

Designed for uptime

Defining data center availability via a tier classification system

MIETEK GLINKOWSKI – All systems can fail – this is a simple fact that every industry must deal with. The paramount concern for the data center industry is the unbroken continuity of systems operations. Industry analysts estimate that a one-hour outage in a data center costs on average \$350,000. And the cost is expected to only go up as more and more business enterprises depend on the storage, networking and processing of digital data, nearly all of it through or inside a data center. Since loss of service for a data center is so costly, even if only for an extremely short time, availability is still the most critical driver for data centers design, operation and maintenance.

The paramount concern for the data center industry is the unbroken continuity of a systems operation.

1 Reliability and availability

Reliability and availability are often misinterpreted and confused with the quality of a system or a product. Reliability is defined as a function of time:

$$R(t) = e^{-\lambda t}$$

where $R(t)$ is reliability, t is time, and $\lambda = Tf/Tp$ is a failure rate. Tf is the total number of failed occurrences during the total period of Tp . The longer the system is operating the lower the reliability. The parameter λ is a reciprocal of MTBF (mean time between failures). Mean time to repair (MTTR), which is the time needed to repair a failed system or device, is another important parameter. Used in combination, MTBF and MTTR determine the inherent availability (A_i) of a system or device:

$$A_i = MTBF / (MTBF + MTTR).$$

If one expands the concept of availability to include the scheduled maintenance downtime the availability changes to the operational availability, A_o .

Reliability and availability are not fixed numbers. They are both functions of specific components of the system as well as the system topology,

which determines how critical these components are to the mission critical function of the data center. Therefore, the reliability has to be evaluated at different points of the system where power is to be delivered to the IT load.

As mentioned, reliability and availability are not the same as quality. Quality refers to the condition of the new equipment when delivered to the customer – ie “out of the box.” Reliability and therefore availability is measured over a period of time. This, besides quality, includes the effects of aging and stress level of the equipment within the system.

Increased reliability can be accomplished by redundancy (of equipment and delivery paths). However, the more equipment the greater the likelihood for one or more of the components to fail. For any system design there is a balance between the level of redundancy and associated complexity and reliability gains. Good system designs need to get the most out of the equipment, utilize their full potential and provide a sufficient level of redundancy and back up for reliable energy supply.

Availability of the data center refers to meeting the uptime expectations of the users. The current high availability of data centers has been achieved mostly through redundancy in design, equipment (both IT equipment and power devices), electricity delivery paths and software → 1. Several classification systems exist in the industry to define data center availability. Rapidly changing technologies, desire to differentiate among themselves, environmental awareness and foremost cost pressures often dictate designs that either fall in between different tier structures or even seek more radical departures. The tier structure from the Uptime Institute, though not always followed, is considered an important industry guideline and thus is the classification referenced in this article. The Uptime Institute defines a four-tier system, where each level describes the availability as a guideline for designing data center infrastructure → 2. The higher the tier, the greater the availability.

The lowest cost and the lowest performance data center, Tier I, has a target availability of 99.671 percent, which translates to 28.8 hours of annual IT downtime. The highest level data center design, Tier IV, has a target of 99.995 percent availability, or 24 minutes of an-

nual IT downtime. The different tier designs are also capable of accommodating different power load densities, from 200 W/m² to 1,500 W/m². For power engineers it is important to realize that the higher the tier the higher the utility voltage supplied to the facility. This is predominantly related to the fact that the availability of power within a power system is generally increasing from low-voltage (LV) area distribution to medium-voltage (MV) distribution to high-voltage (HV) transmission systems. The closer one is to the infinite bus of a large power system the less the likelihood of a disturbance or blackout.

Tier I

This architecture is the simplest and therefore offers the lowest availability and lowest IT load power density. This design concept is called N, reflecting the fact that “n” IT loads need “n” sets of UPS units and gensets. → 3 identifies the basic components of a data center, as described below.

Utility source

The utility source component in a Tier I classification feeds an input transformer stepping down from MV to LV.

Title picture

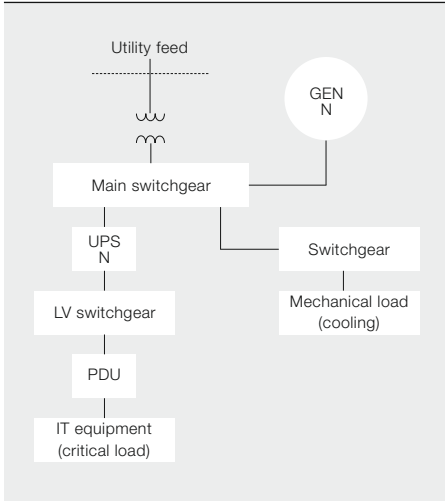
What sort of designs do data centers follow in order to meet the high demands for availability?

