

1

Data Centre Overview

Eduard Oró and Jaume Salom

Catalonia Institute for Energy Research – IREC, Spain

1.1 Data Centre Infrastructure

1.1.1 Introduction

The rapid increase of cloud computing, high-performance computing (HPC) and the vast growth in Internet and Social Media use have aroused the interest in energy consumption and the carbon footprint of Data Centres. Data Centres primarily contain electronic equipment used for data processing (servers), data storage (storage equipment) and communications (network equipment). Collectively, this equipment processes, stores and transmits digital information and is known as information technology (IT) equipment. Data Centres also usually contain specialized power conversion and backup equipment to maintain reliable, high-quality power as well as environmental control equipment to maintain the proper temperature and humidity for the IT equipment [1]. Data Centres are commonly divided into 3 different spaces (Figure 1.1):

- **IT room or data hall.** The IT room is an environmentally controlled space that houses equipment and cabling directly related to compute and telecommunication systems which generate a considerable amount of heat in a small area. Moreover, the IT equipment is highly sensitive to temperature and humidity fluctuations, so a Data Centre must keep restricted power and cooling conditions for assuring the integrity and functionality of its hosted equipment. Thus, many manufactures call the IT room as whitespace.
- **Support area.** These areas are those where different systems such as power, cooling and telecommunications are located. Basically, these areas are electrical gear and uninterruptible power supply (UPS) control,

2 Data Centre Overview

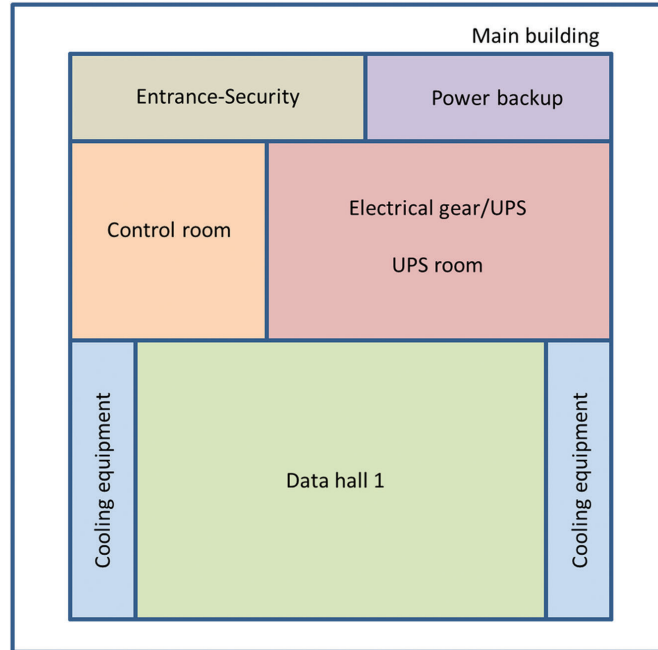


Figure 1.1 Scheme of a typical Data Centre layout.

UPS room, power backup system, cooling equipment areas, control room, etc.

- **Ancillary spaces.** These are those such as offices, lobby and restrooms.

1.2 Main Subsystems

Data Centres typically comprise three main subsystems: IT equipment, which provides services to customers; power infrastructure, which supports the IT and cooling equipment; and cooling infrastructure, which removes the generated heat by the IT and power equipment.

1.2.1 IT Equipment

IT equipment is the main contributor to electricity consumption of a Data Centre, representing about 45–55% of the total and can be described as the IT work capacity employed for a given IT power consumption [2]. Thus, the IT equipment power includes the load associated with all of the IT equipment, such as compute, storage and network equipment, along with supplemental

equipment such as monitors, workstations/laptops used to monitor or otherwise control the Data Centres. Table 1.1 shows the IT equipment components and its classification.

1.2.2 Power System

The power distribution system of a Data Centre is intended as the equipment in charge of bringing electrical power to the loads of the system fulfilling adequate power quality levels and security of supply. In this sense, reliable power supply is one of the cornerstones of a successful Data Centre. Since the public grid could have short voltage dips which can cause servers to malfunction or even crash completely, it is essential to stabilize the power supply. Additionally, there is always the danger of longer power outages, which can crash the entire Data Centre. Figure 1.2 shows a scheme of the power supply for a standard Data Centre from the main grid to the IT equipment. Between the public grid and the infrastructure there is first the transformer which is responsible to transform electrical energy from high/medium voltage to low voltage. After this electrical element, there is the switchgear, which is used to control, protect and isolate electrical equipment. In the framework of Data Centres, it is responsible to connect the IT equipment to the main grid or to the backup diesel generator. As mentioned, usually there is a backup diesel generator to supply power in case that the public grid presents a problem and the battery or flywheel is usually designed to be sufficient to keep the IT equipment running until the diesel generators can be started (few seconds). Moreover, there are other components that support the IT equipment such as the UPS, switchgear and power supply units (PDU), and other miscellaneous component loads such as lighting and filters.

Table 1.1 IT equipment components

Servers	Compute devices	Servers
Networking	Network devices	Switches
		Routers
		Printers
	IT support systems	PCs/workstations
		Remote management
	Miscellaneous devices	Security encryption, storage encryption, appliances, etc.
Storage	Storage	Storage devices – switches, storage array Backup devices – media libraries, virtual media libraries
	Telecommunication	All Telco devices

4 Data Centre Overview

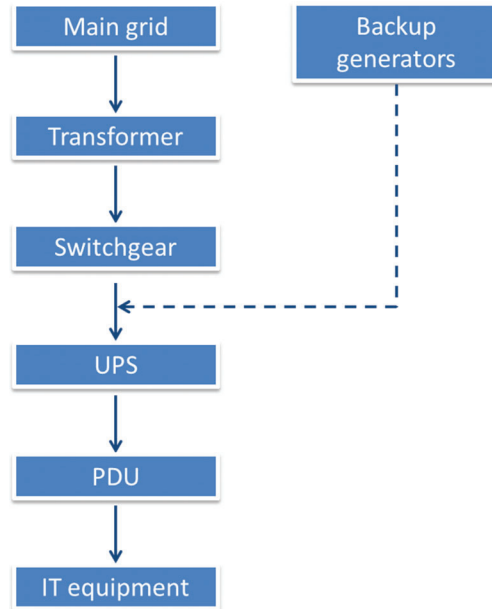


Figure 1.2 Scheme of the power supply in a standard Data Centre.

There are many different UPSs which can vary greatly in physical size, weight, storage capacity, supported input power, cost, etc. Here, the most known ones will be presented and described being divided between static and rotary UPS. The UPSs which do not have any moving parts throughout its power path are the so-called static UPS. On the other hand, when there are rotating components such as a motor generator within the UPS to transfer power to the load, they are called rotary UPSs. Another fundamental difference between them is that static UPS uses chemical stored energy (in the batteries) in the event of a power cut whereas rotary UPS uses the stored kinetic energy of a spinning flywheel. Rotary UPSs have relatively low redundancy (typically between 15 and 30 seconds) on their own compared to the static UPS. However, this low redundancy allows sufficient time for the generator to start up and accept the full load without affecting the output to the Data Centre. The efficiency performance of static UPSs is higher than for rotary ones, especially at partial loads. This means that the rotary UPSs sustain higher fixed losses such as the energy utilized to power controls, flywheels and pony motors associated with the rotary UPS at zero load and the energy utilized to preheat the engine coolant and lubrication. In some cases, at full load, the efficiency of the rotary UPS can be higher [3]. However, technology is moving

faster and the efficiency of these elements is improving day by day; therefore, for each particular case, a comparison study between both technologies should be done.

