

# EQUIPMENT PROTECTION

Enclosures • Thermal • Shielding • Power • Equipment Reliability  
Shock/Vibration • Contamination • Noise • Preventative Maintenance

www.EquipmentProtectionMagazine.com

Serving OEM Design  
Engineers, System Engineers,  
And Other Professionals  
Involved in Equipment  
Protection Technologies  
And Testing.

October/November 2005  
Volume 3, Issue 5  
a webcom publication

## Scalable Data Center Infrastructure It Starts With A Rack

**Martin T. Olsen**  
Product Director  
Wright Line LLC


Successful businesses in today's global economy have a clear commitment to their information technologies and operating systems. These mission critical applications directly impact the productivity of almost every function in the organization. Therefore, today's networks must operate at peak performance on the enterprise level, within workgroups, and at each individual workstation for businesses to deliver results. Accordingly, today's data centers must be organized, maintained, and remain adaptable to advancements in technology.

On a daily basis, new equipment is deployed in the data center, resulting in rapidly increasing power consumption and escalating heat output. Data center managers are constantly moving, adding and changing IT equipment and electrical circuits - all while trying to reduce or maintain costs. This rate of change, combined with the pressure to control costs, often forces data center professionals to do more with less.

Successful data centers utilize enclosure systems that accommodate expansion and effectively manage equipment migrations and upgrades. Conservative estimates place costs of designing, building and maintaining data center environments at four times the cost of similar office space. Therefore, data centers must make the most effective use of available floor space and data center managers must look at standardization of enclosure platforms to ensure minimal and cost-effective changes on the raised floor.

Today, the data center is the central nerve system of the organization.

Without this piece of precious real estate the company would simply not exist. One only needs to consider Internet Service Providers (ISPs) and e-commerce companies to realize this fact. If the servers in an e-commerce company are down the company vanishes off the face of the earth until connectivity to the web is re-established. In fact, 43 percent of companies experiencing data center disaster outages never reopen and 29 percent close within two years.'



**Wright Line's Paramount Enclosure System can be custom configured to address almost any storage, power, thermal and security need.**

Keeping the dependency on the data center in mind emphasizes the importance of deploying, maintaining and running it with the utmost care and consideration. To help IT professionals accomplish this task the industry has stepped up to the plate and many professional associations like AFCOM, 7 by 24 Exchange and ASHRAE have emerged or expanded their reach into this industry over the years. Additionally, the importance of the data center to today's businesses has increasingly become the primary editorial focus of facilities and IT publications alike. From Equipment Protection and other IT focused publications to facilities-based publications like

ARCHI-TECH, Building Operating Management and Today's Facility Manager, these magazines are bridging the gap between facilities and information technology professionals.

As a result of all this emphasis and press, many companies and manufacturers of data center products have developed and enhanced technologies and solutions that help drive availability up and complexity down. However, although many of them tout open platforms and infrastructure that will help the customers achieve flexibility and scalability, many solutions are inherently proprietary and targeted to "lock-in" the customer ensuring a continuous revenue stream for the supplier for years to come.

### The Fluid Data Center

Infrastructure solutions with labels such as "pre-packaged" and "pre-configured" scream the word proprietary. Although pre-configured insinuates standardization, it is standardization within its own design and structure. Add third-party components to the mix and you immediately start cutting features and benefits out of the solution. A true scalable infrastructure is open and adapts effectively with any third party equipment regardless of the manufacturer. In the operating system world, Linux is a great example of an open platform that adapts to the applications the user wishes to run regardless of the software developer. Only a few manufacturers of enclosures and data center infrastructure solutions can honestly claim this type of third-party compatibility. Just as liquid cooled solutions have come full circle, it appears that the argument for OEM server manufacturers to accept third party racks is making a come back too -

Reprinted from the October/November issue of Equipment Protection Magazine

©Webcom Communications Corp., 7355 E. Orchard Road, Suite 100, Greenwood Village, CO 80111, U.S.A, Phone 720-528-3770

for the simple reason that they offer much more scalability and are not tied to any specific equipment or vendor.

Moves, additions and changes in the data center are happening on a daily and sometimes hourly basis and a great deal of resources are spent on things that should have only been minor modifications like that of a simple server swap. Instead, it turns into larger changes in the electrical wiring, in-rack power distribution, thermal management and racks because the thought of switching from one OEM server or network equipment provider to another was not considered when the data center was originally designed. Although changes can vary in size from the simple addition of a port to accommodate a new employee to a complete network upgrade across the enterprise, the important thing to note is that change is a constant on the data center floor and increasingly IT professionals are finding that what looked like a great idea 12 months ago is now hopelessly outdated.

