12

Data Centres

Angelo Baggini and Franco Bua

A data centre can consume up to 100–200 times as much electricity as standard office premises. With such a large energy consumption, they are prime targets for energy efficient design measures that can save money and reduce electricity use.

However, the reliability and the high power density capacity required by a data centre puts many design criteria far above energy efficiency.

Designing an energy efficient data centre therefore requires attention, skill and investment, but if done correctly it can provide substantial benefits. It is important to develop a holistic strategy and management approach to the data centre. This approach should grant the desired reliability, economic, utilization and environmental benefits.

The three main factors to consider when designing a data centre are:

1. reliability, a feature that is often guaranteed by the redundancy of equipment;
2. scalability, which can be achieved by using modular components that allow one to adjust to situations that may change over time and to avoid unnecessary over sizing;
3. the choice of high efficiency components.

These features must be taken into account in the selection of each component, whether it is IT equipment, components of the power supply infrastructure, or part of the HVAC system.

12.1 Standards

A variety of activities on the topic of Green Data Centres are currently ongoing in standardization organizations. An overview of the main initiatives in this field is presented here.

With regard to general data centre standardization, ANSI/TIA-942:2005 represented the first activity in the field of 'telecommunications infrastructure', providing a number of key definitions but with some content in relation to energy efficiency. Work is now being undertaken

Electrical Energy Efficiency: Technologies and Applications, First Edition. Andreas Sumper and Angelo Baggini.

in the USA by BICSI to produce a Data Centre standard that builds on and expands the content of ANSI/TIA-942.

ISO and IEC are working on this subject with ISO-IEC-JTC 1. This joint working group identified the energy efficiency of Data Centres as a significant topic in the industry and has established a Study Group on Energy Efficiency of Data Centres (EEDC) to investigate market requirements for standardization, initiate dialogues with relevant consortia and to identify possible work items for JTC 1.

The telecommunication world has published with ETSI a series of documents entitled ‘Broadband Deployment – Energy Efficiency and Key Performance Indicators’ of which ETSI TS 105174-2-2 specifically addresses operators’ data centres.

CENELEC has set up a specific working group to discuss the potential standardization actions in relation to energy efficiency within data centres. This working group (CLC BTWG 132-2) has reviewed the development and, where completed, outcomes of work relating to data centre energy efficiency undertaken by the European Commission, standards bodies external to CENELEC and recognized international fora. The BTWG 132-2 recommended to CEN, CENELEC and ETSI that they establish a joint CEN-CLC-ETSI group to manage and coordinate European activity in this field.

More specifically within CENELEC, TC 215 has initiated work on the EN 50600 series of standards, which addresses facilities and infrastructures to support effective operation of telecommunications cabling and equipment within data centres. Amongst other issues, these standards will specify:

- relevant measurement methods of parameters that may be used to determine energy efficiency;
- infrastructures necessary to enable the measurement of those parameters.

With reference to de facto standardization, the primary industry forum is Green Grid™ an open industry consortium of end users, policy-makers, technology providers, facility architects, and utility companies that has the aim of improving the resource efficiency of data centres. Also in this field the European Commission DG-JRC has instituted the Code of Conduct on Data Centres, which is a voluntary scheme targeted at improving the energy efficiency of data centres and which comprises an important document concerning best practices.

12.2 Consumption Profile

One of the basic operations to achieve high energy efficiency for a data centre is the measurement or estimation of consumption of the equipment that compose it.

In the case of a new data centre it is advisable to estimate its energy consumption in the design stage, so that the choice of equipment is as correct as possible. Very often the data centre is placed inside a building that is not completely dedicated to it; therefore it is not so easy to detect the energy consumption of a data centre if energy meters are not installed for this purpose.

In addition to total energy consumption, it is also important to understand and monitor how energy is distributed within the data centre, by the various energy flows and correlate them with areas or a single piece of equipment.

