



Tuesday, 22 September
2009

Design Best Practices for Data Centers

Written by Mark Welte

Tuesday, 22 September 2009

The data center industry is going through revolutionary changes, due to changing market conditions and technological demands; these include the inability to secure large sums of capital dollars to fund data center projects, the high cost of energy, and the pursuit of a “low carbon diet,” in the wake of recent cap-and-trade legislation.

But these are not the only issues driving these transformations. Vital new information technologies – such as cloud computing and virtualization – force fundamental changes in the design, operation and planning of a data center. On the other side, the data center containerization industry (DCCI) aims to have the capital investment associated with power and cooling infrastructure be a direct match with computing. Even the UPS industry is in competition for the most efficient and scalable products, with containerized power and cooling solutions to consider. Having a firm design foundation that can adapt to change is paramount to a facility’s success.

Start with the Design Objective

Whether it’s new construction, a facility upgrade, or operating an existing mission-critical facility, it is important to carefully plan the work and to work the plan. In the design phase, it is crucial to consult with design teams experienced in the design of the particular type of mission-critical facilities that your organization requires.

Without this expert input, the desired design is too often not fully-developed or is too vague to specific design and business objectives. A well-created design document will start with a narrative of the project background and provide a detailed discussion of the required objectives and goals. Required functional uses and requirements of the mission-critical facility should be clearly explained. While often difficult to quantify, performance and maintenance criteria should also be addressed. Finally, the expected lifespan and overall quality of the project needs to be defined.

A comprehensive plan is crucial for all involved parties to succeed in achieving maximum uptime in a mission-critical facility. Without it, too much interpretation by the various teams involved, can lead to compromises and miscommunications that weaken the integrity of everything that follows. Upfront planning can be tedious, exacting work; however, it is crucial in ensuring the finished facility will turn out precisely as intended, both initially and in the long term.

Fundamental to the successful design and maximum uptime of the facility is how the energy and cooling systems are handled. And now with green initiatives – driven by both public and political pressure – getting this right can be far more complex. However, there are new tools that significantly aid in proper design and also help to predict performance behaviors, in order to maximize the highest level of uptime while incorporating energy-efficiency strategies.

New design tools to the rescue

When an electrical power system is designed, it is routinely created by a power systems engineer using any of a variety of commercial, off-the-shelf CAD software (like AutoCAD) to create a design. High-end CAD modeling tools are significantly enhanced, making it easier to achieve the optimal integrity and safety of complex electrical infrastructure. As examples, they can simulate and help to resolve issues like arc flash, power flow, power quality, protective device coordination, and dozens of other power considerations iteratively enabling the engineer to produce a model that is “perfect on paper.”

Historically, the CAD model was printed out as an electrical, one-line diagram and “thrown over the wall” to a construction firm that, at the completion of the project, archived the drawings in printed and electronic form. But despite the wealth of information and professional expertise encoded in these models, they historically served no future purpose – unless, of course, operational problems occurred that sent everyone “back to the drawing board” to see where the design model was flawed.

A new class of high-end analytical software is now available that previously was the domain of systems designers for nuclear power plants, naval war ships and other mission critical systems. Now, power analytics, a robust set of software tools that designs the electrical scheme, and can be redeployed in “live” mode to perform predictive analysis and ‘what if’ scenarios in order to prevent any surprises and costly downtime once the facility is operational.

