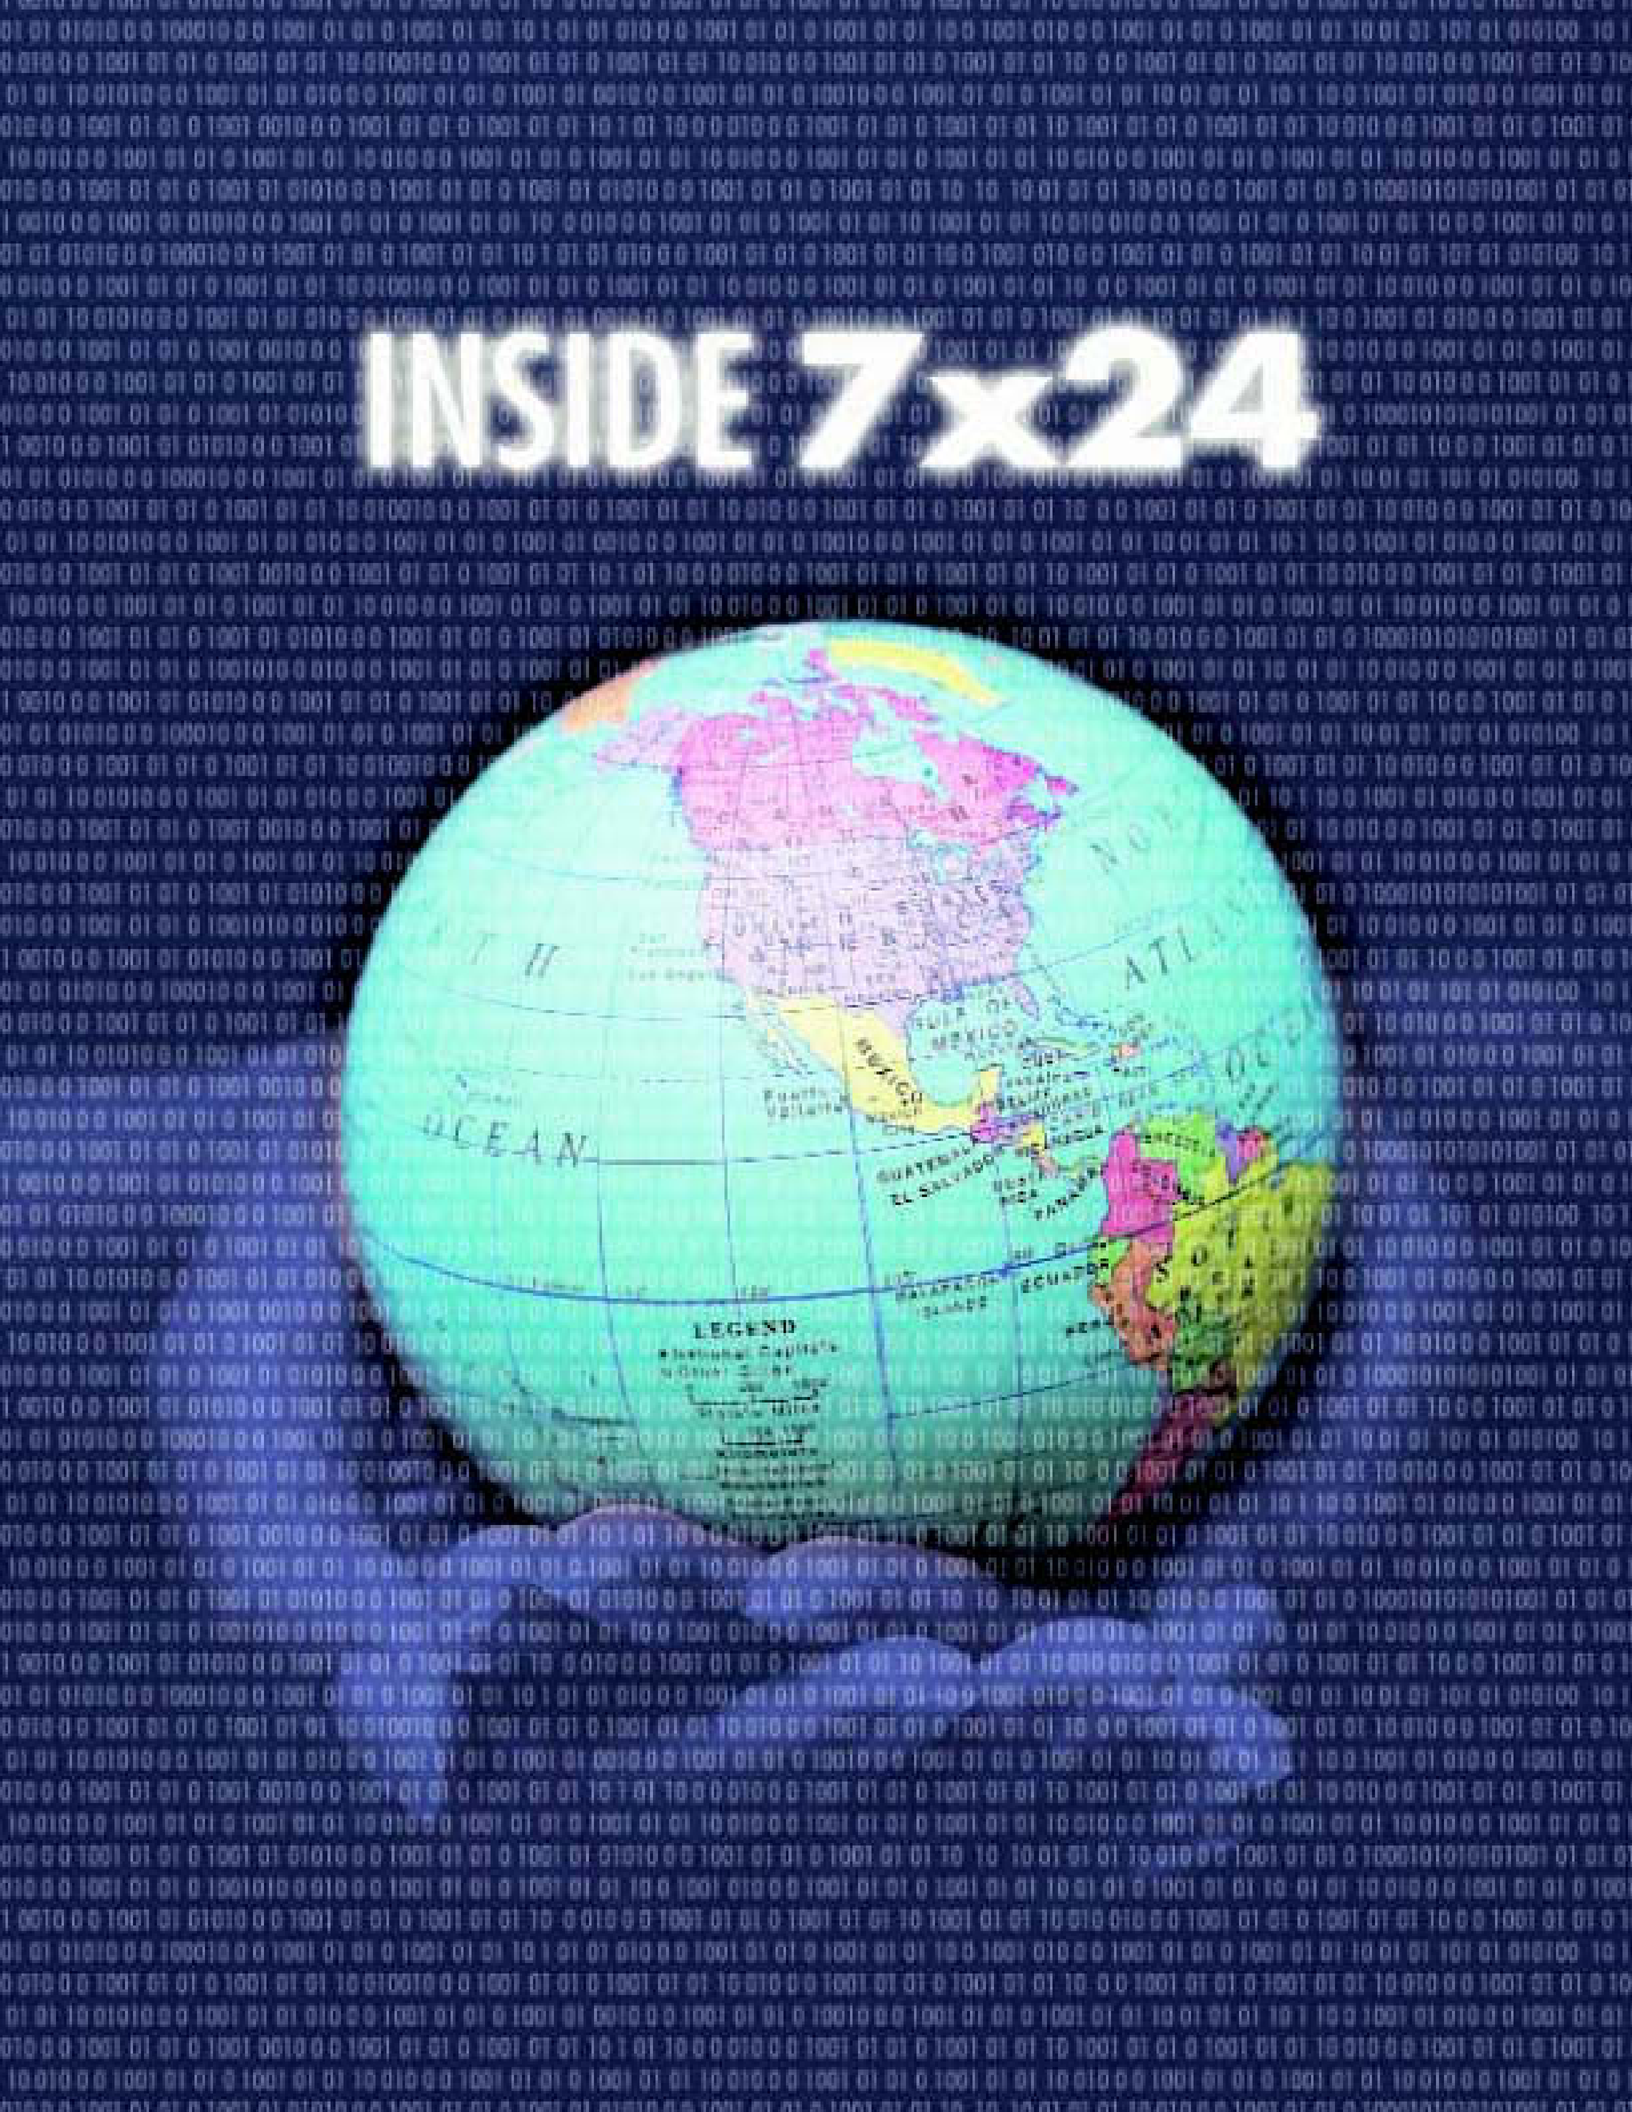
INSIDE 7x24

2007 Fall Conference Highlights

The Fall Conference themed "End-to-End Reliability: Mission Critical Systems" will be held October 28-31 at the Gaylord Texan in Grapevine, Texas. The Fall Conference will feature compelling keynotes, concurrent and tutorial sessions, sessions on the greening of data centers, a spectacular vendor event, and more...

The opening conference keynote address entitled ***"Failure is Not an Option"*** will be delivered by **Gene Kranz, former NASA Director of Mission Operations**. The Apollo 13 crisis pushed Kranz and his team to the brink of fear and doubt, but they refused to give in to these emotions or to succumb to panic. Instead, under his leadership, they worked together to save the imperiled spacecraft, and brought the ship and crew safely home. Kranz will share the challenges and