90V231.01 - Unrestricted

Report

Data Centres Infrastructure Energy Efficiency

State of the art

Author(s)
Igor Sartori

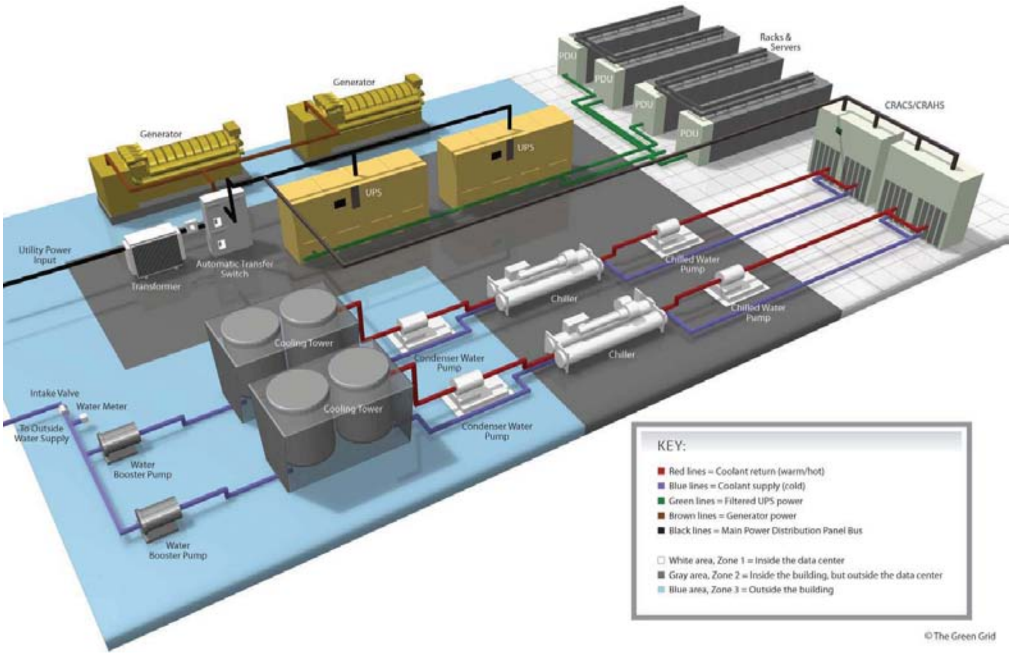

© The Green Grid

SINTEF Building and Infrastructure
Buildings
2010-12-17

SINTEF

SINTEF Byggforsk
SINTEF Building and Infrastructure
Address:
Postboks 124 Blindern
NO-0314 Oslo
NORWAY
Telephone:+47 73593000
Telefax:+47 22699438
byggforsk@sintef.no
http://www.sintef.no/Byggforsk/
Enterprise /VAT No:
NO 948007029 MVA

Report

Data Centres Infrastructure Energy Efficiency

State of the art

KEYWORDS:
Data center infrastructure, energy efficiency

VERSION
1.1

DATE
2010-12-17

AUTHOR(S)
Igor Sartori
Evt flere forfattere

CLIENT(S)
Telenor Real Estate

CLIENTS REF.
Johanne Melø

PROJECT NO.
90V231.01

NUMBER OF PAGES/APPENDICES:
36+ Appendices

ABSTRACT

The state of the art report focuses on the Heating, Ventilation and Air Conditioning (HVAC) systems energy efficiency. For planning and operational purposes the data centre infrastructure is seen in a physical perspective and subdivided in four levels: 1) Inside racks and cabinets, 2) In the data centre room(s), 3) Inside the building and 4) Outside the building. At each level the technological solutions are analysed in consideration of: Components & Systems, Controls & Strategies, Emerging technologies. An important observation is that if Telenor plans to upgrade a data centre capacity to an overall heat load density of 3 $\mathrm{kW/m^2_{floor\ area}}$, from today's reference value of 2 $\mathrm{kW/m^2_{floor\ area}}$, and when this would imply a heat density at rack level > 5 kW/rack, it is advisable to consider adopting liquid cooling systems. For measuring and reporting purposes the data centre is seen in a functional perspective and subdivided into: IT services and infrastructure services. The ratio between energy use for IT services and infrastructure services gives the Power Usage Effectiveness (PUE), which is the internationally recommended metrics to measure and report overall data centre infrastructure efficiency. How to measure and report the PUE is discussed. The most important consideration for data centre infrastructure energy efficiency are summarised in a checklist.

PROJECT MANAGER
Tore Wigenstad

SIGNATURE

CHECKED BY
Tore Wigenstad
Mark A. Murphy

SIGNATURE
Mark Murphy

APPROVED BY (name, position)
Kim Robert Lisø, Research Director "Energy and Architecture"
KRISTIN HOLTHE, for

SIGNATURE
Kristin Holthe

REPORT NO.
90V231.01

ISBN

CLASSIFICATION
Unrestricted

CLASSIFICATION THIS PAGE
Unrestricted

Document history

VERSION	DATE	VERSION DESCRIPTION
0.1	2010-12-13	First draft
0.2	2010-12-16	Second draft
0.3	2011-01-03	Completed version
1.0	2011-01-05	Completed version, language corrected
1.1	2011-01-06	Final version, quality assured

Table of contents

APPENDICES

Data Centres Infrastructure Energy Efficiency

State of the art

1 Introduction

This state of the art report is performed as part of the SINTEF – Telenor collaboration agreement and specifically in the area 2: Power efficiency, activity: Data centre infrastructure. The state of the art focuses mainly on the Heating, Ventilation and Air Conditioning (HVAC) energy efficiency, and looks at the data centre infrastructure from a physical perspective.

