

Our energy working for you.™



Power topic #9020 | Technical information from Cummins Power Generation

Data center design decisions and their impact on power system infrastructure

■ White Paper

By Richard Hallahan, Senior Manager – Strategic Accounts

When designing an enterprise data center, owners, architects and engineers have to understand and decide on everything from site to signage. Some of the most critical decisions revolve around the best power distribution infrastructure to suit the facility's needs. This white paper reviews three key factors to consider when determining the power distribution infrastructure: the size of the system, the reliability architecture that is selected, and the operational complexity of delivering power to the critical loads. Once the team knows these factors, they will have a solid foundation on which to design and build a comprehensive, integrated electrical power distribution system.

Factor one: the size of the power distribution infrastructure

It may seem hard to believe, but the heat generated by IT equipment in a data center is a critical data point in defining the power system infrastructure for the entire facility. (Typically this equipment is in a room with a raised floor for air distribution, although some data centers have a technology room without a raised floor. For the purposes of this paper the technology room will be generically referred to as the raised floor area, but the following applies to both styles of technology room design.) Because the technology equipment and the mechanical infrastructure to keep it cool are the largest consumers of electricity in the building, their maximum electricity consumption defines the power needs of the facility.

Calculation of the electrical consumption of the technology equipment

Data center facility managers have traditionally used watts per square area (foot or meter) when describing the heat load of the equipment within the raised floor area. Recently, however, some managers have started to use kilowatts per rack. Either way, in order to calculate a useful electrical load, one must convert these numbers to watts.

Both methods are described and illustrated below:

Method 1 — raised floor heat load x raised floor area

Once the size of the raised floor area where the technology equipment will reside has been defined, it's a simple calculation of multiplying the number of watts per square foot (or watts per meter squared) times the number of total square feet (or meters) in the raised floor area to determine the watts required to support the technology equipment.

Example:

Raised floor heat load 150 watts per square foot (w/ft²)
Raised floor area x 48,000 total square feet

Electrical consumption 7,200,000 watts = 7,200 kW = 7.2 MW

Similar example in metric units:

Raised floor heat load 1,500 watts per meter squared (w/m²)
Raised floor area x 4,500 total meters squared

Electrical consumption 6,750,000 watts = 6,750 kW = 6.8 MW

Method 2 — technology rack heat load x total number of racks in raised floor area

The watts-per-rack calculation is somewhat area-independent. It focuses on the heat generated at each rack or cabinet as a kilowatt-per-rack value. This approach requires defining the number of technology racks that will reside within the raised floor area when the data center is fully utilized. (This value will lead to the raised floor area size, once row and aisle sizes are determined.) Multiplying kilowatts per rack times the number of racks then provides the power consumption value for the technology equipment in the raised floor area, similar to method one.

Example:

Technology rack heat load 12 kilowatts per rack (kW/rack)
Racks in raised floor area x 600 total number of racks

Electrical consumption 7,200 kW = 7.2 MW

Calculation of the electrical consumption of the data center

Once the kilowatt load of the technology space has been determined, a rule of thumb is to double it to obtain the approximate maximum electricity consumption load for the entire data center. In the case of the earlier examples, the approximate load of the whole data center facility would be 2 x 7.2 MW or 14.4 MW. This rule of thumb assumes that the mechanical systems that cool the technology equipment, e.g., the chiller, pumps, computer room air conditioners and handlers (CRACs/CRAHs), consume

approximately the same amount of power as the technology equipment itself in the raised floor area. These two loads combined account for over 95% of the electrical power consumption of a typical greenfield data center.

Factor 2: the reliability architecture

The second key decision concerns the reliability architecture. In the mid-1990s, the Uptime Institute™ published the first version of its white paper “Tier Classifications Define Site Infrastructure Performance.” Regularly updated, this document has provided the guiding principles for infrastructure reliability at many data centers; it also serves as a commonly used metric to compare data center infrastructures.

