



Designing modular data centers

When planning for modular data center design, the engineer should focus on attributes such as system efficiency and operational characteristics.

Bill Kosik, PE, CEM, LEED AP, BEMP
12/21/2016

Learning objectives

- Provide a high-level understanding of topics related to modularization in data centers and other critical facility types that have sophisticated and complex power and cooling infrastructure.
- Present arguments on the aspects of financial and operating outcomes for a range of modular design approaches, ranging from a minimal application to a highly modular design with multiple levels of redundancy.

Share



In the design of power and cooling systems for data centers, there must be a known base load that becomes the starting point from which to work. This is the minimum capacity that is required. From there, decisions will have to be made on the additional capacity that must be built in. This capacity could be used for future growth or could be held in reserve in case of a failure. (Oftentimes, this reserve capacity is already built into the base load). The strategy to create modularity becomes a little more complex when engineers build in redundancy into each module.

In this article, we will take a closer look at different parameters that assist in establishing the base load, additional capacity, and redundancy in the power and cooling systems. While the focus of this article is on data center modularity with respect to cooling systems, the same basic concepts apply to electrical equipment and distribution systems. Analyzing modularity of both cooling and power systems together-the recommended approach-will often result in a synergistic outcome.

What is modular planning?


When planning a modular facility, such as a data center, there are three main questions that need to be answered:

- What is the base load that is used to size the power and cooling central plant equipment (expressed in kilovolt-amp, or kVA, and tons, respectively)? In the initial phase of the building, if one power and cooling module is used, this is considered an "N" system where the capacity of the module is equal to the base load.
- In the base load scenario, what is the N that the central plant used as a building block? For example, if the base cooling load is 500 tons and two chillers are used with no redundancy, the N is 250 tons. If a level of concurrent maintainability is required, an "N+1" configuration can be used. In this case, the N is still 250 tons but now there are three chillers. In terms of cooling in this scenario, there would be 250 tons over the design capacity.



"CHOOSING CAT MEANS I GET MORE POWER AND BETTER RELIABILITY."


BUILT FOR IT.



Events

Recent News

- Your facility is mission critical: Resiliency, electrical commissioning
- Grounded and ungrounded electrical and power system design
- Lighting webcast: Designing lighting systems with LEDs
- Selecting and sizing transformers to achieve energy efficiency
- Transformer case study
- Weekly merger, acquisition, deal update: June 16
- How NFPA 99-2015 changes affect mechanical design
- Your questions answered: Combined heat and power
- Smart lighting industry set to transform the commercial and residential landscape
- Transformers: ensuring efficient electrical system




is determined to be linear and predictable (which is a rare scenario), the day one module will be replicated and used for future growth. However, when the growth is not predictable or the module design has to be changed due to changing loads or reserve-capacity requirements, there has to be a strategy in place to address these issues. This is where the module-in-a-module approach can be used.

Module-in-a-module

Each module will have multiple pieces of power and cooling gear that are sized in various configurations to, at a minimum, serve the day one load. This could be done without reserve capacity, all the way to systems that are fault-tolerant, like $2N$, $2(N+1)$, $2(N+2)$, etc. So the growth of the system has a direct impact on the overall module. For example, if each module will serve a discreet area within the facility without any interconnection to the other modules, the modular approach will stay pure and the facility will be designed and constructed with equal-size building blocks. While this approach is very clean and understandable, it doesn't take advantage of an opportunity that exists: sharing reserve capacity while maintaining the required level of reliability.

If a long-range strategy includes interconnecting the modules as the facility grows, there will undoubtedly be opportunities to reduce expenditures, both capital expense and ongoing operating costs related to energy use and maintenance costs. The interconnection strategy results in a design that looks more like a traditional central plant and less like a modular approach. While this is true, the modules can be designed to accommodate the load if there were some type of catastrophic failure (like a fire) in one of the modules. This is where the modular approach can become an integral part in achieving high levels of uptime. Having the modules physically separated will allow for shutting down a module that is in a failure mode; the other module(s) will take on the capacity that was shed by the failed module.