Figure 1.3 shows different approaches for distributing power to IT racks seen in Data Centres today. Mainly, there are three kinds of power distribution: panel-board distribution, traditional PDU distribution and modular distribution. In panel distribution, the main Data Centre power is distributed to multiple wall-mounted panel boards. Wall mount panel boards are a very low-cost power distribution strategy, made up of parts that an electrician can acquire in days, not weeks, to implement quickly. In traditional power distribution systems, the main Data Centre power is distributed to multiple PDU throughout the IT space. Moreover, there are two main categories of traditional PDU systems: field-wired using power cables in cable trays or in flexible or rigid conduit, distributed beneath the raised floor or above the IT racks. Finally, in order to meet the modern IT requirements, alternative power distribution approaches are appearing in Data Centres. These approaches are more flexible, more manageable, more reliable and more efficient. Panel-board distribution and field-wired traditional PDU distribution systems are shown to be the best approach when low first cost is the highest priority, when the IT space has unique room constraints, and when IT changes are

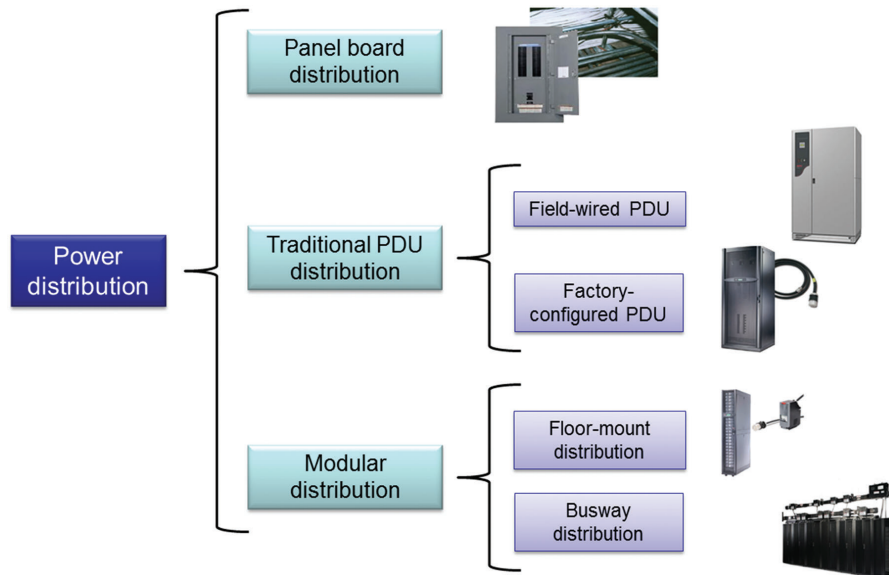


Figure 1.3 Five approaches for distributing power to IT racks.

not likely. Factory-configured traditional distribution systems are shown to be the best approach when a Data Centre requires portability of its equipment, when additional pods may be added in future, and when low first cost is still a priority. Modular distribution approaches allow for more flexibility, more manageability, increased reliability and increased efficiency, to better address the IT needs of many Data Centres today. Bus way is optimal when floor space is a constraint, when there is a large open floor plan with a well-defined IT equipment layout, and when there is a high confidence in the final capacity need. It is also ideal when there is a high frequency of IT equipment turnover requiring new circuits. Floor-mount modular distribution, on the other hand, is optimal when the Data Centre has an uncertain growth plan and locations are not precisely defined in advance, because it offers flexibility to place units in specific locations when the need arises. It is also best for Data Centres that are being retrofit with additional capacity (i.e. addition of a high-density zone) [4].

1.2.3 Cooling System

The physical environment inside the Data Centres is rigorously controlled, and therefore, the design of the cooling system is critical to the efficiency and reliability of the whole Data Centre. There are mainly two types of cooling: air-cooled systems and water-cooled systems.

The majority of existing Data Centres are refrigerated through air. The server racks are normally arranged in cold and hot aisle containments to improve the air management. In these systems, the chilled air produced by the CRAH unit is driven into the cold aisle, either through the floor plenum and perforated tiles or through diffusers in the ceilings. The warm air in hot aisles is captured and returned to the intake of the CRAH. The heat in the CRAH unit is absorbed by water into a chiller system. There is another air-cooling system which consists on in-row cooling, so the CRAH unit (and therefore the chilled water) is embedded next to the racks and the hot air does not flow through the whitespace but inside the containment. Figure 1.4 shows different kind of air cooled Data Centres: cold aisle containment, hot aisle containment and in-row cooling systems.

Some new Data Centre designs have power loadings to levels that are difficult to remove with CRAH units. Therefore, other cooling techniques, such as on-chip cooling (either single- or double-phase liquid systems) and submerged cooling systems are also used. When on-chip cooling is adopted, not all the heat from the server is absorbed by the liquid, and therefore,




Figure 1.4 Cold aisle containment, hot aisle containment and in-row cooling [5].

an auxiliary air-cooling system is also needed. The use of this technology raises a concern about leaks and can cause irrevocable damage if it comes into direct contact with IT equipment. This is an issue because maintenance, repair and replacement of electronic components result in the need to disconnect and reconnect the liquid carrying lines. To overcome this concern, it is possible to use nonconductive liquid such as refrigerant or a dielectric fluid in the cooling loop for the IT equipment. Thus, the processors can be directly immersed in these liquids, which improve thermal efficiencies and often results in simplified liquid cooling systems. Therefore, submerged cooling systems are able to absorb all the heat from the servers through a dielectric refrigerant. These cooling technologies can provide high performance while achieving high energy efficiency in power densities beyond air-cooled equipment and simultaneously enabling use of waste heat when supply facility water temperatures are high enough. Moreover, liquid cooling can also offer advantages compared to air in terms of heat exchange (since it has much higher heat conduction coefficient of $0.6 \text{ W}/(\text{m}\cdot\text{K})$), lower noise levels and closer temperature control. Figure 1.5 shows a picture of each of the liquid cooled technologies available nowadays.




Figure 1.5 On-chip water cooling, on-chip two-phase cooling system and submerged cooling system [6].

1.3 Data Centre Archetypes

To provide a better overview of the Data Centre market, Data Centres are itemised based on the following characteristics:

- Function or objective of the Data Centre
- Size
- Location and surroundings

Figure 1.6 shows the distribution of a Data Centre in function of the business model, size and location and surroundings.