problems that they successfully faced and overcame, giving you the inspiration that they need to face down challenges and adversity in their own lives and careers.

Gene Kranz

The conference will feature various sessions on energy efficiency and the greening of data centers with topics such as:

- *Peeling the Green Onion: What Does Green Really Mean for Data Center Power?*
- *Next Generation Data Center Efficiency Ideas and Considerations*
- *US Government Programs to Advance Data Center Energy Efficiency*
- *Can Mission Critical Servers & Storage Be Green?*

For the complete Fall Conference program and registration information please visit 7x24exchange.org or call (646) 486-3818.

In addition to enhanced programming 7x24 Exchange International presents Roadhouse Ranch, a rodeo and biker barbecue. There is no better place than Texas to get a custom tattoo, ride the mechanical bull, try your hand at Texas hold 'em, black jack, roulette and experience a full blown live rodeo while you chow down on Circle R's authentic "Texas Road Kill Buffet". Upon arrival at the Ranch our attendees will be greeted by a line up of custom choppers and Harley's that would make anyone's head turn. All this while the Roadhouse Ranch band sings your favorite country songs. Never fear we did not forget it is Halloween Eve so the Circle R "Haunted Maze" will be open for anyone who dares... *Special thanks to the following organizations for making this event possible:* ABB, APC/MGE, ASCO Power Technologies, Chloride, ComRent, Cummins, Cyberex, Data Aire, Data Space Advisors, Eaton, EDSA, Emerson Network Power, Enviroguard, Kling Stubbins, Kohler Power Systems, Mitsubishi Electric, MTU Detroit Diesel, PDI, Russelectric, SIEMENS, Starline Track Busway, TAC.