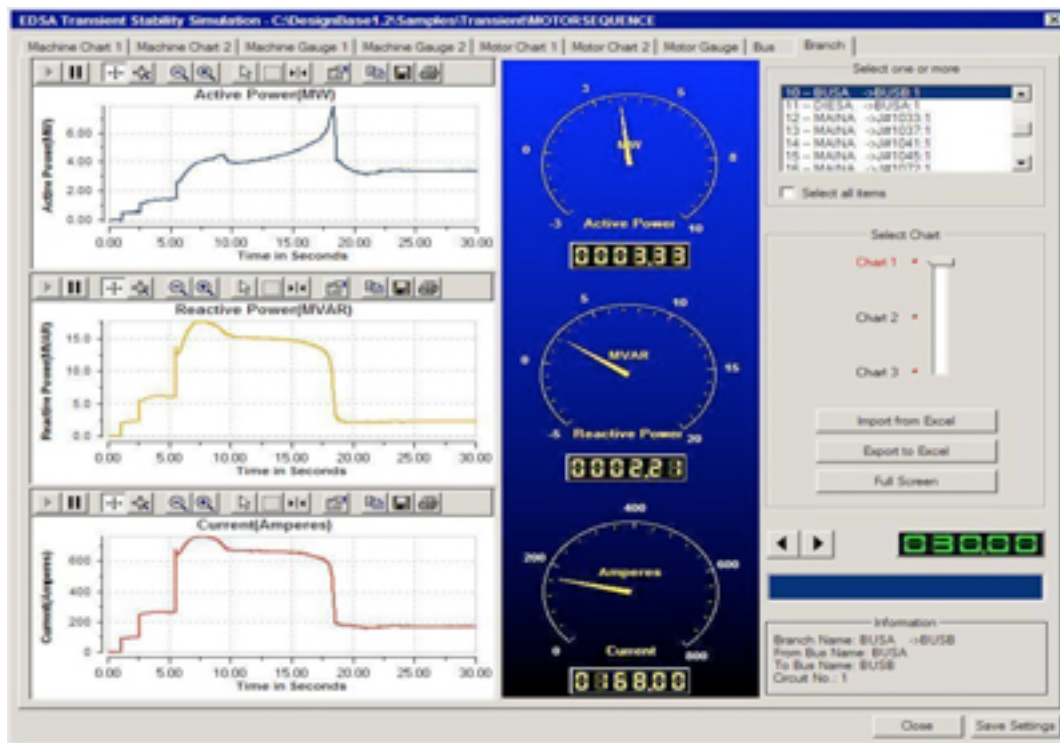
Unlike conventional design/CAD packages, a power analytics’ CAD model remains in electronic form, and is – rather than archived -- retained in on-line mode meaning that all of the components and their specifications “go live” to provide a benchmark for how the system should be performing in its ideal state and what variations may exist between the “ideal” and “actual” states. Electrical engineers can now quickly create a robust electrical “design base” – a detailed design and knowledge base of the all the components, processes, and performance specifications of the entire electrical distribution system.

The first of these power analytics tools, Paladin DesignBase® and Paladin Live® from EDSA, are now in use in extremely demanding customer sites – such as financial data centers, air traffic control, and power distribution networks -- to simulate and analyze systems from a variety of static or dynamic perspectives. They include the ability to model and embed detailed control logic of intelligent electronic devices that control power flow throughout the system.

These tools also provide a way to thoroughly analyze and conduct “what if” testing prior to the construction phase. This is a critical step in the success of the overall design and a crucial tool in optimizing Power Usage Effectiveness (PUE) and Data Center infrastructure Efficiency (DCiE.) These widely accepted benchmarking standards – put forth by the Green Grid – can show just how energy efficient a data center is and monitor the impact of the efficiency efforts.

Once the facility is up and running, power analytics dives deeper by acting as a “collision avoidance” system – one that predicts and prevents electrical power problems from occurring to begin with, potentially saving the organization thousands of dollars per second in lost financial transactions, communications, customers and good will. Planning an efficient power infrastructure – and possessing the capability to continually monitor the system – will help ensure that precious dollars are not lost to inefficiency.

Power analytics systems act as an on-board electrical power system expert, to intelligently filter the power system sensory data, help owner/operators understand the real-time health of their electrical power system, as well as diagnose whether that health is stable, deteriorating, or becoming overloaded (figure 1).



(Fig. 1 Paladin DesignBase power systems engineering modeling platform allows for the development of design models that are re-deployable in real-time.)

Thus, business-impacting decisions rooted in the health and reliability of electrical power can be made well before a costly problem occurs. Incorporating artificial, intelligence-like attributes, power analytics can help the facility or IT manager consider important issues such as:

- ¿ How much more capacity can our existing facilities accommodate, before it becomes necessary to make arrangements for new facilities?•
- ¿ What would be the operational impact of adding new equipment, changing configurations or adopting new technology?•
- ¿ If we were to outsource our manufacturing – or we needed to monitor our suppliers' facilities to ensure the integrity of their systems – how could we do so?