The document establishes four performance levels or tiers associated with the infrastructure of a data center facility — from a facility with the minimum or basic infrastructure needed to support a data center (Tier I) to a fault-tolerant architecture with twice the infrastructure needed to support a data center (Tier IV). With each successive tier the “representative site availability” increases. Or to put it conversely, the average amount of time offline decreases with each tier level from I to IV. The higher the tier level, theoretically, the more reliable the data center facility. While this is not the only way to establish the reliability architecture of a data center, it is a widely accepted way to set parameters that define the desired infrastructure characteristics.

Once the reliability architecture of the data center is selected, certain characteristics of the power infrastructure will become better defined. For example, if a reliability architecture that meets Tier II guidelines is chosen, certain power system elements will have redundant components, i.e., an N+1 configuration. See Figure 1. Accordingly, a facility that needs one 2000 kW generator to support the base electrical load would purchase and install two 2000 kW generators. The second generator would represent the redundant component.

Continuing with the same example, this reliability architecture decision would also drive the power system infrastructure; instead of using a transfer switch to integrate the generator, the choice would most likely be to use paralleling switchgear to integrate both generators into the power infrastructure.

The cost of an installation rises proportionately with the degree of reliability. In general, as the facility is designed to be more reliable (i.e., meeting higher tier levels), more infrastructure equipment is needed to meet the reliability architecture needs. More equipment means greater initial costs associated with the construction of the data center facility, as well as greater costs for maintaining that equipment.

The choice of operating voltage for the power infrastructure will also impact the power system architecture and the initial cost of the facility. The impact on the equipment cost can be quite variable, however, and there is no rule of thumb relating operating voltage to initial equipment cost. If operating voltage is a variable in the power infrastructure design, it is recommended that a cost comparison of designs at different operating voltages be performed, to enable the most informed decision about what is best for the data center.

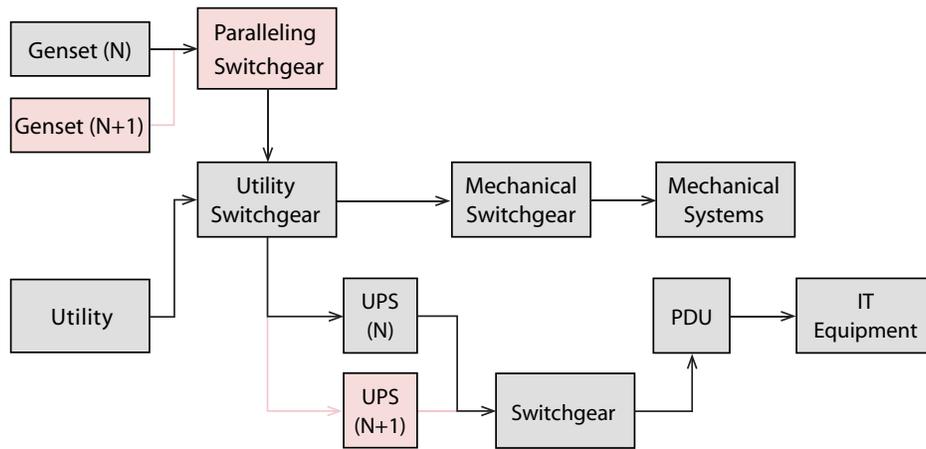


Figure 1 - Sample Tier II topology, with redundant capacity components and single, non-redundant distribution paths.

Factor 3: operational complexity

The third factor that influences the power system infrastructure design is the operational complexity of the system. If a higher tier level is chosen, that choice drives other decisions, such as implementing paralleling switchgear and other automatically controlled devices within the power system architecture. As a result, there is added complexity in the sequence of operations — the automatically initiated program of directions within the power distribution equipment to react to a change in normal operating conditions, such as a loss of utility power.

This sequence of operations implements the transfer to the redundant power path interconnections of the electrical distribution. These interconnections keep the facility operational during the potentially mission-disabling event, such as a power failure. There are many choices associated with operational complexity, from simple designs with just two power paths, to complex designs that control multiple redundant power paths. Once the reliability architecture decision about which tier level is made, the range of choices is reduced, but still a number of options remain.

Commonly used power system architecture designs are one way to simplify the choices associated with operational complexity. The one that best suits the needs of a particular data center design will be determined by earlier decisions in the process, as noted above. These power system architecture configurations are commonly known by names such as isolated bus, multiple transfer pairs, or main-tie-main split generator bus, to name a few. See Figure 2. The standard control blocks — paralleling controls, power transfer controls and system controls — are typically integrated into the larger power distribution system for the facility, which may include additional supervisory control in the form of a SCADA system.