2 Tier similarities and differences

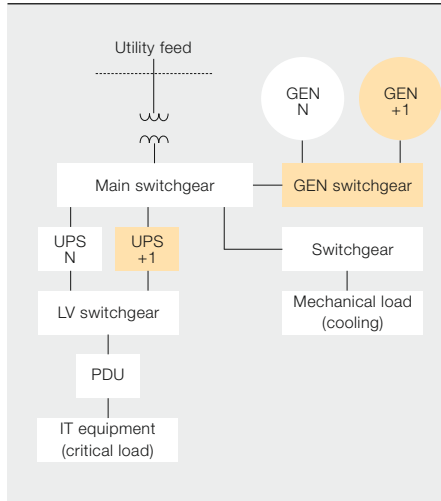
	Tier I	Tier II	Tier III	Tier IV
Number of delivery paths	Only 1	Only 1	1 active 1 passive	2 active
Redundant components	N	N + 1	N + 1	2 (N + 1) or S + S
Utility voltage	208, 480	208, 480	12-15kV	12-15kV
Annual IT downtime due to site	28.8 hours	22.0 hours	1.6 hours	0.4 hours
Site availability	99.671%	99.749%	99.982%	99.995%
© The Uptime Institute				

The current high availability of data centers has been achieved mostly through redundancy in design, equipment, electricity delivery paths and software.

3 Tier I design N



4 Tier II design (N+1)



Genset

A genset is an emergency power generator, typically with a diesel engine, that provides a long-term power backup in the event of a utility outage. Long-term is defined by the amount of fuel stored in the tank and can vary from 24 to 72 hours. Having a high-priority fuel delivery contract can extend the time. The generator is in the form of a synchronous machine with power ratings of few hundred kW to 2 to 3 MW.

Automatic transfer switching

Using a specialized automatic transfer switchgear (ATS) with control and protection logic allows for a seamless switch from the source between the utility and the genset under a number of different conditions. Most of the time the switch from the utility to the generator is open-before-close – ie, when the utility power is lost the utility breaker is open and the genset is closed only after the genset has started properly, reached the desired rpm and excitation, and is synchronized. The starting of the genset can take a few seconds. With multiple

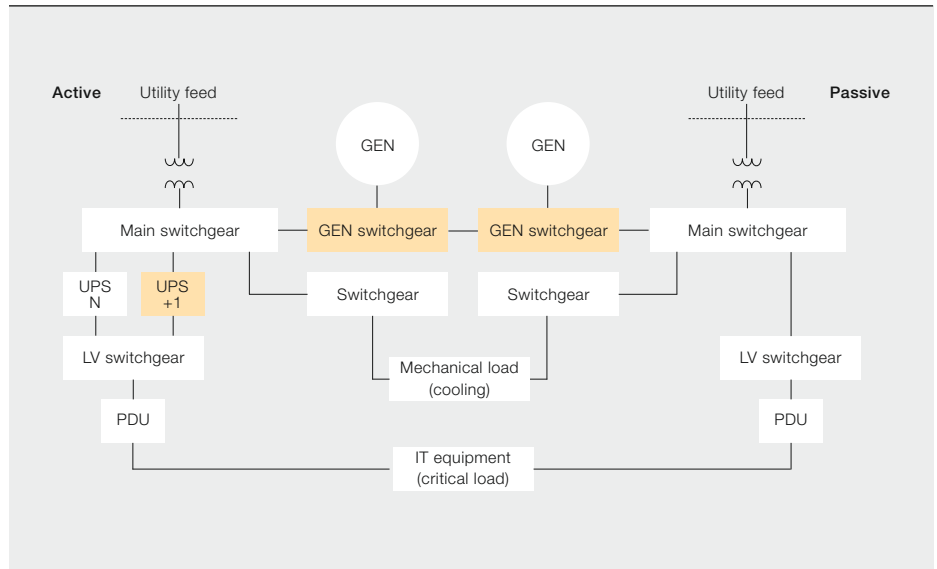
gensets this time can increase to up to a minute.

Uninterruptible power supplies

There are primarily three types of uninterruptible power supply (UPS) technologies – standby, line interactive and double conversion. By far the most popular is double conversion, where all the power flowing through the UPS is rectified from AC to DC, inverted back to AC and therefore fully conditioned and cleaned from all utility-side disturbances, transients, voltage sags and swells, and other power quality (PQ) effects. The DC bus in the middle is also connected to the battery bank, which, in the event of power loss, provides short-term power. The switch between the utility AC power and the internal battery power is seamless and instantaneous. Short-term power is determined by the size of the battery bank and typically varies from 2 to 3 min to 7 to 10 min.

The addition of the passive delivery path significantly raises the cost of the entire system and also complicates the control, coordination and maintenance.

5 Tier III active-passive design; no UPS in the passive path



Switchgear

A variety of switchgear is needed in data centers to distribute the power to the many different rows of IT equipment (critical loads) as well as cooling equipment (pumps, fans, valves, compressors, etc.) and other auxiliary loads. The circuit breakers in the switchgear also provide protection against faults and other abnormal conditions. In the Tier I facility all of the switchgear is low voltage (less than ~1 kV).

Power distribution unit

Power distribution units (PDUs) are comprised of circuit breakers, metering units and, in North America, LV transformers, to further distribute the power to the IT racks as well as provide protection and measure the power (voltage and current) to the individual loads.