1.3.1 Function or Objective of the Data Centre

Based on the European Standard EN 50600-1:2012, there are three types of Data Centres: enterprise Data Centres, co-location Data Centres and hosting Data Centres. The main difference between these types of Data Centres is the owner of the Data Centre facility and the IT, and with that the possibility

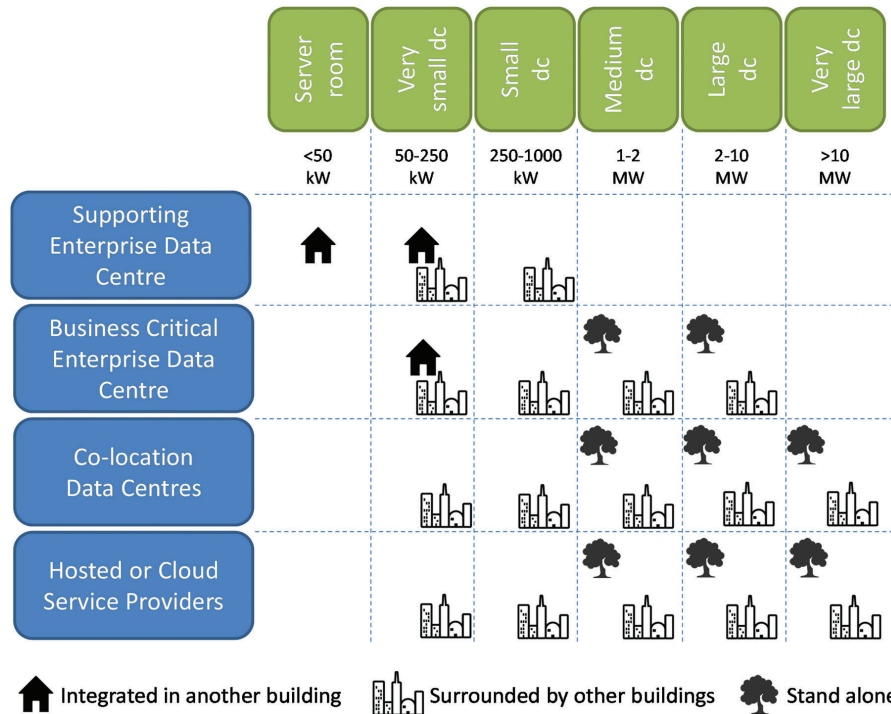


Figure 1.6 Overview of Common Data Centre Archetypes based on size (IT power), Data Centre function or objective and surrounding [7].

to synchronize the Data Centre systems and the IT. However, the enterprise Data Centre group can also be divided into two subgroups for a more clear differentiation: the business supporting Data Centres and the business critical Data Centres. Thus, this results in the following four main groups:

1. **Enterprise** Data Centres. These are built and operated by the end user, which are banks, hospitals and retailers, including HPC at universities or research institutions. For enterprises, a Data Centre can have two distinctive functions as will be described below.
 - (a) **Business supporting** Data Centre. Firstly, the function of the Data Centre is to support the activities of the firm. For instance, a contractor uses a Data Centre to save and backup all the designs, drawings, 3D modelling activities and administration. Similarly, a university uses its Data Centre to perform among others, (high performance) research activities, backup and storage facilitation and administrative computing. In general, these Data Centres will provide safe, secure and reliable hosting facilities for the firms core IT systems. Since the Data Centres are not leading, but supporting, they are most frequently situated close to the actual firm or organisation, and therefore at short distance of the actual activities.
 - (b) **Business Critical** Data Centre. They are an integral part of the main business process. These are, for example, the commercial telecom Data Centres and Data Centres of financial institutions. A bank uses its Data Centre for all the transactions, or telecom providers use the Data Centres for their telecom data transmissions. For these types of enterprises, the Data Centre is a leading necessity within their business process. Therefore, these Data Centres are situated at locations that are beneficial for the IT processes, based on criteria such as (not limited) distance to the customers, distance to a (large) power plant, cost and availability of land, (transatlantic) glass fibre connectivity or carrier neutrality options.
2. **Co-location** or commercial/multi-tenant Data Centres. In a co-location Data Centre multiple customers are housed with their own IT networks, servers and equipment. The co-location Data Centre provides the supporting infrastructures to its customers (mainly security, power and cooling).
3. Hosted or **Cloud** service providers. Third-party organisations built and operate these facilities. They provide direct computing services to other

businesses. This class of Data Centres is important as it includes some of the largest names in IT (Google, Microsoft, Amazon) and some of the largest facilities in terms of square meters and numbers of servers. For these Data Centres, the IT is completely leading, so when deciding the location, factors such as the energy supply, Internet connection and costs of the plot are leading.

1.3.2 Size

IT power capacity is commonly used as the way to express the size of the Data Centre. The following breakdown based on size is identified:

1. Server room <50 kW
2. Very small Data Centre 50–250 kW
3. Small Data Centre 250–1000 kW
4. Medium size Data Centre 1–2 MW
5. Large Data Centre 2–10 MW
6. Very large Data Centre >10 MW

1.3.3 Location and Surroundings

The location and the surroundings of the Data Centre have influence on the possibilities of collaborating with its surroundings. When a Data Centre is close to offices, this creates more possibilities for these offices to use the heat of the Data Centre. However, when a Data Centre is not surrounded by any buildings, there is generally more space and less regulatory conditions for large-scale application of renewables such as windmills and solar panels. When looking at location and surroundings, the following parameters are distinguished:

- **Integrated** in another building (e.g. office)
- **Not integrated** in another building:
 - **Surrounded** by other buildings
 - **Stand alone**

1.3.4 Archetypes Definition

Based on this overview of Data Centre characteristics, the following archetypes of the Data Centres are detailed. Figure 1.7 shows the specified Data Centre archetypes. These are further specified in the text below:

1. Very small (less than 250 kW) Data Centre or server room inside a building. The Data Centre supports the activities of the enterprise.
2. Very small (less than 250 kW) Data Centre in a separate building surrounded by other buildings. The Data Centre supports the activities of the enterprise.
3. Small (between 50 and 1000 kW) Data Centre inside a building or in a separate building surrounded by other buildings. The Data Centre is critical to the business.
4. Medium to large size (between 1 and 10 MW) Data Centre completely stand-alone. The Data Centre is critical to the business.
5. Medium to large size (between 1 and 10 MW) Data Centre in a separate building surrounded by other buildings. The Data Centre is critical to the business.
6. Small (between 50 and 1000 kW) co-location or hosting Data Centre surrounded by other buildings.
7. Medium to very large size (between 1 and >10 MW) stand-alone co-location or hosting Data Centre.