Standard architectures have advantages over custom-designed ones. The most important advantage is that the programming associated with a standard sequence of operations has been used in multiple applications of this type of power system architecture. The associated sequence of operations has been subjected to many hours of operation, which should reveal any major software issues that might cause operational instability. To put it another way, standard power system architectures can be considered more reliable than custom-designed ones because the programming associated with the sequence of operation is more stable, thanks to many hours of operation.

Interchangeability is another benefit associated with standard power system architectures. Most components of the power system design would be “off the shelf” if they are used in a standard design. Should any component need to be replaced for any reason, it should be readily available and easily integrated into the power system infrastructure, reducing potential downtime. An inability to easily replace bad components is one of the main drivers that impact the serviceability of the equipment.

In addition, custom-designed, operationally complex systems have some less obvious costs associated with them. Complexity increases hardware costs because of the custom nature of the product, but complexity also adds cost associated with verifying the operational capability and long-term maintenance costs. These aspects should also be taken into account when deciding on the level of complexity of the power distribution system.

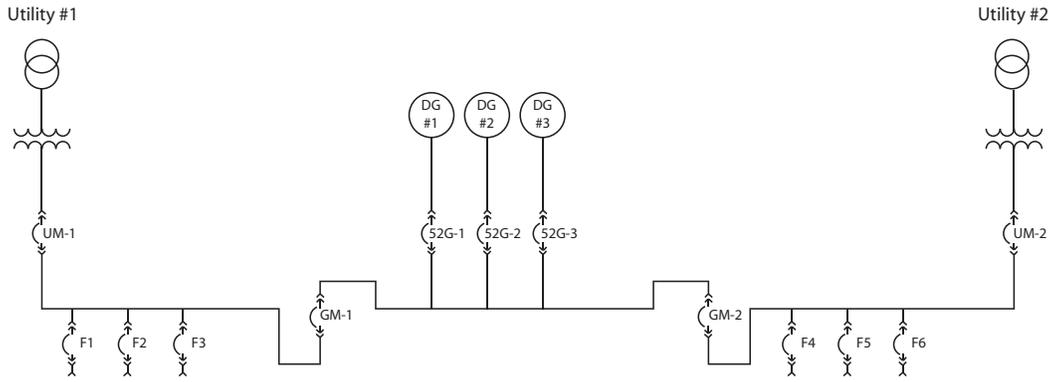


Figure 2 - Example of a commonly used paralleling configuration for data centers, with multiple transfer pairs.

Recommendations

In order to choose the best power system foundation for an enterprise data center, decision makers should carefully consider the three factors discussed in this paper:

- Understand the raised floor area, heat load and square area or number of racks, because that will help estimate the total power consumption needs for the facility. The outcome will tell you the approximate size of the required power system.
- Define the chosen reliability architecture for the facility, using guidelines such as the Uptime Institute's tier levels to drive the power system architecture decisions. This will determine how power gets to the critical loads.
- Implement standard power distribution system architectures, especially when paralleling multiple generators, because this will increase reliability and enhance serviceability.



About the author

Rich Hallahan has been employed by Cummins Power Generation since 2008. His current role is as a Senior Manager – Strategic Accounts. He has been in the electrical construction industry since the early 1990s. Prior to Cummins, Rich worked for EYP Mission Critical Facilities, now a part of HP Mission Critical Services, where he managed teams through the design and construction of multiple enterprise data centers. He also spent over 10 years in the New York metropolitan area as

an electrical engineer at AKF Engineers and as project manager for an electrical contracting firm. Rich is a licensed engineer with a bachelor's degree in engineering from Lehigh University and an M.B.A. from the University of Minnesota's Carlson School of Management.



Our energy working for you.™
www.cumminspower.com

©2011 Cummins Power Generation Inc.
All rights reserved. Cummins Power Generation and Cummins are registered trademarks of Cummins Inc. "Our energy working for you.™" is a trademark of Cummins Power Generation.
PT-9020 (9/11)