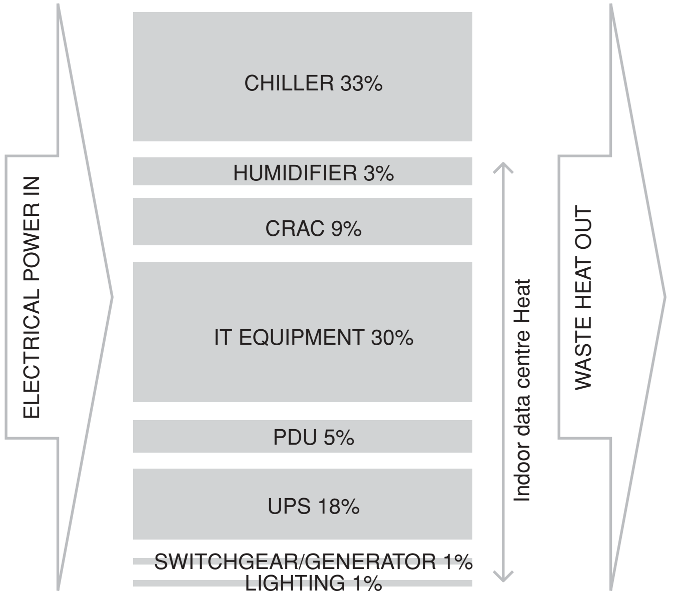

Figure 12.1 Typical data centre energy balance (Elaboration of data from Green Grid)

In this way, it is possible to identify not only the areas or equipment with the highest energy consumption, but also those whose functions are unnecessary and can be switched off or put on standby.

This is an assessment that can be time-consuming, given the large number of measurements needed, but it is rather critical, and fundamental in order to build the energy map of the data centre.

A typical data centre energy balance can be divided into three main parts: power, cooling and IT equipment. (See Figures 12.1 and 12.2.) The useful work is represented by data processing

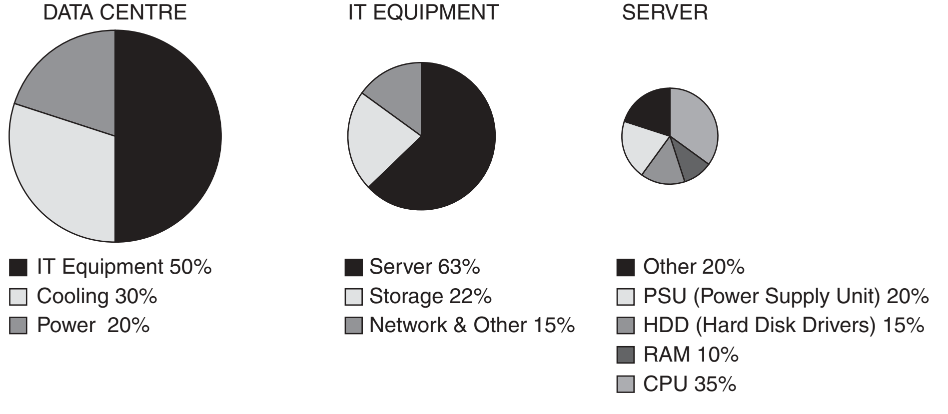

Figure 12.2 Typical values for single data centre components (Elaboration of data from ENEA)

in IT devices whose consumption is about half of the total energy used by a data centre, which must also cover the needs for HVAC and ancillary systems (e.g. fire safety, intrusion, etc.).

The energy used by the IT system is converted into heat, which must be dissipated; increasing the efficiency of these devices implies a direct decrease in the energy consumption as well as a reduction in cooling (about one third of the total).

12.2.1 Energy Performance Index

Evaluating the energy efficiency of a data centre is a complex operation. In principle, the efficiency should be calculated based on the useful work done, but it is difficult to have a standard measure of the useful work for a data centre.

Internationally, a performance indicator has been proposed by Green Grid™: PUE (Power Usage Effectiveness) and its inverse DCIE (Data Center Infrastructure Efficiency). These indexes were initially based on the measurement of electrical power (kW); the tendency is now to shift from these indexes to others based on the measurement of energy (kWh). In early 2010 four categories of PUE were defined: zero, one, two and three depending on whether you refer to power or energy, depending on where the measurement is made.