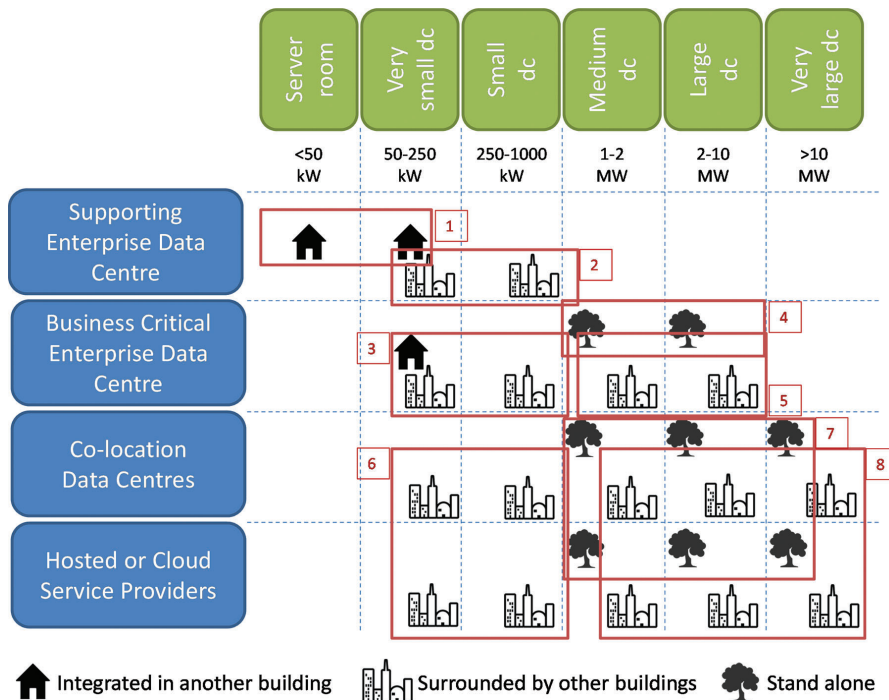


Figure 1.7 Data Centre archetypes defined by RenewIT project [7].

8. Medium to very large size (between 1 and >10 MW) co-location or hosting Data Centre surrounded by other buildings.

1.4 Workload Typology

The power consumption of a Data Centre depends on the current computing load and therefore the usage patterns of the IT equipment. In real Data Centres, there are disparities in performance and power characteristics across servers and different scheduling, task migration or load balancing mechanisms have effects on the IT power consumption and therefore cooling consumption. For simplicity, and according to the service that each Data Centre type provides, three main homogeneous IT workloads can be identified: web, HPC and data workload.

1.4.1 Web Workloads

Web workload has real-time requirements, the users of such workload need to get a response to their petitions in few seconds (i.e. Google search, Facebook surf, etc.). Therefore, the evaluation of web application resource consumption requires realistic workload simulations to obtain accurate results and conclusions. Although networking workloads are very heterogeneous and they change dramatically over relatively short time frames, they also have general traits that are common to all workloads, such as the daily cycle of activity. Characterizing the workload on the worldwide web has two components: the data available on servers and how clients access that data. There is no typical resource consumption profile for web workloads but they may use CPU, memory, network or disk in several proportions. In order to make a complete characterization of both client workloads and resource consumption, real production logs were taken from a top Online Travel Agency (OTA) in [8], observed in the 35+ physical node cluster in which the application was deployed. Figure 1.8 shows two traffic patterns for a dataset, including number of hits relative to peak load (a) and number of user sessions started (b) over one week, grouped in 30-minute periods. As it can be observed, the traffic decreases over the night, until it starts growing again soon in the morning. It keeps growing until noon, when it slightly decreases. Finally, the workload intensity starts increasing again over the afternoon until it reaches its maximum around 9:00 pm. Over the night, the traffic volume decreases until it finally reaches the beginning of the cycle again. Therefore, web workload has the particularity to follow a daily/weekly pattern.

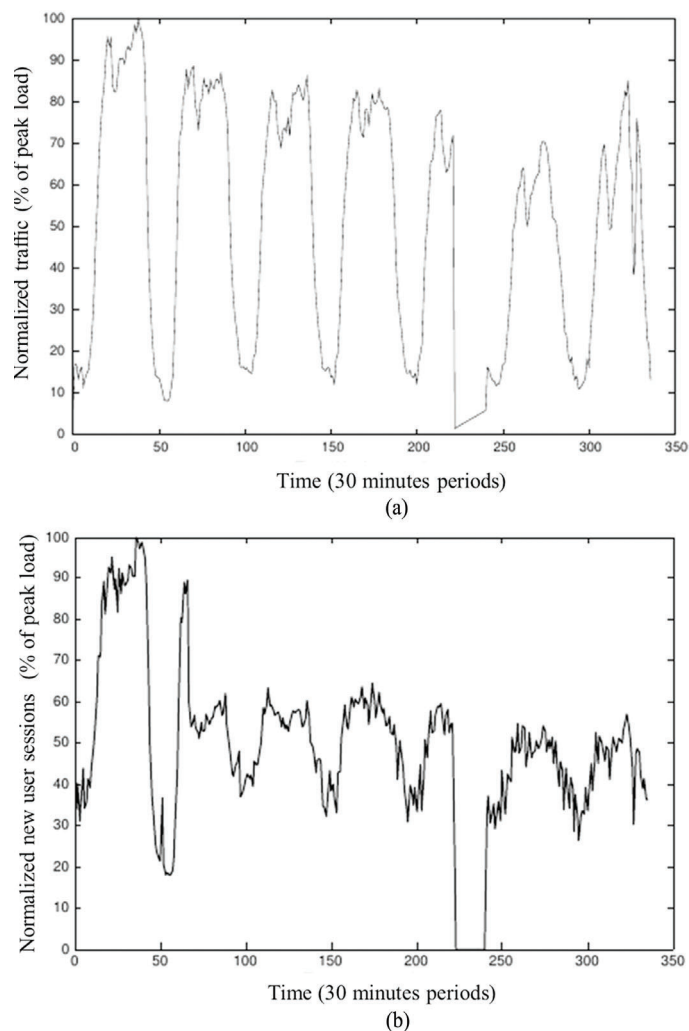


Figure 1.8 Typical profile for the traffic volume intensity during one week [8].

1.4.2 HPC Workloads

HPC workload is typically CPU intensive. They perform a large amount of floating-point operations for scientific calculations. Because HPC workloads may last for hours, or even days, they do not have real-time requirements, and they are usually allocated in job queues that may execute them hours or days after they are submitted by the users. Figure 1.9 shows power profiles of

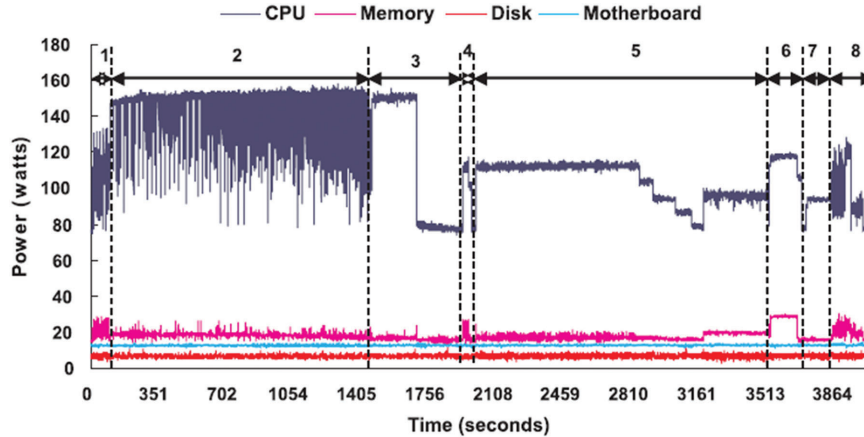


Figure 1.9 Snapshot of the HPC power profile. The entire run of HPC consists of eight micro benchmark tests in this order: (1) PTRANS; (2) HPL; (3) Star DGEMM + single DGEMM; (4) Star STREAM; (5) MPI RandomAccess; (6) Star RandomAccess; (7) Single RandomAccess; (8) MPI FFT, Star FFT, single FFT and latency/bandwidth [8].