This is the situation more and more IT managers find themselves in when procuring their rack, power and cooling products from server OEMs. Server OEMs are successful companies because they build great network and server products - not physical infrastructure. However, server manufacturers view the racks, power and cooling products as additional revenue that will help them increase the average selling price to their customers. Conversely, many IT professionals look to the server manufacturer as a one-stop shop and try to simplify their busy lives by purchasing the "pre-configured" package. This approach was viable three to five years ago when space, power and cooling were unlimited in resources and equipment density levels were much lower.

Today the outlook looks dramatically different and customers are often asking for a minimum of 8 to 10 kW of cooling capacity, possibly more, per enclosure. OEM racks were never designed to meet these requirements, not to mention the data center they were deployed in. It is important to take the extra time and do your due diligence when designing or retrofitting your data center. Be honest about the fact that changes will happen although it may not seem likely at the time of deployment and be prepared to defend your position with management. When choosing the right rack solution it is obviously important to make sure the server fits in the enclosure but it is equally important to focus in on power, thermal management and security features and benefits. This is where only a few companies stand out.

### Feeling the Heat

The prominent question of how to effectively cool the rising heat loads in today's data centers is the subject of a lot of discussion and it ranges from the simplest configuration concepts up to the \$25,000 to \$30,000 per liquid cooled enclosure solutions. Solutions that promise a lot in terms of thermal management, but options that literally change the world of the IT manager forever. A key concept in the 1970's for the cooling of mainframes and supercomputers, liquid cooling solutions are the latest buzz in many IT circles as we see more centralized computing and consolidation within data centers.

Before you throw in the towel and decide to completely re-design your existing data center due to increasing heat loads, consider your options carefully. The notion of thermal management can be very daunting and it is a very intangible concept. It is not as cut and dry as power management. With power you know that if you pull a 30 Amp feed you can supply 30 amps of load on the other end. When it comes to cooling the data center you may be feeding the equipment loads with 20 tons of cool air, but still you can hardly keep ten tons of heat load cooled on the other end. Why is that?

For a long time the theory on cooling has been 160 CFM of chilled air per kilowatt of heat in the data center. This number was established by, and is commonly used by precision cooling suppliers and consultants in the industry. This is the number they use to size their CRAC units and determine their heat loads per enclosure, but it is not necessarily very accurate.

If you study the dynamics of a data center from a thermal management standpoint it doesn't make sense to only look at cold air distribution versus heat removal. Simply providing cold air feeds to the data center is not the answer to your thermal management issues when you consider these factors:

- Obstructions under the raised floor creating non-uniform airflow through floor tiles.
- Re-circulation of hot air from the rear of the enclosure to the front of the server intakes.
- Lack of a hot/cold aisle strategy where enclosures are all facing the same direction, "drinking" each others heat exhaust.
- Stratification across the front face of the enclosure that leads to failures in top mounted network and server equipment.
- Exhausted heat from the rack enclosure that never makes it back to the CRAC units.

The real key is heat removal from the enclosure as well as the data center itself. Consider heat as the by-product or waste from the whole process. If the heat is not removed effectively you can continue to feed as much cold air in the room as you wish and you will still run into unacceptable heat levels within the enclosures. Think of it this way, if you can remove all the heat generated by each enclosure in the data center you won't have any thermal management problems because in an ideal world the cold air distribution would always be at a constant 68 to 77°F.

### Cooling Down

So, now that we know heat removal is the primary factor in keeping the loads in optimal condition and avoiding hot spots that lead to downtime, there are still some common misconceptions. The most common, but not necessarily the most effective, used air removal solutions look to the temperatures in the rear cavity of the enclosure. Using simple math, vendors figure that with a 20°F temperature rise across the server combined with the original 68 to 77°F inlet temperature, the exhaust should be kept around 88 to 97°F.

However, temperature can be a dangerous factor to rely on. It is very common that different server brands and models are deployed in the same enclosure. One server may have been designed for a 20°F temperature rise while another in the same enclosure may have been designed for 30°F. The dilemma starts right there. Which temperature do you rely on to actively remove the heat? If you choose the lower, then you impede the internal thermal dynamics of the circuit boards in the one with the higher temperature rise possibly overdriving its fans. On the other hand if you follow the higher temperature rise then the one with the lower will run too hot. If you think this throws the calculation for a loop then consider the fact that some of them push 30 CFM at a temperature rise of 20°F and others push 45 CFM at 30°F.

The proper way to look at heat removal is volume of air moved, not what temperature you think you want to keep it at. Again, the temperature rises across servers vary from as low as 20°F in multi-U servers up to 40°F or more in newer blade server technology.