1.1 HVAC energy efficiency

Data centres typically host equipment that provide IT services (also named data equipment) and telecommunication services; altogether such equipment is usually referred to as IT equipment (or datacom equipment). However, operating IT equipment is not the only energy demanding activity in a data centre. Energy is also required for auxiliary electrical services, such as switchgears, Uninterruptible Power Supply (UPS), Power Distribution Units (PDU), etc., and for keeping the indoor environment within specified conditions by means HVAC equipment. In particular, in data centres the most significant HVAC energy demand is associated with cooling (including de/humidification) and ventilation. Energy demand for auxiliary electrical services and for HVAC equipment together is referred to as infrastructure energy demand (as opposed to the IT energy demand).

In past years, as the demand for more computing capacity increased rapidly, the primary issue for data centres has been space availability. Now, more and more data centres run out of power before they run out of space and cooling requirements for dense servers are driving power demand [1]. Recent surveys report that infrastructure energy is about 50% of the total demand throughout the period from 2000 to 2006 [2], while average HVAC energy demand was about 35% in 12 data centres monitored in 2007 [3] as shown in Figure 1.

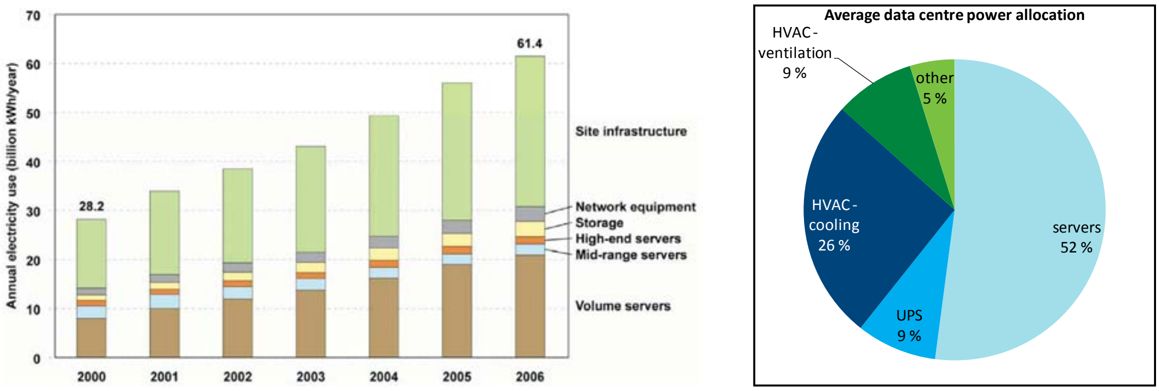

Figure 1. Data centre energy allocation by end use, source left) [2] and right) [3].

This state of the art report focuses on the HVAC energy demand and related efficiency measures because of their importance in the overall efficiency of a data centre.

1.2 Perspectives: Functional and Physical

A data centre hosts a variety of technologies that need to interact, often in complex configurations. In general, it is possible to distinguish between the main different functions performed so that the equipment can be grouped accordingly, i.e. IT-, auxiliary- and HVAC-services. This means adopting a functional perspective over the data centre. However, while the space taken by IT-services may be identified with the sum of the racks space, infrastructure services are not so well physically separated as they are functionally; there is not such a thing as an HVAC room or an auxiliaries room that contain all the HVAC or auxiliary equipment. Rather, such equipment is distributed in all of the infrastructure spaces, from inside the IT-racks to outside of the building as shown in Figure 2.