Power supply units

Power supply units (PSUs) are part of the IT equipment. Similar to the power supply of a desktop computer these units transform the 220 V or 110 V input power to the DC voltage distributed to the various IT equipment: servers, network and storage systems. The most popular PSUs are transformer-less switched mode power supply (SMPS). Due to the redundancy of the power distribution for Tier III and IV more and more PSUs are now provided with dual AC inputs and can function from either of the two.

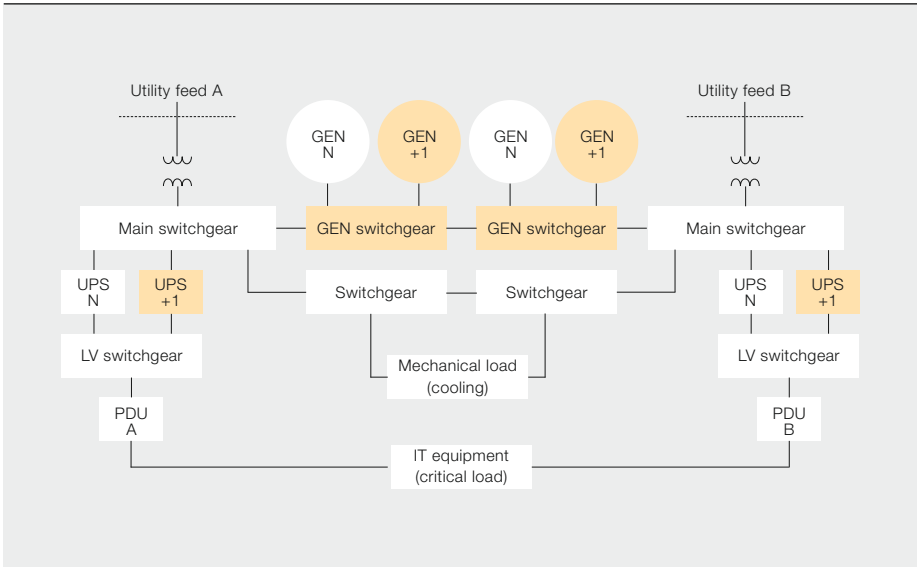
Tier II

This design is known as N+1 → 4. The primary difference between a Tier I and Tier II classification is the presence of an

additional genset and UPS. This provides some degree of device redundancy of the most critical components of the system for short-term and long-term backup. All other components of the system are basically the same. Even with this redundancy there are still several different single points of failures in the path to deliver power to the IT load.

Tier III

Tier III is referred to as an active-passive system → 5. In a Tier III classification the power delivery path has to be doubled. Besides the redundant critical components there has to be a second path parallel to the critical IT load in case the primary path has failed. This second path could be passive, ie, used only in case of emergency. A Tier III classification also requires a second utility connection. The addition of the passive delivery path significantly raises the cost of the entire system and also complicates the control, coordination, maintenance, etc. There is also an additional switchgear and motor control center (MCC), which should allow the full operation of the data center from the passive path. The IT equipment can now take full advantage of the dual supply paths and therefore utilize dual PSUs for each server, for example. As a result the number of single points of failure is significantly reduced. However, the passive delivery path does not require UPS so during the emergency conditions the system is vulnerable to utility conditions, therefore potentially exposed to utility power quality issues or even power outages.



For any system design there is a balance between the level of redundancy and associated complexity and reliability gains.

Tier IV

Referred to as a 2N+1 system, the Tier IV classification is also considered the Cadillac of data center design → 6. A relatively small number of data centers in

For example, during one year, 10 short power interruptions at the server power supply lasting 50 ms each will have a much more detrimental impact on the operation of the servers than one

longer interruption of 500 ms during the same period of time. Although both will result in the same annual availability (total of 0.5 s of lost power) the first one will cause the servers

Tier IV designs are fully redundant, complete dual systems running actively in parallel.

the world are certified as Tier IV designs. They are fully redundant, complete dual systems running actively in parallel. By virtue of the redundancy the rating of each path has to be 100 percent of the load and therefore the maximum utilization of the two paths under normal operating conditions is at maximum 50 percent. In addition, some Tier IV designs will have N+1 of UPSs and gensets in each path, further increasing the complexity and cost but at the same time gaining the valuable fraction of a percent (0.01 percent to be exact) for availability. The target for Tier IV availability is to allow a maximum of 24 min per year of the annual site-caused end-user downtime (representing one failure every five years).

to reboot and possibly lose some data 10 times during the year; the second one will result in only one reboot a year.

Highly skilled engineering resources are needed to design, implement, and optimize the entire data center ecosystem for their availability and reliability. The traditional way of thinking about availability and reliability is changing rapidly. Increased system voltages, more sophisticated switching schemes, wider operating regimes for IT equipment, and foremost the advent of failure-resilient software and cloud computing introduce new dimensions to data center reliability. So, stay tuned.

Changes to come

Tier structure availability and downtime are not the only factors to consider. Impact of the interruptions on the operation of the mission critical facility can vary.

Mietek Glinkowski
 ABB Data Centers
 Raleigh, NC, United States
 mietek.glinkowski@us.abb.com