the HPC benchmarks in a cluster made up of nine dual-core, Opteron-based server nodes. The power consumption of major components per compute node was obtained when running a full HPC benchmark suite using eight cores. These profiles are obtained using the problem size where HPL achieves its maximum performance on two nodes. In this test, the problem size fits the peak execution rate for HPL. For LOCAL tests, a benchmark on a single core with three idle cores was performed. Power consumption is tracked for major computing components including CPU, memory, disk and motherboard. These four components capture nearly all the dynamic power usage of the system that is dependent on the application. Notice that HPC workloads do not follow a given pattern, and they will depend on the access policy and dimension of the Data Centres of each institution.

1.4.3 Data Workloads

Data workload is usually both memory and disk intensive, while they can also use a high rate of CPU operations for data analysis. Despite the fact that data workloads may have real-time requirements (i.e. a search query in Google), data workloads can also have no real-time requirements (i.e. background data analytics for business intelligence applications). Obtaining accurate information on disk drive workloads is a complex task, because

disk drives are deployed in a wide range of systems and support a highly diverse set of applications. The dynamics of file access in distributed systems are interesting in the sense that caching may cause the characteristics of the workload to look different at the client and at the server level. For example, locality is reduced near the servers because caches filter out repetitions, while merging of different request streams causes interference [9]. Notice that data workloads do not follow a given pattern, and they will depend on the access policy and dimension of the Data Centres of each institution.

1.4.4 Consumption versus Workload Typology

In order to predict the Data Centre consumption from the IT load, a relationship between server usage (IT load) and server consumption is needed. Firstly, it is needed to define a proper workload. In the framework of the RenewIT project, three homogeneous workloads were studied. Web workload is a real pattern collected from the access log of an ISP within the UPC [10], while HPC and data workloads patterns are extracted from the CEA-Curie Data Centre which are publicly available in the Parallel Workloads Archive [11]. Figure 1.10 presents the three homogeneous IT load profiles.

Secondly, the definition of IT load is an additive function that considers the load rates of CPU, main memory, disk and network, pondered according to the measured impact of each term in late 2012 servers [13]. Firstly, different types of micro-benchmark for fully stressing the system were executed in

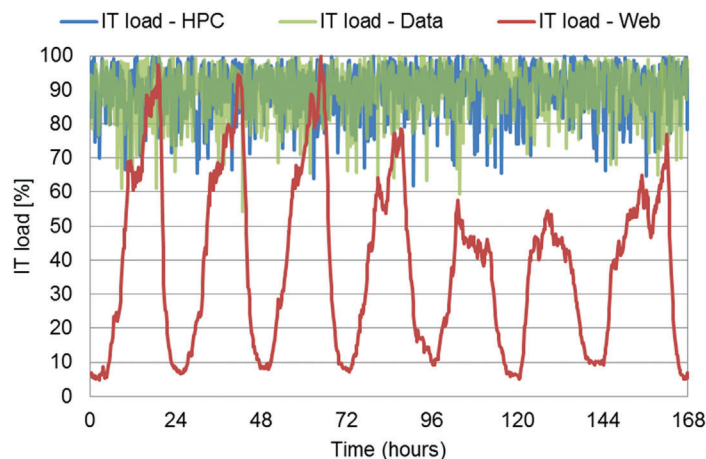


Figure 1.10 Different IT workload profiles (web, HPC and data) during a week [12].

order to reach the maximum real power of the system. These benchmarks included Ibench suite [14], Stress-ng [15], Sysbench [16], Prime95 [17], Pmbw [18], Fio [19] and Iperf3 [20]. After this initial process, different benchmarks based on real-world software stacks from CloudSuite [21] for web and data benchmarks, and NAS Parallel Benchmarks [22] for HPC were also executed.

With this experimentation, the relation between IT load and power consumption can be drawn for different server architecture. Notice that for its further adaptation to other hardware, these correlations were normalized. Figure 1.11 shows the results of the experimentation and the regressions to predict different consumptions in function of the IT load. The y-axis shows the percentage of the power referred to the maximum detected. The x-axis shows the percentage of the IT load referred to the maximum executed.

The variability in the power/load measurements shows that there is not a generic power profile for software, because all the components of a host (CPU, memory, disk and network) do not work independently. They must be coordinated because there are dependencies between data and procedures (and the usage of resources is variable across the same execution, depending of the availability of their required inputs at a given time). Finally, Figure 1.12 shows the results from RenewIT project compared to a standard function [23]. Notice that at low IT load, the power consumption is similar between the standard regression and the ones found in the project. However, at higher loads, especially above 80% of the IT load the difference becomes important. This means that while some researchers or Data Centre operators assume that at 100% of the IT load, for instance 100% CPU usage, the server consumes 100% of the nominal power, the reality is different.

1.5 Redundancy Level

1.5.1 Basic Definitions

Redundancy is an operational requirement of the Data Centre that refers to the duplication of certain components or functions of a system so that if they fail or need to be taken down for maintenance, others can take over. Redundant components can exist in any Data Centre system, including cabling, servers, switches, fans, power and cooling. It is often based on the “N” approach, where “N” is the base load or number of components needed to function. N+1 means having one more component than is actually needed to function, 2N means having double the amount of total components, and 2N+1 is having