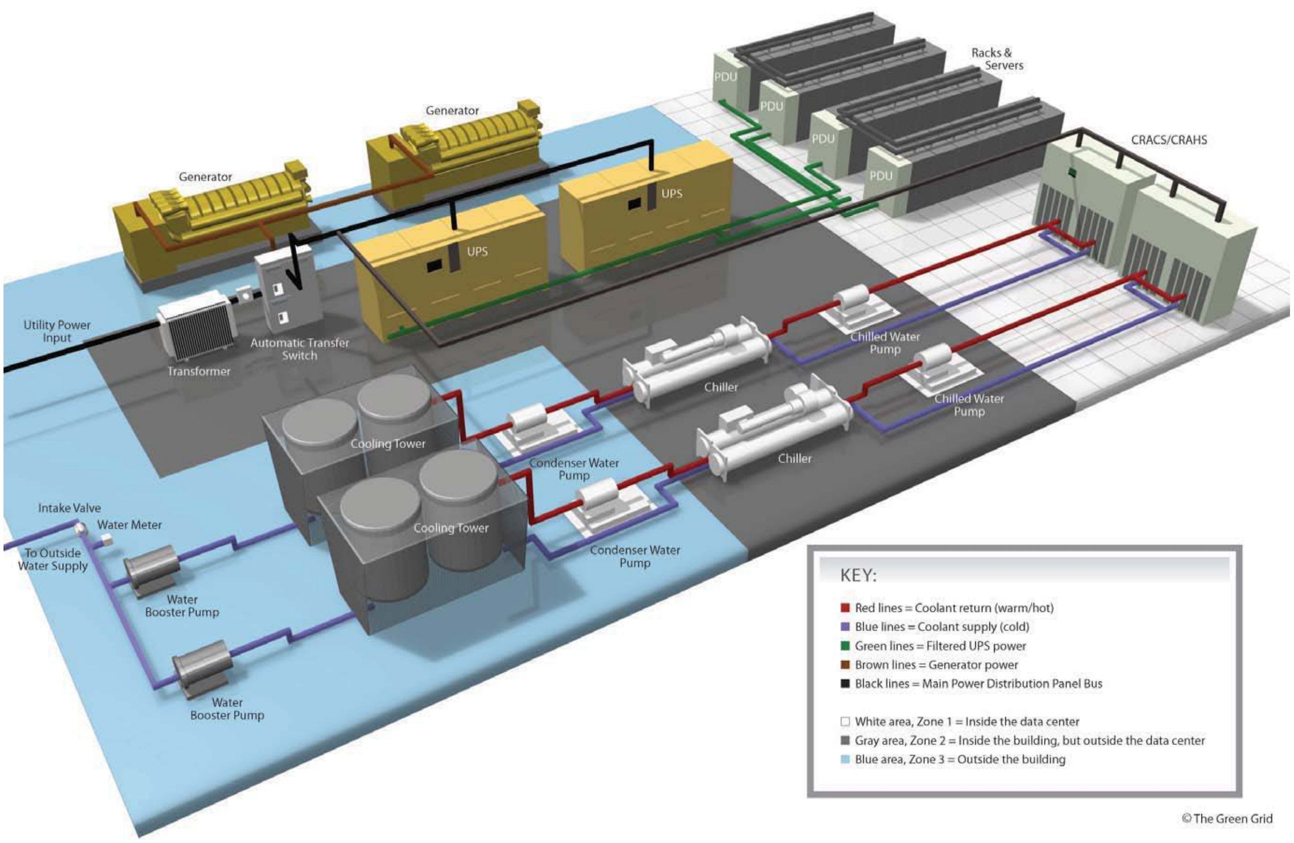

© The Green Grid

Figure 2. Main zones partition in a physical perspective of a data centre.

Looking at how the equipment is distributed in the space means to adopt a physical perspective over the data centre. For the purposes of this state of the art report the attention is put on the HVAC equipment and not on the electrical auxiliary equipment.

1.2.1 Functional perspective

From a functional point of view the data centre activities are often grouped into two main categories:

- IT services, that comprises both the computing and communication tasks. This is considered the net work performed by a data centre, as well as its core business.
- Infrastructure services, which support the IT services providing the necessary conditions for IT equipment operation. This includes:
    - electrical auxiliary services, such as transformers, switchgears, UPS, PDU, cabling and emergency generators
    - HVAC services, such as ventilation, cooling and de/humidification of IT equipment and rooms, including heat rejection to the outside and all intermediate equipment.

The functional perspective is often reflected also in the data centre organisation, where IT and infrastructure are two separate departments. The energy bills are often competence of the infrastructure departments, but energy efficiency is an issue to tackle holistically and that’s why the physical perspective is preferred in this report. Furthermore, non coordination between the two departments can be part of the causes of some energy inefficiency and/or an obstacle to implement energy saving measures.

The functional perspective is suitable for measuring and reporting purposes, hence §3 of this report adopts a functional perspective, with the separation between IT and infrastructure services.

1.2.2 Physical perspective

From a physical point of view, IT services are normally confined within the racks of a data centre while infrastructure services are spread at all levels, from the electronics within the rack (e.g. PDU and cooling fans) to the outside of the building (e.g. emergency generators and heat rejection devices). Hence, a holistic approach to energy efficiency needs to take into consideration all physical levels, including the IT equipment itself.

The physical perspective is suitable for planning and operation purposes, hence §2 of this report adopts a physical perspective, with the separation into the following physical levels:

- Inside racks and cabinets
- In the data centre room(s)
- Inside the building
- Outside the building

At each level, the investigated technologies and solutions are analysed with consideration of:

- Components & Systems
- Controls & Strategies
- Emerging technologies

This state of the art report focuses on the HVAC services, while non-HVAC (i.e. either IT or auxiliary) energy efficiency measures are simply mentioned when relevant.

2 Planning and operation

2.1 Inside racks and cabinets

2.1.1 Components & Systems

Heat load

The heat load density of a data centre depends on the heat release by the IT equipment. The first consideration is then about the historic and expected trend for heat load of IT equipment. This is normally expressed in W/m^2 of footprint area. To avoid any misunderstanding on the meaning of such data it is first indispensable to specify exactly what these measures refer to. When it comes to heat load [W] this refers to the typical heat rejection in full configuration and not to the nameplate value (see discussion below about proper sizing and Figure 7 and Table 1). For the footprint area, it is important not to confuse it with the total area of the conditioned space, which includes also the space for hot/cold aisles and other spare space in a datacom room. The footprint area is defined as the width * length dimension of the packaging; for equipment mounted on a 2-post rack – typically telecommunication equipment – the width includes the two posts either side of the packaging, see Figure 3. Another way of giving heat density data is in W/rack; here it is normally assumed that a rack occupies a space of approximately 0.6 m^2 (50 * 120 cm).