Attendees network at the 2007 Spring Conference.

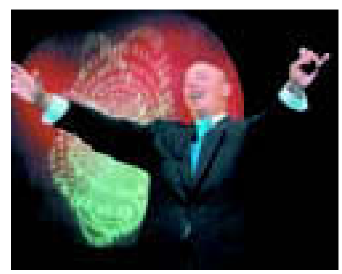

Fran Dramis, former EVP and CIO of BellSouth Corporation, delivers the keynote address at the 2007 Spring Conference.

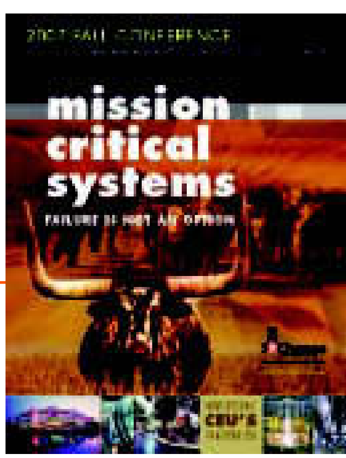

Fall 2007

End-to-End Reliability: MISSION CRITICAL SYSTEMS

October 28-31, 2007

Gaylord Texan
Grapevine, TX

For information about sponsoring a 7x24 Exchange event please contact Brandon Dolci, CMP at (646)486-3818 x108

SAVE THE DATE!!!

SPRING 2008
June 1-4, 2008
Boca Raton Resort & Club
Boca Raton, FL

FALL 2007
November 16-19, 2008
JW Marriott Desert Springs
Palm Springs, CA

Interested in presenting at 7x24 Exchange? Download the Call for Presentations www.7x24exchange.org or call 646-486-3818 x104

CORPORATE LEADERSHIP PROGRAM MEMBERS *(at press time)*

2007 Fall Conference

End-to-End Reliability: MISSION CRITICAL SYSTEMS

Silver Partners

ABB

ACTIVE POWER

APC MGE — Critical Power and Cooling Services

PDI — creating the perfect wave

CATERPILLAR®

Cummins® Power Generation

mtu

DETROIT DIESEL

EATON

SIEMENS

SQUARE D

SYSKA HENNESSY GROUP

Key Partners

ASCO Power Technologies®

CHLORIDE

ComRent® Load Bank Rentals

CYBEREX

DATA AIRE INC.

DataSpace ADVISORS

EDSA

EMERSON Network Power

ENVIROGUARD — COMPLIANCE, SAFETY & SECURITY SOLUTIONS

KLING STUBBINS

KOHLER POWER SYSTEMS

MITSUBISHI ELECTRIC UPS Division

Russelectric — POWER CONTROL PEOPLE YOU CAN RELY ON

SKANSKA

STARLINE TRACK BUSWAY

t.a.c®

Media Partner The DATACENTER Journal

INTERESTED IN PRESENTING AT THE SPRING CONFERENCE?

Visit www.7x24exchange.org and download the Call for Presentations

Deadline: January 11th

SUBMIT AN ARTICLE FOR THE SPRING NEWSLINK

Visit www.7x24exchange.org and download the Call for Articles

Deadline: January 9th

BECOME INVOLVED IN YOUR LOCAL 7x24 CHAPTER

INTERNATIONAL CHAPTERS INCLUDE:

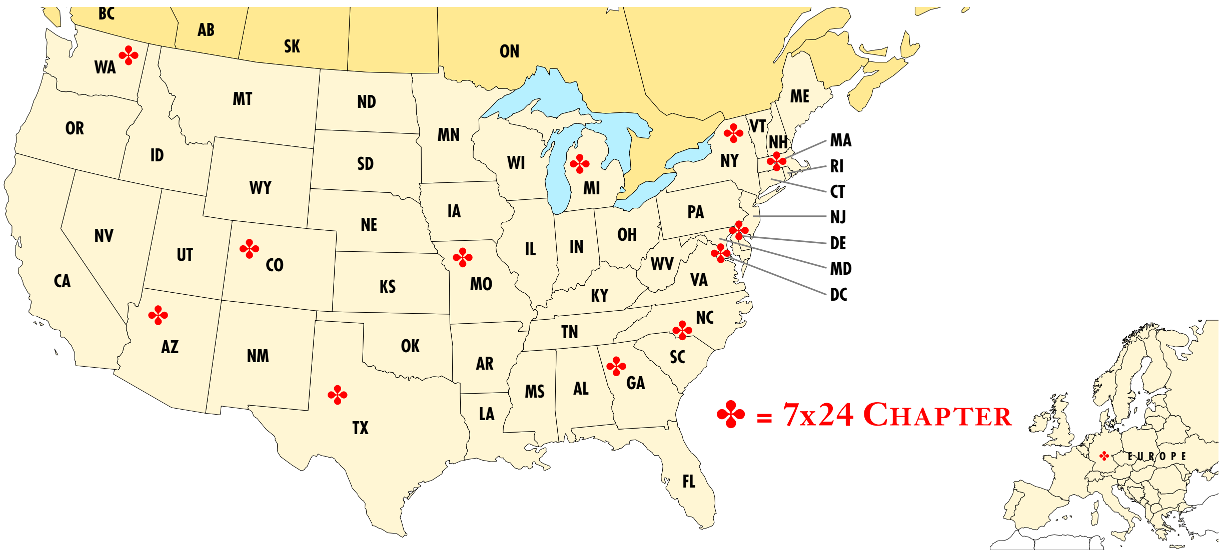

- Arizona
- Atlanta
- Boston
- The Carolinas
- Delaware Valley
- Europe (London, UK)
- Lake Michigan Region
- Lone Star (Dallas)
- Midwest (Kansas City)
- Metro New York
- Northwest (Seattle)
- Rocky Mountain
- Greater Washington DC

Visit www.7x24exchange.org today to participate in your local chapter!

7x24 Exchange International

The end-to-end reliability forum.

THE PREMIER ORGANIZATION SERVING THE UPTIME INDUSTRY

MEMBER FORUM:

Post ideas...solutions...ask industry related questions!
Both end users and vendors are welcome to use this discussion forum that was built with the 7x24 Exchange Mission in mind, to promote end-to-end reliability among those individuals who design, build, use and maintain mission critical enterprise information infrastructures.

CAREER CENTER ONLINE:

7X24 Exchange Career Center is the premier electronic recruitment resource for the industry. Here, employers and recruiters can access the most qualified talent pool with relevant work experience to fulfill staffing needs.

- **POST A JOB**
  Target your recruiting and reach qualified candidates quickly and easily. Simply complete our online Registration Form and start posting jobs today!

- **POST YOUR RESUME**
  Whether you're looking for a new job, or ready to take the next step in your career, we'll help you find the opportunity that's right for you.

Visit www.7x24exchange.org and Take Advantage of These New Services Today!

NEWSLINK OPPORTUNITIES

Following are the Editorial Guidelines for Newslink together with the Member Advertising Rate Card. Advertisers interested in placing an ad may fax the insertion order to 7x24 Exchange at 212.645.1147 or email to jeremy@7x24exchange.org. Questions? Please call Jeremy O'Rourke at 646.486.3818x109.

NewsLink

Member Advertising Rate Card

BLACK AND WHITE RATES

Size	1X	2X	3X
Full Page	$1,500	$1,300	$1,100
2/3 Page	1,100	1,000	900
1/2 Page Island	900	800	700
1/2 Page	700	600	550
1/3 Page	600	550	500
1/4 Page	500	450	400

COVERS & PREMIUM POSITIONS – INCLUDES 4 COLOR

Size	1X	2X	3X
DPS	$5,000	$4,500	$4,000
2nd / 3rd Cover	2,500	2,200	2,000
4th Cover	3,500	2,750	2,500

COLOR RATES

Process Color (4/c)	$900
PMS Colors (add per color)	$600
Process Colors (add per color)	$500

Revisions and Proofs: $50

Position Guarantee: 15% premium

*Non-Members add 40% to all rates

NON-BLEED AD DIMENSIONS

Size	Width	Length
Full Page	7.5″	10″
2/3 Horizontal	7.5″	6.5″
2/3 Vertical	5″	10″
1/2 Island	4.875″	7.5″
1/2 Horizontal	7.5″	4.875″
1/2 Vertical	3.625″	10″
1/3 Horizontal	7.5″	3.25″
1/3 Vertical	2.5″	10″
1/4 Horizontal	4.5″	3.25″
1/4 Vertical	3.25″	4.5″