The PUE zero, the first to be defined, is the ratio between the total electric power consumption by a data centre and the power consumption by IT equipment; the theoretical ideal would be 1, but values up to 1.5 are considered good.

Existing data centres have average values of around 2.5; this indicates that there is a significant degree of improvement and suggests the opportunity for investment in energy efficiency measures.

12.3 IT Infrastructure and Equipment

IT equipment can be considered as the end user of a data centre. Their function, however, can only be ensured by the presence of an adequate supply system and cooling infrastructure. The supply system has to ensure electrical continuity and quality, while the cooling infrastructure has to dissipate the heat produced by the IT equipment and to avoid overheating, with consequent damage to the block of assets, or even permanent failure of the equipment.

The choice of IT equipment in a data centre should be seen not only in terms of its pure IT performance such as, for instance, speed of calculation; when it comes to energy efficiency the indicator operations per second in relation to absorbed watts (Ops/W) should be taken into account. This is an indicator that SPEC (Standard Performance Evaluation Corporation) calculates as a benchmark for the IT equipment market.

The choice of efficient IT equipment brings indirect benefits: lower heat production and less need for electric power to ensure quality and continuity of service.

Voluntary programmes that classify and label the most efficient IT equipment can help in this selection. Energy Star is the most popular energy efficiency programme for servers, PCs, desktops, monitors and printers, while the 80 PLUS programme is for power supplies.

12.3.1 Blade Server

This is a particular type of server that includes a high concentration of computer components (CPU, RAM, storage, etc.) that share the auxiliary components (power supplies and fans).

This type of server can be either introduced in new racks, or in place of the classic-type ones. For example 14 blade servers occupy the space of seven classic. This solution thus allows the elimination of classic-types of units and replaces them with blade-types, which can be considered high density and high performance computing components. Increasing the power of computing in a smaller space will result in large increase of heat per unit area, up to 28 kW per rack, so special cooling systems need to be provided. This type of server provides an overall saving compared with the classic setup of about 10%, with the same computing power.

12.3.2 Storage

The choice of technology storage systems, solid state storage (SSD Solid State Drives) technology versus traditional HDDs (Hard Disk Drives) can give advantages both in terms of efficiency and speed of data selection.

The SSD have lower power consumption and produce less heat. The speed of data selection in a storage system is measured in input and output operations per second (IOPS Input/Output Operations Per Second). In the case of HDDs, with a rotation speed of 15 000 rpm, 300 IOPS are possible, while SSDs can reach up to 30 000 IOPS per unit. In order to have the same speed input/output data of an SSD 100 HDDs are needed.

According to Sun/Oracle estimates, the ratio IOPS/watt for the SSD is around 10 000, while in the case of the HDD it is about 20. (Figure 12.3) At the same HDD IOPS consumption is about 500 times.

SSDs diffusion is still hindered by high costs compared with traditional HDDs, even if this difference should decrease in the future.

12.3.3 Network Equipment

The new generations of network equipment pack more throughput per unit of power; there are active energy management measures that can also be applied to reduce energy usage as network demand varies. Such measures include idle state logic, gate count optimization, memory access algorithms and Input/Output buffer reduction.

Peak data transmission rates continue to increase, requiring dramatically more power; increasing energy is required to transmit small amounts of data over time. Ethernet network

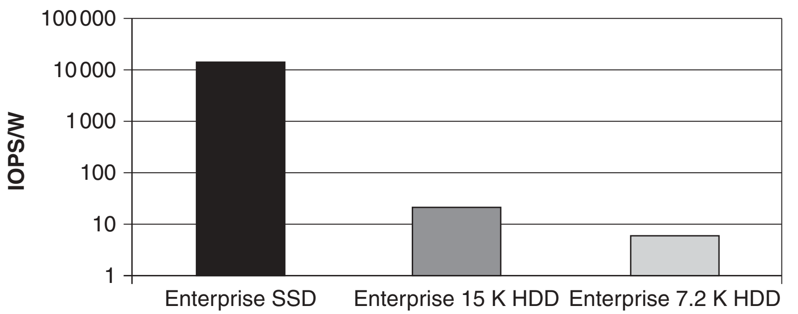

Figure 12.3 Comparison of IOPS/W ration for different storage devices

energy efficiency can be substantially improved by quickly switching the speed of the network links to the amount of data that is currently transmitted.