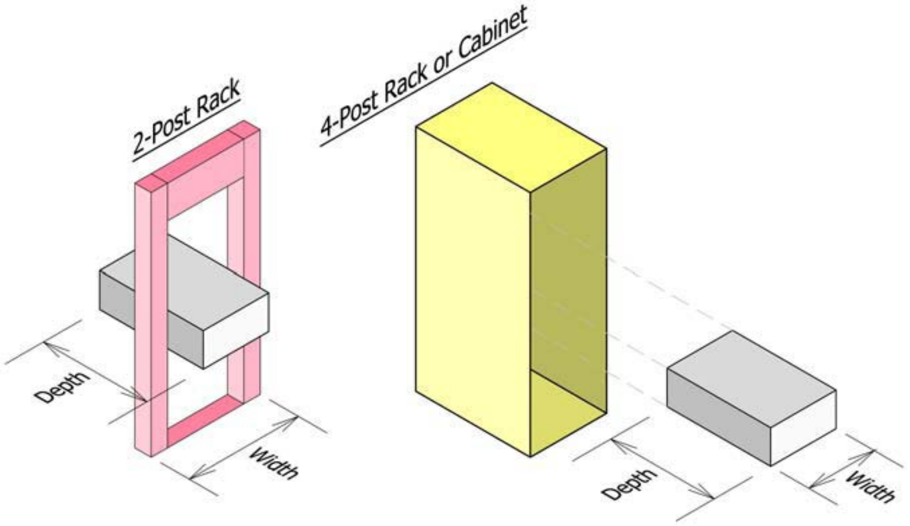

Figure 3. Graphical representation of width × depth measurements used for equipment footprint definitions, source [4].

The heat load density of IT equipment has increased dramatically in the last twenty years, and it is expected to continue growing also in the coming five years. The increase is due to the increasing computing capacity and reduced dimensions of electronic chips and boards; this in turn results in denser packaging of components within a rack. Figure 4 shows the historic and projected heat load density in W/m^2 footprint, split for different type of equipment; note that the y-axis is in logarithmic scale. Figure 5*a)* shows the same data but with a linear scale, while Figure 5*b)* presents the values in kW/rack.

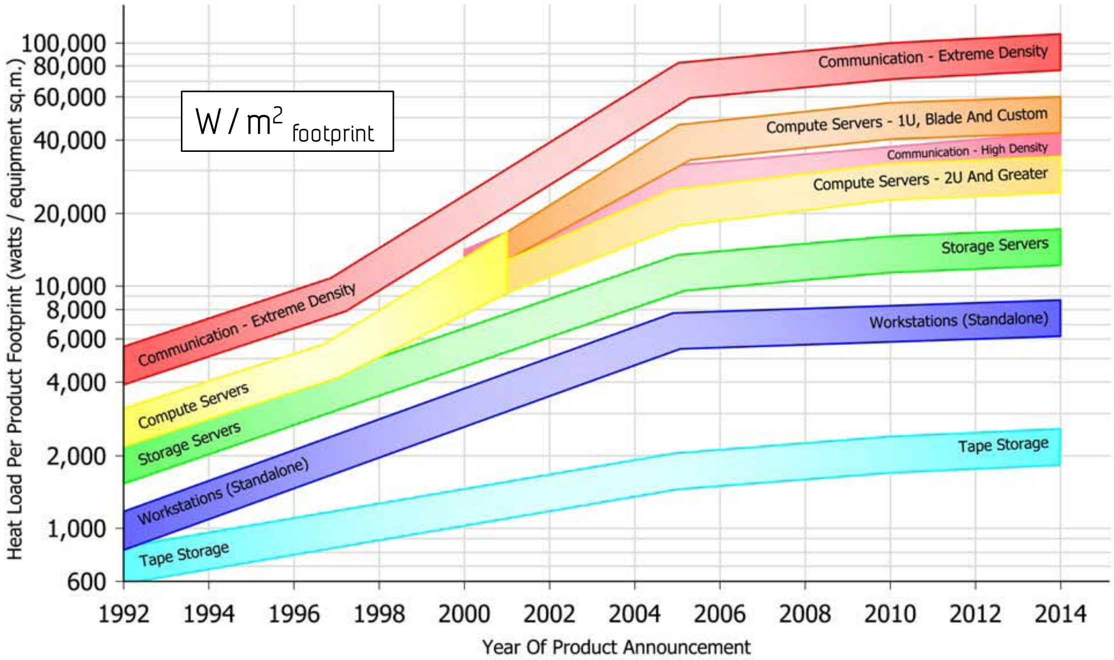

Figure 4. ASHRAE power trend chart in logarithmic scale, source [4].