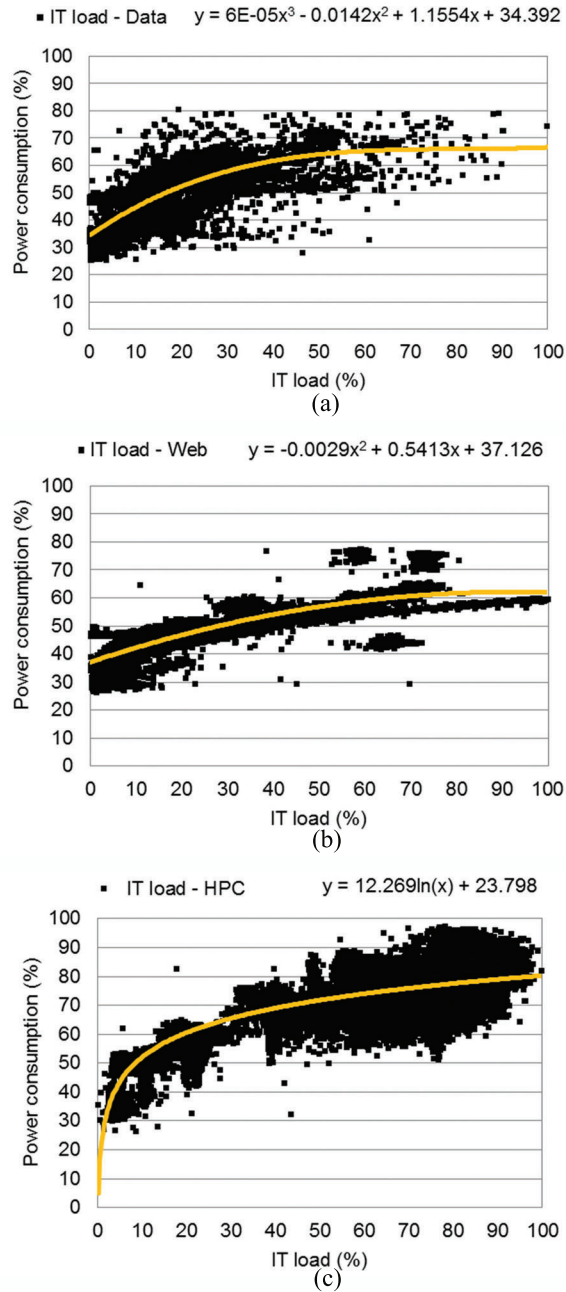


Figure 1.11 Curves for (a) data, (b) web and (c) HPC workload and their corresponding regression. Note that y means power consumption while x IT load in %.

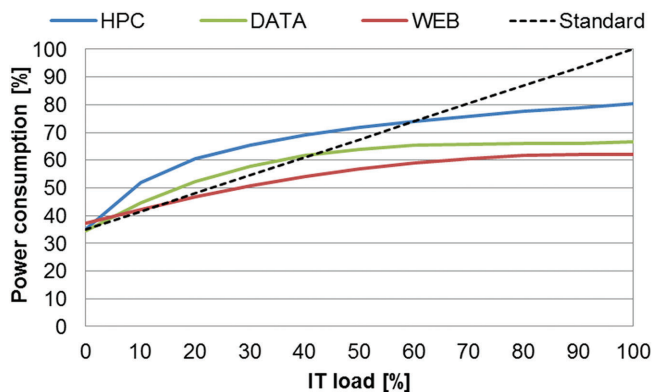


Figure 1.12 IT load relationship with power consumption.

double the amount plus one. While redundancy can be viewed as a “planned” operational function, availability is based on “unplanned” downtime and relates specifically to how many minutes or hours can be tolerated. For example, with 8,760 hours in a year, an availability of 99.9% indicates the ability to tolerate 8.76 hours of downtime per year. If the estimated cost of downtime within a specified time unit exceeds the amortized capital costs and operational expenses, a higher level of availability should be factored into the Data Centre design. If the cost of avoiding downtime greatly exceeds the cost of downtime itself, a lower level of availability should be factored into the design. Finally, reliability is the combination of redundancy and availability that go into determining reliability, and there are a variety of industry ratings that indicate reliability: Uptime Institute, TIA-942 and BICSI-002-2010. The Uptime Institute’s Tier rating is the most common standard used and therefore it is described below.

1.5.2 Tier Levels

The Uptime Institute [24] defines a system of “Tiers” based upon business objectives and acceptable downtime which determines the redundancy of energy and cooling equipment. Moreover, TIERS levels can be described as follows:

- **Tier I:** Basic site infrastructure (N). It has non-redundant capacity components (components for cooling and electrical storage) and a single, non-redundant distribution path serving the computer equipment. It is estimated an availability of $\sim 99.671\%$ and *28.8 hours of downtime per year*.

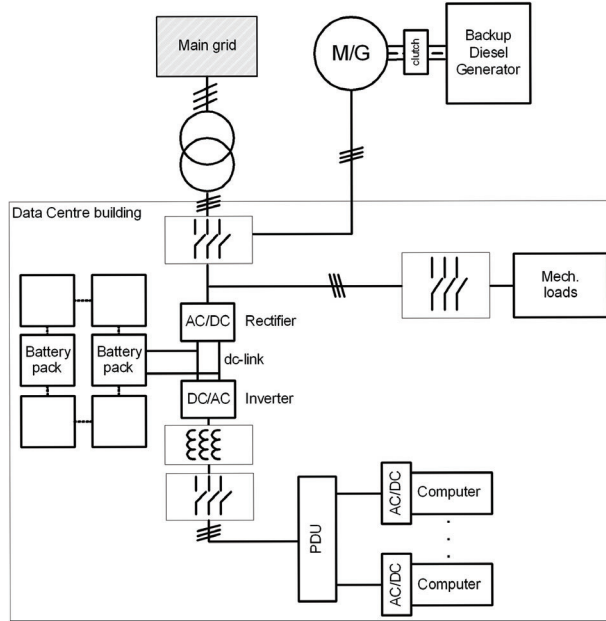


Figure 1.13 Power distribution system for Tier I considering static UPS (battery pack) [25].

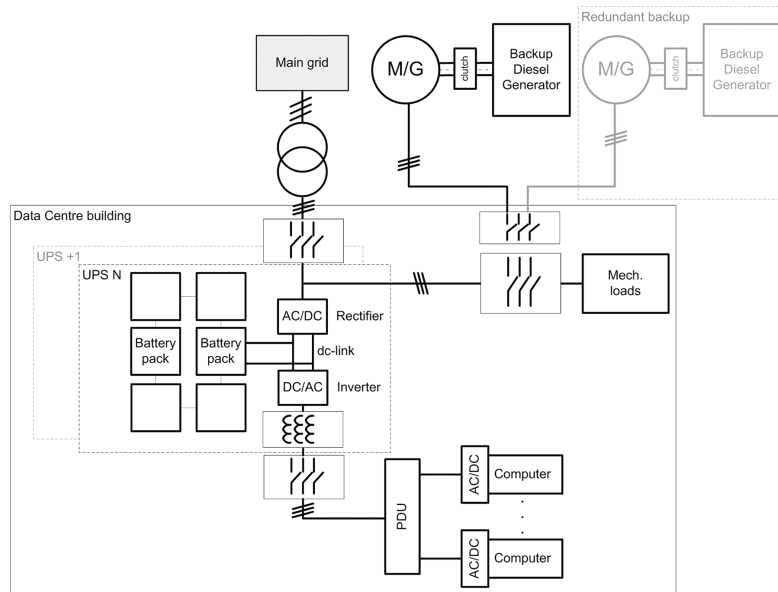


Figure 1.14 Power distribution system for Tier II considering static UPS [25].