Using the interconnected approach can reduce the quantity of power and cooling equipment as more modules are built, simply because there are more modules of N size installed (see Figure 1 through 4). Installing the modules with a common capacity and reserve capacity will result in a greater power and cooling capacity for the facility.



If uncertainty exists as to the future cooling load in the facility, the power and cooling equipment can be installed on day one, but this approach deviates from the basic design tenets of modular data centers. And while this approach certainly provides a large "cushion," the financial outlay is considerable and the equipment will likely operate at extremely low loads for quite some time.

Equipment capacity, maintenance, and physical size

When analyzing the viability of implementing a modular solution, one of the parameters to understand is the size of the N and how it will impact long-range costs and flexibility. To demonstrate this point, consider a facility with a base load of 1,000 tons. The module could be designed with the N being 1,000 tons. This approach leaves little reserve capacity or the ability to maintain the equipment in a way that minimizes out of range temperature and humidity risks to the IT systems. In this N configuration, taking out a major piece of HVAC equipment will render the entire cooling system inoperable (unless temporary chillers, pumps, etc., are activated during testing or maintenance).

design
More Featured
Content



**"CHOOSING
CAT MEANS
I GET MORE
POWER AND
BETTER
RELIABILITY."**


BUILT FOR IT.



Going to the other end of the spectrum yields an equipment layout that consists of many smaller pieces of equipment. Using this approach will certainly result in a highly modular design, but it comes with a price: All of that equipment must be installed, with each piece requiring electrical hookups (plus the power distribution, disconnects, starters, etc.), testing, commissioning, and long-term operations and maintenance. This is where finding a middle ground is important; the key is to build in the required level of reliability, optimize energy efficiency, and minimize maintenance and operation costs.

Chiller power multiplier based on overall building cooling load

Cooling load distributed across chillers - 75% load



Factory-built versus site-erected modules


When considering how the modules are constructed, there are a few options: site-erected, hybrid, or factory-built. Each of these options has its own set of advantages and constraints. Consider:

- The location of the facility immediately influences the type of module design approach. For example, when facilities are located in sparsely populated areas where skilled piping, sheet metal, and electrical design and construction experts are hard to come by, it will be beneficial to use a factory-built, tested, and commissioned module that is delivered to the site-probably in multiple sections-assembled, and connected to the other systems and facilities. It's a bit more complex and detailed, but for this type of scenario, the factory-built option makes sense.
- Oftentimes, facilities must be built in geographical areas without manufacturer support for start-up, commissioning, and maintenance of new power and cooling equipment. This will require long-distance travel by the manufacturer's technical teams-not desirable in cases of operating anomalies or equipment failures. If there is no choice on the location, upfront planning and special requirements can be written into the specifications to proactively address these concerns. While there will be an increased cost from the equipment vendor, purchasing spare parts upfront and stipulating maximum response time in case of an operating anomaly will lessen the impact of an equipment failure.

- The construction schedule of data centers and other critical facilities typically is driven by a customer's needs, which is often driven by revenue generation or a need by the customer's end-user (e.g., the community, business enterprises, government agencies) to use/occupy the proposed facility as soon as possible. When analyzing the best approach to the construction of the overall facility, it is advantageous to have the module built offsite, in parallel with the construction of the facility. The module can be shipped to the site and installed even if the facility is not complete. Because all of the equipment, piping, and electrical in the module have been installed, tested, and commissioned, the overall time to build the facility can be reduced. Additionally, commissioning and testing of the equipment in a factory setting can be more effective-especially when the people who built the module are onsite with the commissioning authority and all are working together to make sure all of the kinks are worked through.

Chiller power multiplier based on overall building cooling load

Cooling load distributed across chillers - 50% load



critical: Resiliency, electrical commissioning

■ Grounded and ungrounded electrical and power system design

■ Lighting webcast: Designing lighting systems with LEDs

■ Selecting and sizing transformers to achieve energy efficiency

■ Transformer case study

■ Weekly merger, acquisition, deal update: June 16

■ How NFPA 99-2015 changes affect mechanical design

■ Your questions answered: Combined heat and power

■ Smart lighting industry set to transform the commercial and residential landscape

IT'S NOT A QUESTION OF IF, BUT WHERE





- Between the two choices of site-built and factory-built is the hybrid approach to constructing a module. As the name implies, the hybrid approach uses a combination of factory-built and site-erected components. There is not one solution for this approach because the amount of work done on the site, as compared to within the factory, vary greatly from project to project. A good example of why a hybrid approach would be used is when there could be difficulty in shipping large pieces of power and cooling equipment that will be installed in a module. The balance of the HVAC and electrical work could still be completed at the factory and take advantage of reducing the overall schedule. And future expansions can be handled the same way, building in quick expansion capability.