12.3.4 Consolidation

Where dealing with an existing data centre the possibility of replacing old machines with new ones should be considered.

The calculus performance of processors and memories evolves so fast – according to Intel estimates – that every four years it is possible to completely replace all machines with the latest models and reap the benefits in terms of energy consumption.

For example the consolidation of nine 2005 racks with a 2009 single rack can allow about a 10 times reduction in power demand.

Server consolidation can be achieved in several ways. Two typical approaches are:

1. combining applications on to a single server and a single operating system instance and
2. virtualizing application workloads in virtual machines and hosting them on a hypervisor.

Consolidation can also be applied to physical space. A data centre with lower supply fan power and more efficient cooling system performance, in fact, can be achieved when equipment with similar heat load densities and temperature requirements are grouped together. Isolating equipment by environmental requirements of temperature and humidity allows cooling systems to be controlled to the least energy-intensive set points for each location.

This concept can be expanded to data facilities in general. Consolidating underused data centre spaces in a centralized location can ease the utilization of data centre efficiency measures by condensing the implementation to one location, rather than several.

12.3.5 Virtualization

Virtualization is a form of consolidation and can be applied not only to servers but also to other components or applications, such as desktop, storage, application, and network virtualization.

Server virtualization is a technology that allows for the consolidation of computing workloads that historically required a dedicated physical server per workload. Organizations can achieve energy savings by reducing physical servers while maintaining an ample level of computing capacity on a single virtualized server or, perhaps, a condensed array of virtualized servers.

Virtualization is a tool that provides greater energy savings, allowing one to maximize the potential of the available machines. This technique allows several 'virtual servers', each of which has specific operating systems for the functions to be performed, inside a physical machine. Hardware performances are therefore maximized with respect to the software employed.

This action significantly reduces energy consumption as it allows the consolidation of some machines with shorter payback times, although it is difficult to quantify as related to the configuration of the system to be virtualized.

Virtualization, while maximizing the potential of the servers, leads to the formation of hot areas inside a room. Sometimes it may therefore be necessary to provide appropriate cooling systems on servers that contain the virtual machines.

Some of the main benefits of adopting virtualization include:

- greater flexibility;
- better management processes;
- reduction in the number of machines;
- reduction of operating costs (power and cooling);
- more space available to the data centre.

12.3.6 Software

The choice of data processing software can also have its effects on energy efficiency. In some cases, software that perform the same operations require a different percentage of CPU utilization. Often, the operations within the data centre are repetitive and software that requires fewer calculations to the CPU, and therefore less energy to perform the same functions, can give interesting results.

It is therefore advisable to check the load of the processors, memory and other server resources for different applications in the evaluation, before making the choice, and take it into account, especially when there is no large difference in terms of computational results.

12.4 Facility Infrastructure

12.4.1 Electrical Infrastructure

IT equipment needs a continuous and high-quality power supply. Very often, however, the electrical supply system is of poor quality and can damage equipment or interfere with its functioning.

The main equipment of the power supply system are as follows:

- UPS (Uninterruptible Power Supply);
- PDU (Power Distribution Unit);
- PSU (Power Supply Unit);
- cables.

12.4.1.1 UPS (Uninterruptible Power System)

The choice of UPS must ensure not only the maximum user protection from power disturbances and discontinuities, but also high efficiency.