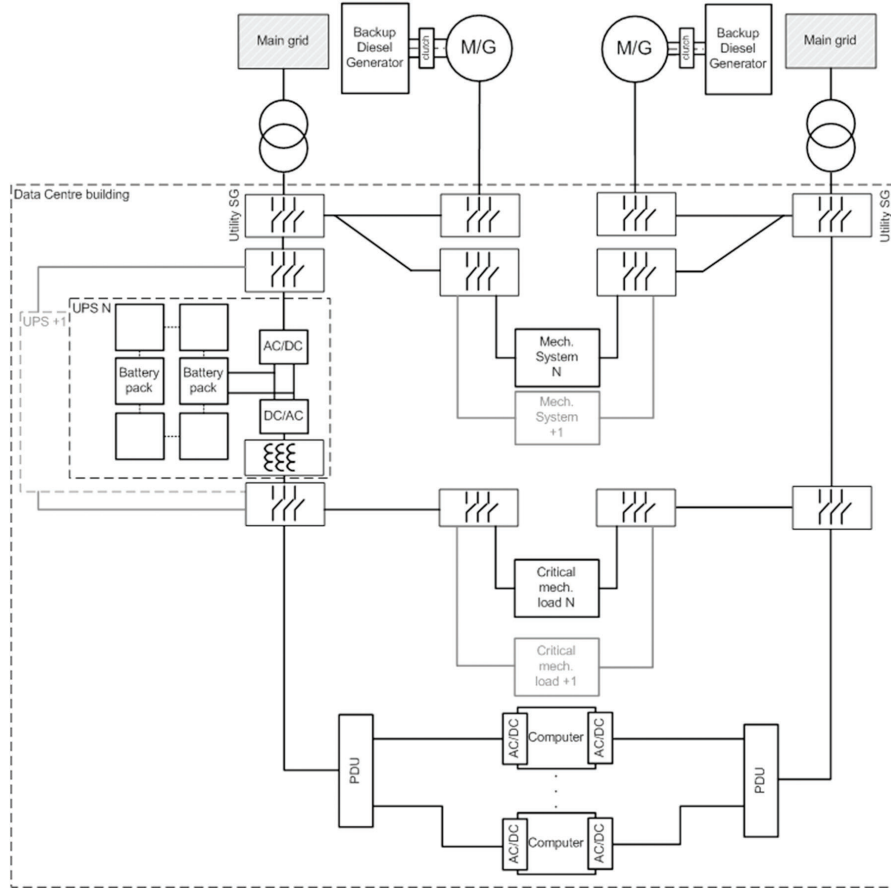


Figure 1.15 Power distribution system for Tier III [25].

- **Tier II:** Redundant site infrastructure capacity components (N+1). It has redundant capacity components and a single, non-redundant distribution path serving the computer equipment. It is estimated an availability of $\sim 99.741\%$ and *22 hours of downtime per year*.
- **Tier III:** Concurrently maintainable site infrastructure (2N). It has redundant capacity components and multiple independent distribution paths serving the computer equipment. Only one distribution path is required to serve the computer equipment at any time. It is estimated an availability of $\sim 99.982\%$ and *1.6 hours of downtime per year*.

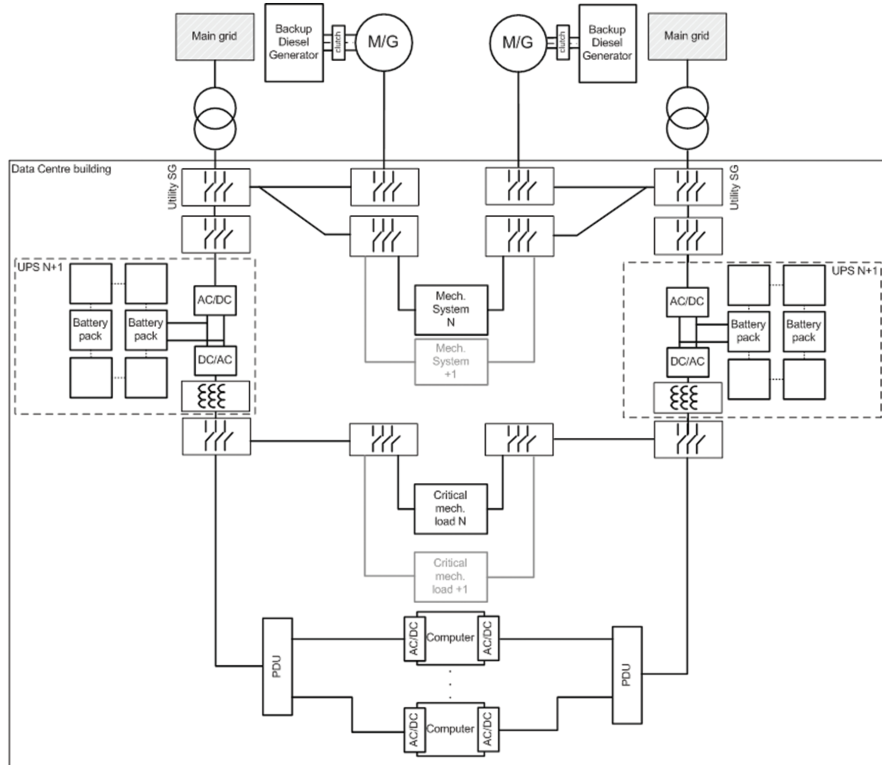


Figure 1.16 Power distribution system for Tier IV [25].

- Tier IV:** Fault tolerant site infrastructure $2(N+1)$. It has multiple, independent, physically isolated systems that provide redundant capacity components and multiple, independent, diverse, active distribution paths simultaneously serving the computer equipment. The redundant capacity components and diverse distribution paths shall be configured such that “N” capacity providing power and cooling to the computer equipment after any infrastructure failure. It is estimated an availability of $\sim 99.995\%$ with 0.4 hours of downtime per year.

In general terms, Tier I corresponds to the basic structure and has non-redundant capacity components and distribution paths. Tier II is quite similar to Tier I, just including redundancies in capacity components but not in power distribution paths. Tiers III and IV contain redundancies both in capacity components and distribution paths, with Tier IV being the only fault tolerant

site infrastructure. Therefore, failure or maintenance of any component will not affect computer equipment. Attending to this brief description of Tier levels classifying power distribution systems of Data Centres, electrical schemes for Tier I, Tier II, Tier III and Tier IV are presented in Figures 1.13–1.16, respectively.

1.6 Future Trends

Global power demand for Data Centres grew to an estimated 40GW in 2013, an increase of 7% over 2012 [26]. Most experts agree that this figure will continue rising in future despite improvements in IT performance as well as Data Centre design and operation. However, a recent report from NREL estimates almost no increasing in the United States Data Centre industry electricity consumption. These predictions always should be taken carefully since it is really difficult to predict the future trends. What for sure will increase is the Data Centre demand due to the arrival of Internet of Things, Internet and Social Media increasing demand from emerging countries.

Even though the most common Data Centres are small company-owned facilities, cloud computing has been driving major changes across the entire IT industry for nearly a decade, and its impact has been accelerating with every passing quarter. Many consultants predict that the impact of cloud and the growing role of hyper-scale operators are existential and severe and more work will go to co-location, hosting and cloud companies. That means there will be more IT and Data Centres in proportionately fewer hands. Figure 1.17 illustrates how capacity is shifting in terms of Data Centre space.




Figure 1.17 Global Data Centre space distribution [26].

A small but growing number of influential Data Centre owners and operators have seen financial, sustainability and availability benefits from increasing the proportion of energy they use from on-site and grid-based renewables such as wind, solar, hydropower and biofuels. On one hand, the average Data Centre is still likely to only derive a small percentage of its energy from grid or on-site renewables, but a handful of large Data Centres companies are skewing that trend. On the other hand, big players such as Google, Apple, Facebook and eBay have a track record of adopting disruptive approaches to Data Centre design and operation. Investment in renewable energy is another example of this disruptive approach in action.