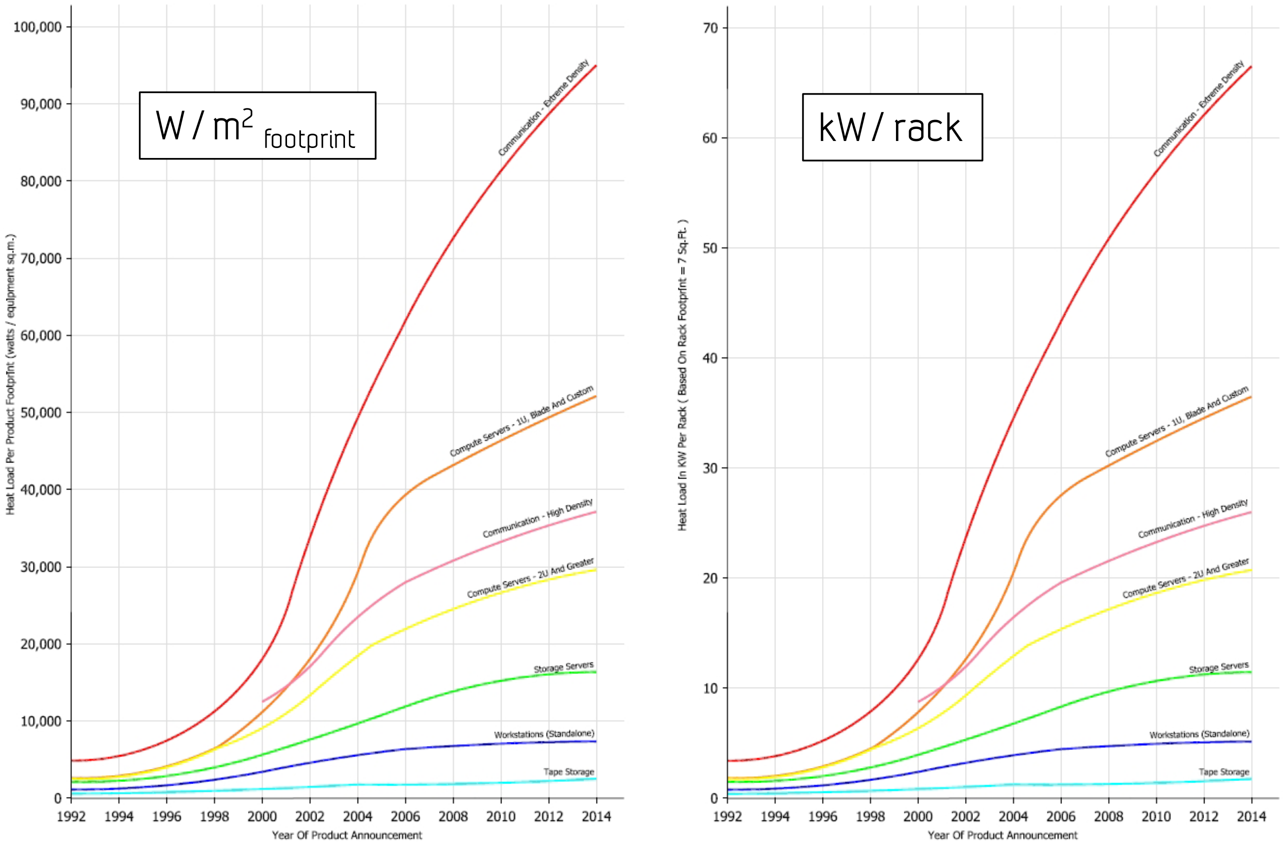

Figure 5. ASHRAE power trends in linear scale: *a)* per square metre of footprint, *b)* per rack, source [4].

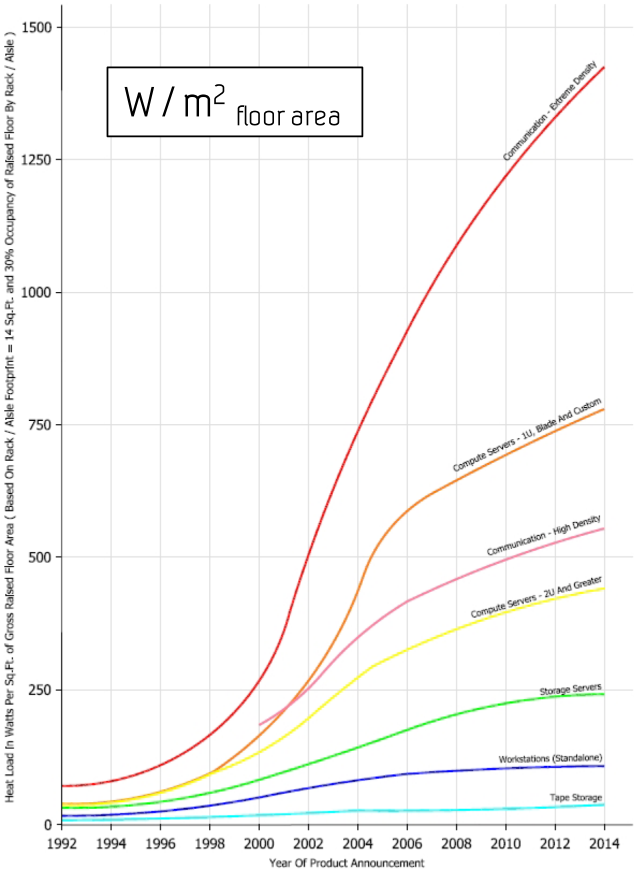

Figure 6. ASHRAE power trend per square foot of gross floor area, source [4]. 1 m^2 = ca. 10 ft^2.

The density per m^2 footprint may be converted into density per m^2 of floor area when it is known the ratio of aisle floor space/rack and the occupancy of rack-aisle over the total floor are of the room. The example in Figure 6 shows the case for an aisle/rack ratio of 2 (1 rack = 7 ft^2 with 1 rack every 14 ft^2 of aisle area) and a rack-aisle occupancy of 30% – hence not a pure server room but a room for mixed use. While pure server rooms have occupancy probably higher than 30%, the aisle/rack value of 2 can be considered generally valid. In other words, it means the floor space area is double the rack footprint area. Recalling the number of 0.6 $m^2_{\text{footprint}}$/rack, we have then 1.2 $m^2_{\text{floor area}}$/rack.