A static UPS consists of rectifier, inverter and battery. The rectifier converts the AC current to DC for recharging the batteries, while the inverter converts the DC to AC supplying the end user. Double transformation of the energy from the network applies.

Static UPS can have three different architectures:

1. *double conversion*: in normal mode of operation, the load is continuously supplied by the converter/inverter combination in a double conversion technique, i.e. AC/DC and DC/AC;
2. *line interactive*: in a normal mode of operation, the load is supplied with conditioned power via a parallel connection of the AC input and the UPS inverter. The inverter is operational

to provide output voltage conditioning and/or battery charging. The output frequency is dependent upon the AC input frequency;

3. *passive standby*: in a normal mode of operation, the load is supplied by the AC input primary power via the UPS switch. Additional devices may be incorporated to provide power conditioning as voltage stabilizer systems. The output frequency is dependent on the AC input frequency. When the AC input supply voltage is out of UPS preset tolerances, the UPS enters a stored energy mode of operation, when the inverter is activated and the load transferred to the inverter directly or via the UPS switch (which may be electronic or electromechanical).

The most appropriate mode of operation to evaluate the energy efficiency of a static UPS is a double conversion (also known as online), which guarantees complete protection against power disturbances.

Each UPS is sized for the load that it has to feed and its efficiency is generally a maximum for load values close to the maximum.

For the latest generation of UPS, although the efficiency is more constant as load decreases, the tendency is to use modular groups that can adapt to the conditions and work load required in order to ensure the highest efficiency.

The energy lost in conversion turns into heat, which must be extracted by the HVAC system; this means that greater efficiency means less cooling required and lower operating costs.

Dynamic or flywheel UPS have the same function as static, and are generally installed without batteries; in this case, however, in order to provide continuity of supply for more than 10–15 seconds or so they must be coupled with gensets or the protection will cover transient and short duration disturbances only.

This type of UPS is formed by a flywheel, which rotates at high speed accumulating kinetic energy that is converted into electricity when a voltage disturbance occurs. When the flywheel UPS is equipped with batteries their useful life will increase as the flywheel reduces the number of charge/discharge cycles to which they are subjected.

Dynamic UPS are characterized by high efficiency, around 96–97%. They also have the advantage of being able to operate in environmental conditions that are less restrictive than those required by a static UPS, with the possible advantages of reduced consumption for air conditioning.

When selecting the type of UPS, the network characteristics of the construction site must be taken into account: few disturbances allows a UPS to work in by-pass mode with high efficiency; if the frequency of power disturbances is high and their duration is limited to few seconds, the use of a rotary UPS could be considered.

The choice should always be made taking into consideration the specific characteristics of each installation and carrying out a life cycle cost analysis, which takes into account the variables such as the cost of cooling, maintenance, etc.

12.4.1.2 PDU (Power Distribution Unit)

PDUs (Power Distribution Units) have the function to distribute power to the various end users in the data centre. The losses introduced are low, but they should not be neglected as they ultimately turn into heat, which has to be removed.

In the design phase it is very important to choose their position properly as it influence the power cables routes; in fact power cables should preferably be placed in an orderly manner to facilitate maintenance and to avoid obstructing the passage of the cooling air. They are generally mounted at the rear of the cabinet, and fixed to the floor or the ceiling.

It is advisable to choose PDUs equipped with instruments that can measure the instantaneous power consumption of the connected equipment, which allows for the monitoring of the system.

12.4.1.3 PSU (Power Supply Unit)

Most data centre equipment uses internal or rack mounted alternating AC/DC power supplies. Historically, a typical rack server's power supply converted AC power to DC power at efficiencies of around 60–70%. Today, through the use of higher-quality components and advanced engineering, it is possible to find power supplies with efficiencies up to 95%.

Using higher efficiency power supplies will directly lower a data centre's power bills and indirectly reduce cooling system cost and rack overheating issues.

There are also several certification programmes currently in place that have standardized the efficiencies of power supplies; for example the voluntary programme 80PLUS, labels power supplies with high levels of efficiency, and can help in the selection process.