Data Centre use of renewables can be broadly categorized as either from the grid or on-site. However, the reality is more nuanced, with a variety of




Figure 1.18 Overview of possible renewable supply options for Data Centre industry [27].

different approaches used to bring renewable energy into facilities. These range from generating renewables directly on-site to purchasing renewable energy credits from grid suppliers (see Figure 1.18). These different approaches can be broken down into four main groups:

- **On-site generation from on-site renewables:** Renewable energy is generated directly on-site through the use of solar panels or wind turbines.
- **On-site generation from off-site renewables:** Renewable energy is generated on-site but the fuel such as biomass, natural gas or biogas has to be transported to the facility.
- **Off-site generation:** Renewable energy is purchased from a nearby renewable energy provider. Technically, this could also be considered on-site generation depending on the proximity of the renewable generation source.
- **Renewables from third parties:** This includes purchasing Renewable Energy Guarantees of Origin (REGOs) in Europe or the US equivalent Renewable Energy Certificates (RECs). Other mechanisms include Power Purchase Agreements (PPAs) and Renewable Energy Tariffs (RETs).

References

- [1] Shehabi, A. e. a. (2016). United States data center energy usage report. Ernest Orlando Lawrence Berkeley National Laboratory, Berkley, Californig. LBNL-1005775.
- [2] IBM, Rear door exchanger planning guide, 2005–2008.
- [3] Cottuli, C. Comparison of static and rotary UPS, White Paper 92. Schneider Electric.
- [4] Rasmussen, N., and Torell, W. (2013). Comparing Data Center Power Distribution Architectures. White Paper 129. Schneider Electric, 2013.
- [5] Belden, “Belden,” [Online]. Available: <http://www.belden.com/blog/data-centers/To-HAC-or-CAC-That-is-the-Question.cfm>.
- [6] “PHYS.ORG,” [Online]. Available: <http://phys.org/news/2012-09-intel-math-oil-dunk-cooler-servers.html>.
- [7] Van der Ha, B.N.B. (2014). Deliverable D7.1 Data Centres: AMrket Archetypes and Case Studies.
- [8] Poggi, N., Moreno, T., Berral, J., Gavalda, R., and Torres, J. (2009). Self-adaptive utility-based web session management. *Comput. Networ.*, 53, 1712–1721.

- [9] Froese, K.W., and Bunt, R.B. (1996). The effect of client caching on file server workloads, in *29th Hawaii International Conference System Sciences*.
- [10] Macías, M., and Guitart, J. (2014). SLA negotiation and enforcement policies for revenue maximization and client classification in cloud providers. *Future Gener. Comp. Sy.*, 41, 19–31.
- [11] [Online]. Available: <http://www.cs.huji.ac.il/labs/parallel/workload/>.
- [12] Carbó, A. O. E. S. J. C. M. M. G. J. (2016). Experimental and numerical analysis for potential heat reuse in liquid cooled data centres. *Energy Convers. Manage.*, 112, 135–145.
- [13] [Online]. Available: <http://www.morganclaypool.com/doi/pdfplus/10.2200/S00516ED2V01Y201306CAC024>.
- [14] Delimitrou, C., and Kozyrakis, C. (2013). iBench: Quantifying interference for datacenter applications, in *2013 IEEE International Symposium on Workload Characterization (IISWC)*.
- [15] [Online]. Available: <http://kernel.ubuntu.com/~cking/stress-ng/>.
- [16] [Online]. Available: <https://launchpad.net/sysbench>.
- [17] [Online]. Available: <http://www.mersenne.org/download/>.
- [18] “Pmbw – Parallel Memory Bandwidth benchmark/measurement,” [Online]. Available: <http://panthema.net/2013/pmbw/>.
- [19] “Fio – Flexible I/O tester,” [Online]. Available: <http://git.kernel.dk/?p=fio.git;a=summary>.
- [20] “Iperf3: A TCP, UDP, and SCTP network bandwidth measurement tool,” [Online]. Available: <https://github.com/esnet/iperf>.
- [21] [Online]. Available: <http://parsa.epfl.ch/cloudsuite/cloudsuite.html>.
- [22] [Online]. Available: <http://www.nas.nasa.gov/publications/npb.html>.
- [23] Oró, E. (2014). Deliverable D4.2 Energy requirements for IT equipment, 2014.
- [24] “Tier Classifications Define Site Infrastructure Performance,” Uptime Institute.
- [25] Shrestha, N. e. a. (2015). Deliverable D4.5 Catalogue of advanced technical concepts for Net Zero Energy Data Centres, 2015.
- [26] Donoghue, A. (2015). Energizing renewable-powered datacenters. 451 Research, London.
- [27] Oró, E. D. V. G. A. S. J. (2015). Energy efficiency and renewable energy integration in data centres. Strategies and modelling review. *Ren. Sustain. Energy Rev.* 42, 429–445.
- [28] “Thermal guidelines for Data Processing environments,” ASHRAE Datacom Series 1 – Third Edition.

- [29] ASHRAE, “Gaseous and particulate contamination guidelines for data centres, Technical Committee (TC) 9.9,” ASHRAE, 2009.
- [30] ASHRAE, “Gaseous and particle contamination guidelines for data centers.,” 2011.
- [31] ASHRAE, “Save energy now presentation series,” 2009.
- [32] Hintemann, R., Fichter, K., and Stobbe, L. (2010). Materialbestand der Rechenzentren in Deutschland. Umweltbundesamt, Dessau-Roßlau.
- [33] Song, S., Ge, R., Feng, X., and Cameron, K. (2009). Energy profiling and analysis of the HPC challenge benchmarks. *International Journal of High Performance Computing Applications*, Vol. 23, pp. 265–276.
- [34] Ren, Z., Xu, X., Wan, J., Shi, W., and Zhou, M. (2012). Workload characterization on a production Hadoop cluster: A case study on Taobao. in *IEEE International Symposium on Workload Characterization*.
- [35] P. V. o. S. S. D. u. IBM, “http://www-03.ibm.com/systems/resources/ssd_ibmi.pdf,” 2009. [Online].
- [36] Gmach, D., Rolia, J., Cherkasova, L., and Kemper, A. (2007). Workload analysis and demand prediction of enterprise data center applications, in *IEEE 10th International Symposium on Workload Characterization*, Washington, DC, USA.
- [37] Chen, Y., Alspaugh, S., and Katz, R. (2012). Interactive analytical processing in big data systems: A cross-industry study of MapReduce workloads. *PVLDB*, 5, 1802–1813.
- [38] Ren, K., Kwon, Y., Balazinska, M. and Howe, B. (2013). Hadoop’s adolescence. *PVLDB*, 6, 853–864.
- [39] 16 December 2013. [Online]. Available: <http://parsa.epfl.ch/cloudsuite>.
- [40] University of Hertfordshire, “TIA 942. Data Centre Standards Overview,” 2014. [Online]. Available: http://www.herts.ac.uk/_data/assets/pdf_file/0017/45350/data-centre-standards.pdf.