The reference value for today's situation in most Telenor's data centres is an overall heat density of 2 kW/$m^2_{\text{floor area}}$, which also according to [9] is a typical value for data rooms with miscellaneous equipment, i.e. a mix of old and new compute equipment, i.e. a mix of old and new compute servers, storage servers and communication equipment. Assuming rack-aisle occupancy of 75% in the data centre room we have that:

- 2 [kW/$m^2_{\text{floor area}}$] =
  = x [kW/rack] / 1.2 [$m^2_{\text{floor area}}$/rack] * 75%
- hence x = 3.2 [kW/rack]

If the above hypotheses are likely true – aisle/rack ratio of 2 and aisle-rack occupancy of 75% – it means that Telenor's data centre per today are equipped with racks having an average heat load density of ~ 3 kW/rack. As discussed later in §2.2.1.1 about air cooling, this is a heat load that can normally be managed by an air cooling system in a server room arranged in hot/cold aisles and with a raised floor for cold air supply, see Figure 14*a)*. However, **if Telenor wants to upgrade its capacity to a level of 3 kW/$m^2_{\text{floor area}}$** that means, in the above hypotheses, an average of ~ 5kW/rack. A value of 5 kW/rack represents the upper limit for an air cooling system – unless special solutions are adopted in order to achieve fully controlled airflow paths, as shown in Figure 14*b)* – and this already considering the adoption of baffles or overhead ducts to minimise air mixing between hot and cold aisles. Therefore, **it is advisable to take into consideration adopting liquid cooling systems** discussed later in §2.2.1.2. However, it must be reminded that while these numbers are given in average the reality of a data centre is more complex and each solutions should be carefully designed for each specific case.

Proper sizing

For a proper sizing of power and cooling loads it should be considered that the lifetime of IT equipment, 3-5 years, is short compared to the lifetime of the data centre infrastructure, 15-30 years, and it is therefore natural to expect more cycles of equipment turnover. Furthermore, some extra space and hence even extra cooling capacity is normally supposed to give room for adding more equipment (in facts, even though new equipment is more efficient in terms of computing capacity per footprint area, the demand for data centre services has shown to increase even more rapidly).

Calculating with reasonable accuracy the heat load of an IT system is not an easy task because it depends on the actual way the system is controlled and what are its normal workload and utilisation levels. However, it is import to remind that nameplate values are not suitable indicators of equipment heat release. As warned in [5] “Nameplate ratings should at no time be used as a measure of equipment heat release. The purpose of nameplate rating is solely to indicate the maximum power draw for safety and regulatory approval.” Tests reported in [6] show that also simple methods of de-rating, such as the quite common “nameplate divided by 2”, could result in incorrect power and cooling load estimates. The study also indicates other, preferable ways of estimating the heat load, in order of increasing preference (if available): from the equipment thermal report, for the vendor’s estimation tools, from real time measurements, see Table 1.

Table 1. Proper sizing of power and cooling loads, source [6].

Never use	Use for proper sizing
• Nameplate values • Simple de-rating method also incorrect	• Equipment thermal report • Vendor estimation tools • Real-time measurements

The heat load densities shown in Figure 4 and Figure 5 are obtained by several manufacturers emulating the most probable level of power consumption assuming a fully configured system. Data for heat release for full configuration are found in the equipment’s thermal report; a fac-simile is shown in Figure 7.

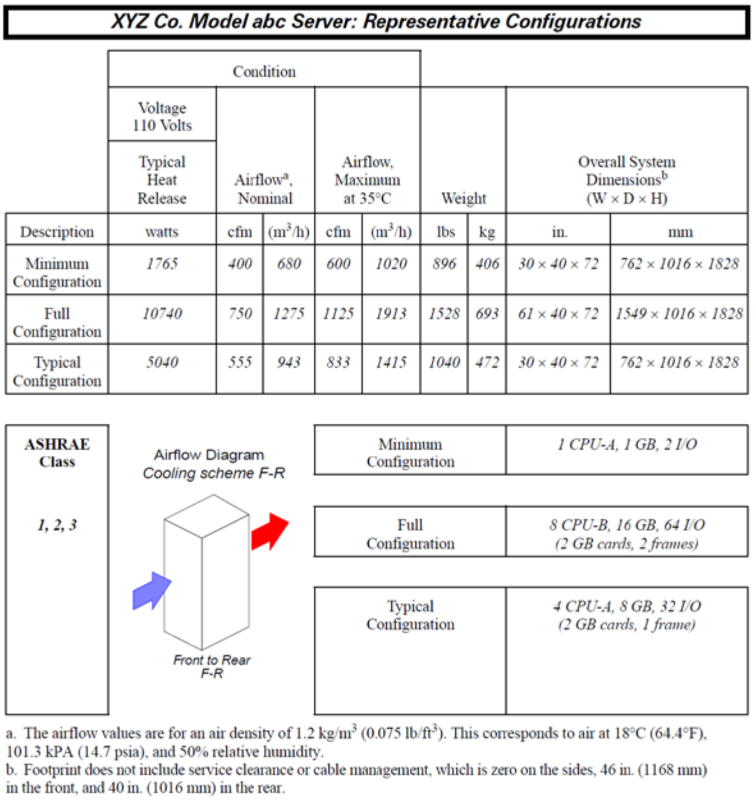

XYZ Co. Model abc Server: Representative Configurations

	Condition								
	Voltage 110 Volts								
	Typical Heat Release	Airflow[a], Nominal		Airflow, Maximum at 35°C		Weight		Overall System Dimensions[b] (W × D × H)	
Description	watts	cfm	(m^3/h)	cfm	(m^3/h)	lbs	kg	in.	mm
Minimum Configuration	*1765*	*400*	*680*	*600*	*1020*	*896*	*406*	*30 × 40 × 72*	*762 × 1016 × 1828*
Full Configuration	*10740*	*750*	*1275*	*1125*	*1913*	*1528*	*693*	*61 × 40 × 72*	*1549 × 1016 × 1828*
Typical Configuration	*5040*	*555*	*943*	*833*	*1415*	*1040*	*472*	*30 × 40 × 72*	*762 × 1016 × 1828*

ASHRAE Class

1, 2, 3

Minimum Configuration	*1 CPU-A, 1 GB, 2 I/O*
Full Configuration	*8 CPU-B, 16 GB, 64 I/O* *(2 GB cards, 2 frames)*
Typical Configuration	*4 CPU-A, 8 GB, 32 I/O* *(2 GB cards, 1 frame)*

a. The airflow values are for an air density of 1.2 kg/m^3 (0.075 lb/ft^3). This corresponds to air at 18°C (64.4°F), 101.3 kPA (14.7 psia), and 50% relative humidity.

b. Footprint does not include service clearance or cable management, which is zero on the sides, 46 in. (1168 mm) in the front, and 40 in. (1016 mm) in the rear.