Performance comparisons



difference between the facility air conditioning

load and the capability of the compressors to reduce energy use. In the analysis, each chiller will share an equal part of the load; as the number of chillers increases, each chiller will have a smaller loading percentage. In general, compressorized equipment is not able to have a linear energy-use reduction as the air conditioning load decreases. This is an inherent challenge in system design when attempting to optimize energy use, expandability, and reliability. The following parameters were used in the analysis:

Curve designation: CentH2OVSD-EIR-fPLR&dT

(This is energy modeling shorthand for a water-cooled centrifugal chiller with a variable speed compressor. EIR is the energy input ratio, which is what the equations solve for. fPLR&dT indicates that EIR is a function of the part load ratio of the chiller and the lift of the compressor-chilled water supply temperature subtracted from the entering condenser water temperature.)

Type of curve: biquadratic in ratio and dT

Equation: $f(r, dT) = c1 + c2*r1 + c3*r1^2 + c4*dT + c5*dT^2 + c6*r1*dT$

Coefficients:

- $c1 = 0.27969646$
- $c2 = 0.57375735$
- $c3 = 0.25690463$
- $c4 = -0.00580717$
- $c5 = 0.00014649$
- $c6 = -0.00353007$


An advantage of using a modular design approach is obtaining a higher degree of flexibility and maintainability that comes from having multiple smaller chillers, pumps, fans, etc. When there are multiple redundant pieces of equipment, maintenance procedures are less disruptive and, in an equipment-failure scenario, the redundant equipment can be repaired or replaced without threatening the overall operation.

In data centers, the idea of designing in redundant equipment is one of the cornerstones of critical facility design, so these tactics are well-worn and readily understood by data center designers and owners. Layering modularization on top of redundancy strategies just requires the long-range planning exercises to be more focused on how the design plays out over the life of the build-out.

To illustrate this concept, a new facility could start out with a chilled-water system that uses an N+2 redundancy strategy where the N becomes the building block of the central plant. A biquadratic algorithm is used to compare the different chiller-compressor unloading curves. These curves essentially show the

Chiller power multiplier based on overall building cooling load

Cooling load distributed across chillers - 25% load



■ Transformers:
ensuring efficient
electrical system
design
[More Featured
Content](#)



**"CHOOSING
CAT MEANS
I GET MORE
POWER AND
BETTER
RELIABILITY."**

BUILT FOR IT.



Each of the scenarios (Figures 1 through 4) were developed using this approach, and the results demonstrate how the efficiency of the chiller plants decrease as the overall air conditioning load decreases.

Summary of analysis:

- The N+2 system (three chillers) has the smallest decrease in energy performance when the overall facility load is reduced from 100% to 25%. This is due to the fact that the chillers are already operating at a very small load. So large swings in cooling loads will not have a large impact on the efficiency of the system.
- The N system (one chiller) shows the greatest susceptibility to changes in facility cooling load. The chiller will run at the highest efficiency levels at peak loading, but will drop off quickly as the system becomes unloaded.
- The N+1 system (two chillers) is in between the N and N+2 systems in terms of sensitivity to changes in facility loading.

When put into practice, similar types of scenarios (in one form or another) will be a part of many data center projects. When these scenarios are modeled and analyzed, the results will make the optimization strategies clearer and enable subsequent technical and financial exercises. The type of modularity ultimately will be driven by reliability, and first and operational costs. Because a range of different parameters and circumstances will shape the final design, a well-planned, methodical procedure will ultimately allow for an informed and streamlined decision-making process.

Bill Kosik is a data center energy engineer. He is a member of the Consulting-Specifying Engineer editorial advisory board.