12.4.1.4 Lighting

In designing the electrical system for the data centre, lighting should also be considered. Data centre spaces are not uniformly occupied and, therefore, do not require full illumination during all hours of the year.

Careful selection of an efficient lighting layout, lamps and ballasts will reduce not only the lighting electrical usage but also the load on the cooling system.

Therefore, it is recommended that one use the most efficient lighting technologies, such as LEDs, combined with an automatic control system zone based on occupancy sensors.

12.4.2 HVAC Infrastructure

The design of the HVAC system must be based on the actual load of IT and electrical equipment, trying to avoid as much as possible unnecessary over sizing and ensuring high reliability.

The main parameters to consider are:

- the total heat to be dissipated developed by all the equipment;
- the spatial distribution of thermal power to be dissipated with the identification of 'hot spots';
- the type of cooling system to be used;
- the operation temperature of IT equipment.

The designer's work can be supported by appropriate simulation software of thermal loads and air conditioning systems. Computational Fluid Dynamics (CFD) may be very useful for studying the movement of air flows.

In the case where a data centre has a very significantly different specific heat load per unit area (kW/m^2) or per unit of rack (kW/rack), different cooling systems can be used according to the needs of individual areas and/or individual systems.

The main technical characteristics of a high efficiency HVAC system are:

- continuous monitoring of temperature and humidity;
- optimized control system based on the actual requests of IT equipment;
- high efficiency air conditioning system components, among which it is worth mentioning:
  - high efficiency motors, pumps and fans with variable speed drives;
  - use the free-cooling system;
  - high EER (Energy Efficiency Ratio) chillers.

12.4.2.1 Cooling Best Practices

Low Density Rooms

The cooling of IT equipment is guaranteed by streams of cold air from the raised floor or ceiling. The cold air is supplied by air conditioning units CRAC (Computer Room Air Conditioning), located on either side of the room, or by centralized systems.

In this case it is of fundamental importance to minimize or eliminate mixing between the cooling air supplied to equipment and the hot air rejected from the equipment. In order to separate the two flows, hot air aisles are created, alternating them with cold air aisles, where air intake vents are placed. To ensure that no mixing of hot and cold air occurs, the hot aisles or the cold ones can be completely closed; panels can be installed to properly delineate the flows.

In addition to a complete separation of the aisles, a series of ducts can be designed with air intake vents to return hot air back to the CRAC. This cooling system is no longer sufficient when the thermal power dissipated per cabinet exceeds 3–4 kW (Figure 12.4).

High-Density Rooms

In the rooms where the installed capacity per cabinet is greater than 5–6 kW the use of dedicated cooling systems, placed near the heat source, has to be evaluated.

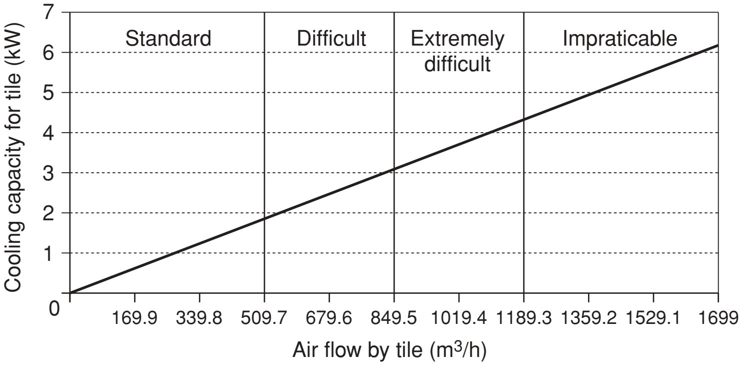

Figure 12.4 Air flow as a function of the cooling capacity

The main technological solutions available as power density increases are as follows:

- cooling systems for aisles: they are air conditioning units that are installed between two racks, integrating with the row; they can be mounted in the vicinity of high-density cabinets to compensate for the cooling that is not guaranteed by the classical configuration;
- cooling systems per unit (or rack): the cooling system (compressor and heat exchangers) is contained in the rack itself, which is hermetically sealed or placed separately on the rack. The heat released in the external environment can be removed or transferred to the HVAC system;
- direct cooling systems on the component within the server. A fluid, usually cool water, passes through the appropriate channels in direct contact with the CPU and electronic components, removing the heat. The main advantage of this system is its high thermal power dissipation; the disadvantages are the danger of loss of fluid in the vicinity of IT equipment and the higher cost.