Figure 7. Example of thermal report, source [5].

Air and liquid cooling

Datacom racks or cabinets are traditionally air-cooled; conditioned air enters the front of the rack, moved either by centralised or local fans, passes through the rack cooling the electronics and the warmed air is finally expelled from the rear of the rack. High density racks, though, may require liquid cooling, as shown in Figure 8. Liquid cooling of racks and cabinets is defined in [7] as follows:

> “Conditioned liquid is supplied to the inlets of the rack/cabinet/server for thermal cooling of the heat rejected by the components of the electronic equipment within the rack. It is understood that

> within the rack, the transport of heat from the actual source component (e.g. CPU) within the rack itself can be either liquid or air based (or any other heat transfer mechanism), but the heat rejection media to the building cooling device outside the rack is liquid.”

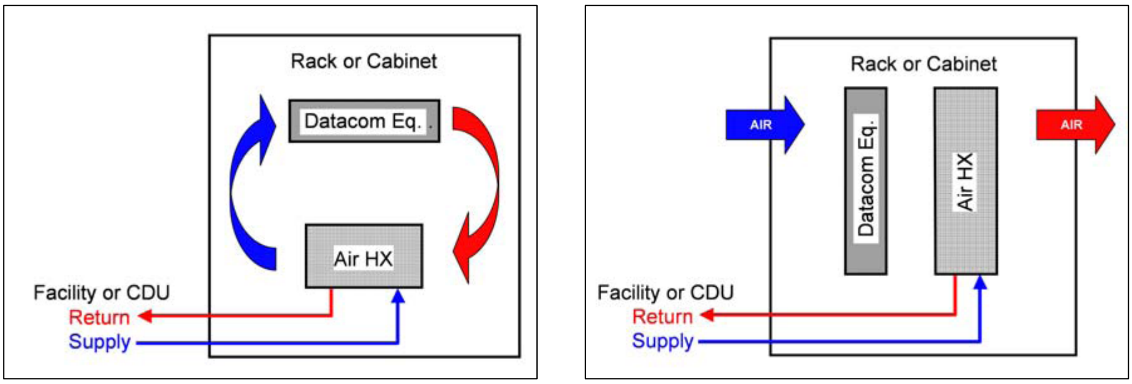

Closed architecture

Open architecture

Figure 8. Liquid cooling of racks, closed and open architecture, source [8].

Both closed and open architectures are possible. In a closed architecture the heat is entirely removed from the rack by the cooling liquid, while in an open architecture the room air still passes through the rack but a cooling coil positioned at the rear of the rack extracts almost all of heat from the room air before it returns to the air conditioning unit.

Liquid cooling is advisable for racks with heat density higher than 5 kW, unless fully controlled air path solutions are in place, and is absolutely necessary for racks with heat density higher than 15 kW [9].

2.1.2 Controls & Strategies

Recommended environmental envelope

The most important concept when it comes to energy efficient control of environment for datacom equipment is summarised in the sentence “colder does not mean better”. For various reasons the temperature in data centres is often kept too low ([10], [11] and [14]), including:

- Misinterpretation of required operating conditions: temperature and humidity ranges refer to inlet air conditions, not to space conditions or return air conditions.
- The room is kept cold to achieve longer ride-through time during a cooling outage. However, while increasing the set point by a few degrees can have a large effect on energy demand; it is unlikely that a few degrees lower set point will offer any significant improvement in ride-through time.
- There is a concern that increasing the IT inlet air temperatures might have a significant effect on reliability; such a concern is not well founded. An increase in inlet temperature does not necessarily mean an increase in component temperatures, as shown in Figure 9. Using variable speed fans allows keeping the component at a constant temperature – within vendor’s temperature specifications – for a wide range of inlet air temperatures.

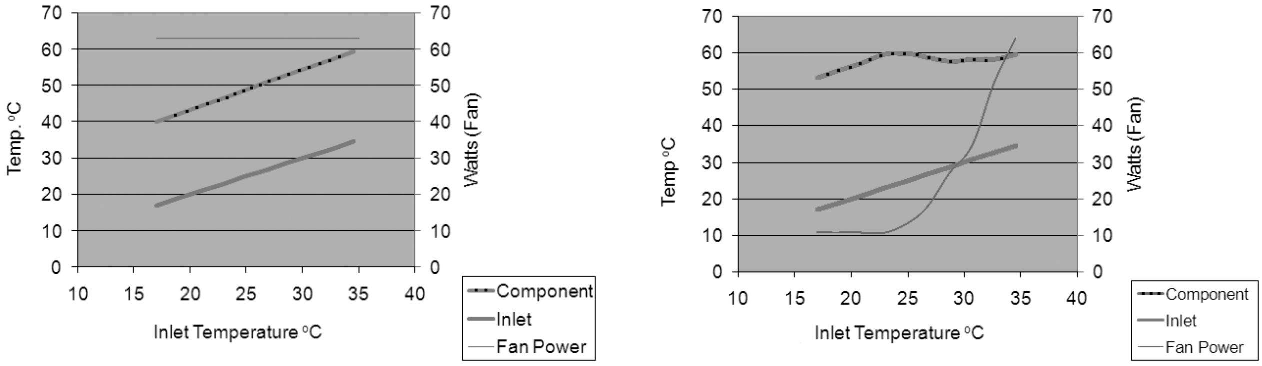

Figure 9. Inlet and component temp. with *left)* fixed fan speed, *right)* variable fan speed, source [10].