8 1/2″ x 11″ MECHANICAL REQUIREMENTS

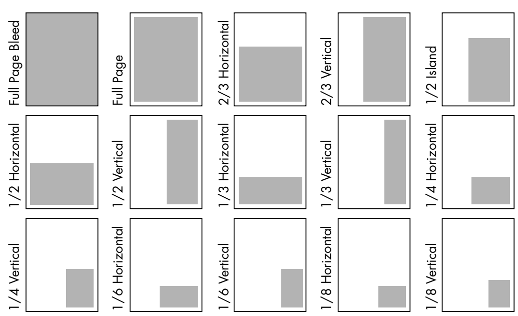

Live Area: 7.5″ x 10″

Trim Size: 8.5″ x 11″

Bleed Size: 8.75″ x 11.25″

Halftone Screen: 133 lines up to 150 lines

DPS Mechanical Requirements:

Live Area: 16″ x 10″

Trim Size: 17″ x 11″

Bleed Size: 17.25″ x 11.25″

Halftone Screen: 133 lines up to 150 lines

Advertiser indemnifies 7x24 Exchange against losses or liabilities arising from this advertising. 7x24 Exchange assumes no liability whatsoever, except to the extent of a one time paid advertisement of the same specification, in the next or similar publication, if any proven or admitted errors or omissions have occurred. Payment is due upon receipt of the invoice. Interest shall be charged at 2% per month compounded to yield 26.82% per year on overdue accounts. Revisions to previously submitted ad copy are subject to additional charges. A charge of $30.00 will be levied for returned checks. In the event of a contract cancellation, the advertiser or agency agrees to repay 7x24 Exchange any discounts granted for multiple insertions less any discount applicable for the number of insertions completed in the contract. All cancellations must be received in writing prior to the advertising sales deadline. All premium positions are non-cancelable. Prices are net of agency commission.

EDITORIAL GUIDELINES FOR NEWSLINK

Manuscript specifications: Feature articles vary in length from 500 to 2,000 words. While Newslink accepts articles in a variety of formats, it prefers to receive materials on CD. All articles must be received by the deadline to be considered for a specific issue. Material submitted after the deadline will be considered for the following issue.

Bylines: All articles should include a brief (1-2 sentence) author biographical sketch at the end of the article, that includes the author's name, title, affiliation, address, and phone number. Photos of authors are never used. Newslink does not pay authors for contributions.

Visuals: Authors are encouraged to submit photographs and charts, graphs, or other illustration that will help readers understand the process being described, though it does not guarantee that visuals will be used with the article. Submit all charts, graphs, and other artwork separately; do not incorporate them in the body of the article. Indicate caption material separately. Newslink reserves the right to publish submitted visuals.

Editorial procedures

All articles are reviewed for suitability. Accepted materials are then edited for grammar and to conform with Newslink's editorial style. All attempts are made to preserve the author's writing style, however, Newslink has the right to edit for style, clarity, and to fit space allotments, and to make final selection on headlines, subheads, and graphic treatment. Manuscript submission implies author agreement with 7x24 Exchange's Editorial Policies.

Copyright

Newslink requires first serial rights for submitted articles. This means the author(s) grant Newslink the right to publish the article for the first time. We also request permission for electronic distribution on 7x24 Exchange's web site, www.7x24exchange.org.

Disclaimer

The responsibility for accuracy remains with the author. The opinions and information in bylined articles in this publication are those of the authors and do not necessarily reflect those of the Officers and Board of Directors of 7x24 Exchange.

There are some things that everyone knows should happen every day...

Some things are not so obvious...

...like monitoring your batteries for ohmic value once a day. It is obvious a system should be quick and easy to install and work reliably once in. It is obvious that it should be able to measure all your batteries (even your generator batteries). So, when you think about it, it should be obvious that something as critical as your UPS battery should be measured every single day for ohmic value. Cellwatch was designed from the start to do just that. Cellwatch uses the lightest measurement "footprint" on your battery so that it can safely check the ohmic value of everyone of your batteries, every single day. It then shows you the information in a way you can actually understand. What does this mean? No days without a functional battery. Seems obvious to us.

NDSL Inc.
5800 McHines Place
Suite 120
Raleigh, North Carolina 27616
USA

Phone: 919 790 7877
Fax: 775 535 0139

NDSL Ltd,
Oakfield Industrial Estate,
Eynsham, Witney OX29 4TS
England

Phone: +44(0) 1865 884288
Fax: +44(0) 1865 884289

Cellwatch - We do it daily.

NDSL

www.cellwatch.com