Keeping Your Cool

Technologies such as blade servers make cooling an even larger consideration in the facilities' design as they significantly alter the cooling landscape. Moreover, green technologies for HVAC are very much load and location dependent, so the answers are not always intuitive. It is important to start the design process with a clean sheet, and to make the initial design decisions based upon sound energy and financial modeling.

Important questions to be considered include:

- 7 Is the HVAC system you are installing today able to handle changes in the future?
- 7 Will you be able to add more Computer Room Air Conditioner (CRACs) or Computer Room Air Handler (CRAHs)?
- 7 Will you be able to add enough to handle the anticipated load? What about chiller loops and cooling tower capacity?
- 7 Can the architecture handle a 25 or a 50 percent increase in growth when required? If not, that could present major challenges down the road. Without proper cooling, equipment will fail and outages will cause havoc on your operations.
- 7 Does the cooling system strike the right balance between capital and operating costs? Has the system efficiency been evaluated and correlated with the expected load growth and utility escalation rates?

Look Before you Leap

For projects with a life expectancy over 10 years, the operating costs of the mechanical system will most likely exceed the initial installation cost. In addition, the initial selection of a cooling system will define all future options for cooling upgrades and expansions.

It is important to properly evaluate all cooling system options at the start of each project. At the plant level, the annual energy consumption of chilled water, water-cooled and air-cooled systems must be evaluated. At the terminal level, the selection of underfloor, overhead, or in-row cooling equipment should be considered. Fortunately, there are several energy modeling software packages on the market to make sense of this.

Modern energy modeling software can calculate system performance and energy consumption on an annual basis at various geographic locations and load profiles. For example, a free-cooling chiller plant in Miami will not have the same annual performance as a similar system located in Chicago. In addition, a system that is optimized for full-load performance may be highly inefficient when operated at partial load. Finally, a system that is highly efficient may not pay back if the project life expectancy is 5 years or less.

System selection must be carefully coordinated with the client expectations for installation cost, operating costs, and project life-cycle. Only when considering both energy and financial considerations together can the best decision be made.


Proper Air Flow

The quantities of CRACs and CRAHs required is initially calculated based upon nameplate data, but Computational Fluid Dynamics (CFD) is the best way to verify the design selections and quantities. CFD provides companies with a detailed 3-D analysis of how cold air is moving through the aisles and racks, identifying potential "hot spots" where equipment is receiving too little airflow. Thermal mapping can also find areas in a data center that are receiving more cold air than needed, wasting cooling and energy. For example, the selection of 50 ton CRAH units for a room with a 24" raise floor may look good on paper, but may result in uneven airflow distribution when installed.

CFD is the best way to verify and optimize CRAH and IT rack layouts and densities. Not all computer rooms are the same; CFD modeling helps to significantly improve in-room designs and aid in maximizing utilization of existing facilities.

For high density environments, it's important to know that blade servers consume, typically 10 times more total power than conventional rack servers, and require at least a four-fold increase in cooling capacity. As a result, according to an Uptime Institute study, power consumption in data centers has increased sevenfold (over 600 percent) in just seven years. As the

graph from ASHRAE demonstrates (Fig.1), the trend for increased heat load will increase exponentially in the future.



2N Design

Today's high density infrastructures require full 2N redundancy supported by uninterruptible power supplies (UPSs), power distribution units (PDU) and automatic transfer switches (ATS) in order to meet the "always-on" requirement of the facility. A key challenge is managing the trade-offs between first cost, energy efficiency and maximum uptime, which can sometimes be competing goals. For example, a 2N cooling plant provides maximum uptime, but is capital intensive to build and operate. And, "going green" is only viable if it makes sense to the bottom line. Being in control of the infrastructure is a smart business decision that will help drive savings on energy and maintenance costs.

There is, of course, more to designing the optimum facility, but having a firm foundation for the energy and cooling systems will go a long way in the overall success of the design and function of the data center. As mentioned earlier, power analytics can be a major advantage in assuring high nines of uptime by predicting failures before they emerge thus paying for itself in short order. Armed with a solid design plan and incorporating new software tools that can help design and perform critical "what if" scenarios, achieving more energy efficiency and higher facility uptime can now be quickly achieved.

About the Author: Mark Welte is with Robert Derector Associates - www.derector.com