These systems can be integrated with a traditional low density cooling system or replaced completely.

New IT systems, almost always coupled with virtualization, lead to the formation of ‘hot spots’; this problem can be overcome by the installation near the load of a dedicated air conditioning units.

In these cases the presence of a raised floor for air distribution is not required; this usually implies a higher ceiling with respect to standard buildings.

The main advantages of a high-density rack are: lower operating costs, due to the dedicated cooling, and an increase in the space available in the data centre.

Free Cooling

Free Cooling (FC) allows one in certain conditions to use external air to directly cool the data centre, without using the HVAC system, except for dehumidification.

There are two types of FC:

1. *direct*: where you enter outside air at a low temperature inside the room; air must be properly filtered and humidified or dried to make it suitable for internal requirements;
2. *indirect*: where cooling is achieved by the single passage of the coolant in the cooling tower, without the intervention of the refrigeration compressor.

In both types, although external conditions are not such as to allow a total FC, there is still a further range of temperatures over which you can perform a free pre-cooling, reducing the compressors work required to maintain the set temperature.

The FC and the pre-cooling FC operating hours vary depending on the climate of the place of construction and the internal temperature settings.

Taking the average temperature in different months of the construction site of the data centre, the hours of operation in free-cooling mode, the operating costs and payback time can be estimated.

The FC, especially the indirect FC, can be also introduced as a retrofit in existing data centres, by adding some components in the system, without radically changing the HVAC system configuration.

ASHRAE 2008

T min	18°C
T max	27°C
UR min	Time condensation 5.5°C
UR max	69% time condensation 15°C

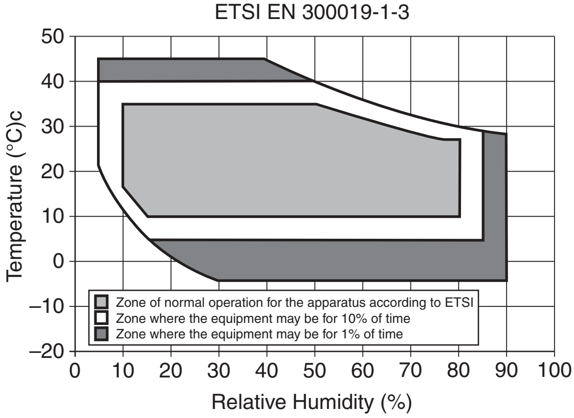

Figure 12.5 Reference values for temperature and relative humidity in data centres defined by ASHRAE and by ETSI

Temperature and Humidity Set Points

The choice of the temperature and humidity set points inside the room are two factors that can significantly affect the efficiency of a data centre.

The energy consumption of HVAC system, which generally accounts for roughly 30% of the total consumption, depends, in fact, on these values.

A rise in temperature may on the one hand allow a reduction in the consumption of the HVAC system and increase the hours of operation per year of the free-cooling system if one is present; on the other hand it reduces the time available for intervention in case of a downshift of the HVAC system and may increase the consumption of IT equipment for the operation of internal fans, which depends on the temperature set point in the bios of the machines.

The temperature set point optimization can be an energy saving opportunity but it requires careful analysis.

The reference values for temperature and relative humidity in the data centres defined by ASHRAE (American Society of Heating, Refrigerating and Air Conditioning) and by ETSI (European Telecomunications Standards Institute) in its standard EN 300019-1-3, are summarized in Figure 12.5.