You can keep each enclosure in perfect thermal condition if you can figure out exactly how much air is exhausted from the rear of the enclosure and accurately and effectively push this heat into a close-coupled system such as a duct or return plenum. Obviously, the most accurate method of doing this at this point is a

close-coupled heat containment enclosure. Typically, this type of enclosure will feature a solid rear door, bottom and side panels with some sort of gasket mechanism to keep the inside air and outside air separated from each other. A variable fan speed system is deployed in the top for practical reasons and a duct provides a sealed coupling to the return plenum that sends the hot air back to the CRAC units.



**This graphic represents a data center layout in a hot-aisle/cold-aisle configuration with active heat removal rear doors that exhaust heat high up and close to the CRAC avoiding re-circulation.**

Consider the analogy of trying to fill up a bath tub with 70°F hot water. Do you monitor the temperature only and continue adding water in the bath tub until the mix is just right and you hit the 70°F mark regardless of whether the tub overflows spilling over its sides?

In an enclosure solution equipped with variable speed fans and effective monitoring systems you're assured that only the amount of hot air exhausted from the IT loads is removed. A monitoring system would actively look at the static pressure in the contained cavity in the rear of the enclosure and throttle the fan system accordingly. This accomplishes two things. First it removes nothing more or less than the air exhausted by the loads and secondly given its close-coupling with the return plenum there are no risks of re-circulation.

#### **Hold Your Water...for Now**

From a cost standpoint, a drop ceiling, if feasible, is far more cost-efficient than multiple chilled water feeds to each water-cooled enclosure in a data center. Some may argue that the cooling capacity of a water-cooled enclosure is much higher than air cooled, but that is not entirely true. Theoretically, the water-cooled enclosure works the same way an air-cooled enclosure does. The difference is in the placement of the chilled water supply. In a water-cooled enclosure the heat exchanger in the bottom or top takes care of the heat removal at a local, yet

expensive, level. In an air-cooled enclosure the exchange happens at a central, and cost effective, level.

The heat exchanger in the bottom of a water-cooled enclosure provides cold air to the front and removes heat from the rear, the exact same concept as an air-cooled solution.

So, what's the big difference? The difference is in the effectiveness. An open air-cooled enclosure environment does not perform that task very efficiently, but as soon as you couple the exhaust to the return plenum in the drop ceiling only the speed of the fans, same as in a water-cooled enclosure, dictates the amount of heat you can remove.



**This graphic represents a data center layout in a hot-aisle/cold aisle configuration with active heat removal rear doors close-coupled with the return plenum through collars achieving double-digit kW heat removal.**

#### **Conclusion**

All the current research indicates that today's leading data centers are in the planning stages for their next generation data center upgrades, re-designs and/or re-builds to address the increasing power and thermal management demands of high density technologies. When considering your space planning issues start with the enclosure and build from there. Deploying a scaleable rack platform that lets you effectively store, cool, power, manage and secure a number of different server brands and models will make all the difference when server and network moves, additions consolidations and changes come around. And, let's face it, as well as you plan, the technology will continue to change and evolve and modifications on the data center floor will be necessary.

With cooling or thermal management being, by far, the biggest challenge in the data center it is all the more important that your enclosure platform is capable of being scaled from a low 0 to 3.5 kW to more than 5 to 8 kW and possibly up to double-digit heat loads per enclosure without having to change

the primary platform. Deploying and re-deploying servers and cabling in enclosures could easily run \$50 to \$75 per component adding up to thousands of dollars for every enclosure change out.

Recovery rates are important to any cost conscious IT manager. The more equipment he/she can recover during moves, additions and changes the better the initial investment. Although the initial investment of a scaleable, multi-vendor compatible enclosure is typically slightly higher than that of a non-scaleable platform, it will pay dividends in the end. Apart from long term savings it also offers an unprecedented level of standardization that drives out waste and mistakes on the data center floor. Regardless of the equipment, the enclosure always stays the same and you know the platform will accommodate whatever equipment you're purchasing. In fact with some enclosure platforms it's guaranteed!

Seek an enclosure platform that can easily adapt to evolving technology requirements and provide a "swappable" solution to the various server and network equipment transitions and upgrades. By doing so, your data center personnel can concentrate on their core competencies including the development of information technology applications that will have a direct impact on the overall business' bottom line. That's a scenario where everyone wins.

*\*Source: Disaster Recover Journal, Spring 2001*

*Wright Line, a designer and manufacture of consoles, enclosures, office and other specialty furniture, has been serving the data center, office, call center, electronic laboratory and high-tech manufacturing environments for over 70 years. The company is headquartered in Worcester, MA and can be contacted by email, [info@wrightline.com](mailto:info@wrightline.com) or by phone 800.225.7348. To view the company's full product line visit [www.wrightline.com](http://www.wrightline.com).*



**The use of blanking panels on any open U space optimizes airflow patterns and significantly reduces re-circulation of hot air.**