12.5 DG and CHP for Data Centres

The use of DG and CHP in data centres result in cost savings to the facility operator in the form of:

- reduced energy-related costs and enhanced economic competitiveness – from reduced fuel and electricity purchases, resulting in lower operating costs;
- increased reliability and decreased risk from outages – as a result of reliable on-site power supply; and
- increased ability to meet facility expansion timelines – by avoiding the need for utility infrastructure upgrades.

Because of their very high electricity consumption, data centres have high power costs. Installation of CHP systems with absorption cooling can often reduce energy costs by producing power more cheaply on-site than can be purchased from the utility supplier. In addition, waste heat from the power generation can drive absorption chillers that displace electric air conditioning loads.

Data centres require both high-quality and extremely reliable power. Of all customer types, data centres, telecommunication facilities, and other mission-critical computer systems have the highest costs associated with poor power quality.

Data centres almost always have UPS systems to condition power and eliminate power disturbances. Battery backup is generally used to provide a short-term outage ride-through of a few minutes to an hour. Longer-term outages are typically handled with standby diesel generators.

On-site power generation, whether it is an engine, fuel cell, microturbine or other prime mover, supports the need for reliable power by protecting against long-term outages beyond what the UPS and battery systems can provide. DG/CHP systems that operate continuously provide additional reliability compared with emergency backup generators that must be started up during a utility outage. Backup generators typically take 10–30 seconds to pick up load in the case of an outage.

Developing on-site power sources gives data centre operators increased flexibility in both the expansion and design of new facilities. Upgrading older, smaller data centres with new equipment can result in a large increase in power demand to the facility that the utility might not be able to meet in the near term. Incorporating continuous prime power DG/CHP options can facilitate expansion and facility development on a more rapid schedule than can sometimes be possible by relying solely on the existing utility grid. Minimizing external power demand also reduces additional utility infrastructure requirements and associated costs that might be required for new or expanded facilities.

12.6 Organizing for Energy Efficiency[1]

To engineer the migration from a traditional energy-consuming data centre to a modern energy efficient data centre properly requires an organizational alignment that facilitates such a migration.

Figure 12.6 illustrates an IT organizational structure that integrates the expertise of personnel who understand both IT systems and physical infrastructure systems. The new organizational wrinkle involves the integration of an IT facilities arm into the rest of the IT organization.

This organizational alignment presents several advantages. For years, IT and facilities departments have operated as separate entities and evolved separate cultures and even used separate languages.

As a result, most data centre design/build or upgrade projects are painful, lengthy, and costly.

This new IT facilities group is a separate group from the traditional ‘building’ facilities group. The IT facilities group acts as a liaison between IT and the facilities building group, but is under the direct control of IT.

[1] Green Grid, White Paper no. 1 *Guidelines for Energy-Efficient Data Centres*.

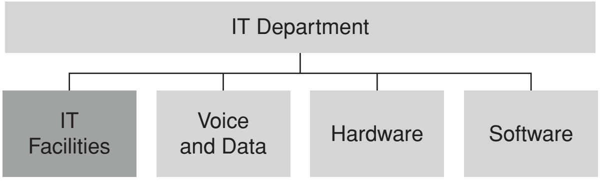

Figure 12.6 IT organizational structure proposed by GreenGrid

The IT facilities group addresses data centre issues specific to hardware planning, electrical deployment, heat removal, and physical data centre monitoring.

This organizational alignment allows a data centre team to rapidly deploy an energy efficient data centre upgrade policy that addresses both IT systems and physical infrastructure systems.

Further Readings

M. Bramucci, D. Di Santo and D. Forni, Linee guida per la progettazione di datacenter ad alta efficienza, FIRE, 2011.
European Commission Joint Research Centre, *2011 Best Practices for EU Code of Conduct on Data Centres*.
The Green Grid. Library and Tools, http://www.thegreengrid.org
PG&E, *High Performance Data Centres: A Design Guidelines Sourcebook*, 2006.
US DoE, *FEMP Best Practices Guide for Energy-Efficient Data Center Design*.