To provide greater flexibility in facility operations, particularly with the goal of reduced energy consumption in data centers, ASHRAE Technical Committee (TC) 9.9 has undergone an effort to revisit these recommended equipment environmental specifications. Reviewing the available data from a number of IT manufacturers, the 2008 expanded recommended environmental envelope is the agreed-upon envelope that is acceptable to all IT manufacturers, and operation within this envelope will not compromise overall reliability of IT equipment [10]. A comparison between previous (2004) and 2008 recommended environmental envelope is shown in Table 2 and Figure 10. The 2004 values were originally included also in the (informative) recommendations in the norm TIA-942 [11], while the 2008 values have been included with the second amendment, TIA 942-2 [12].

Table 2. Comparison of 2004 and 2008 recommended environmental envelope data, source [10].

	2004	2008
Low-End Temperature	20°C (68°F)	18°C (64.4°F)
High-End Temperature	25°C (77°F)	27°C (80.6°F)
Low-End Moisture	40% RH	5.5°C DP (41.9°F)
High-End Moisture	55% RH	69% RH and 15°C (59°F DP)

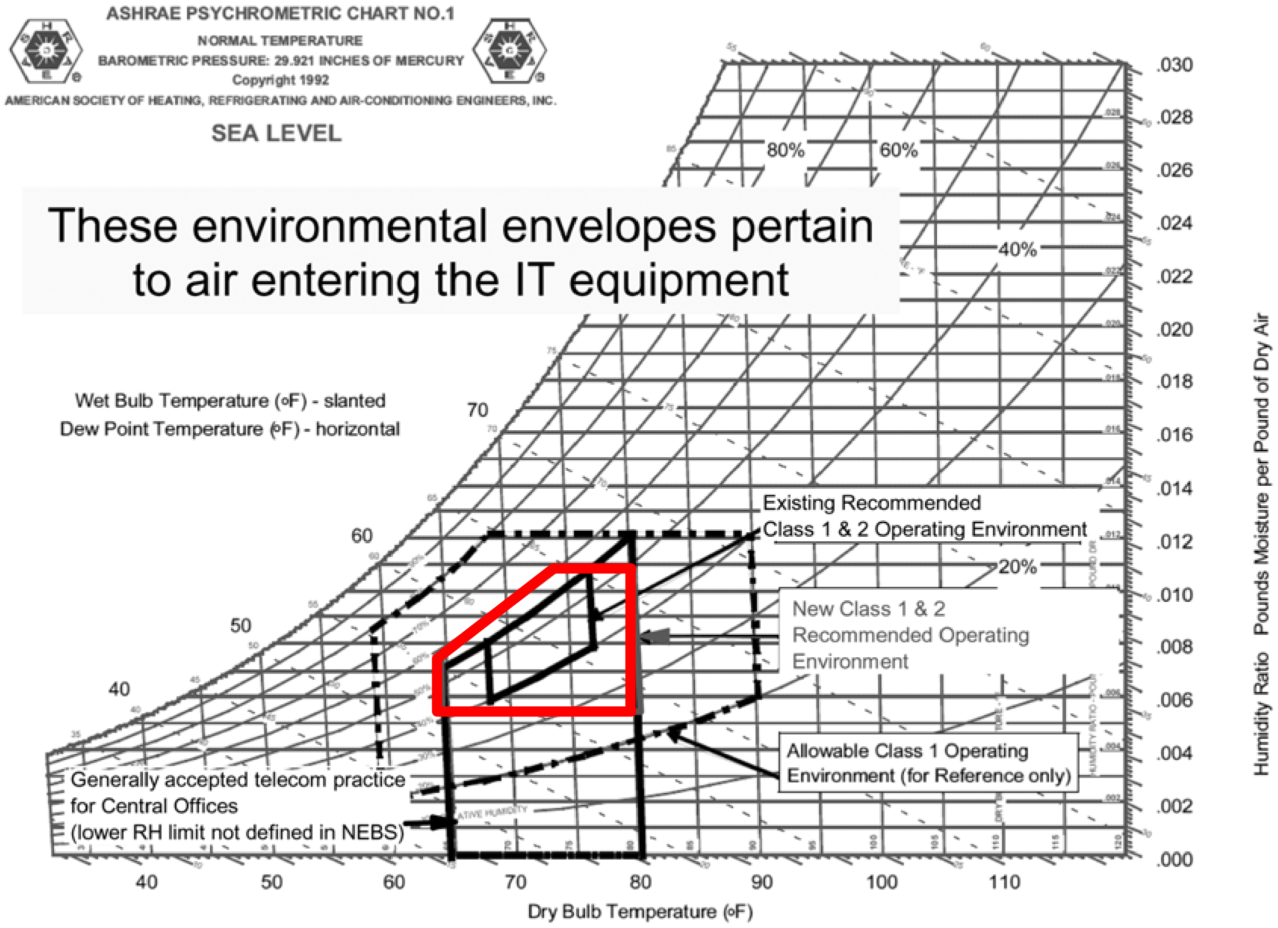

Figure 10. ASHRAE 2008 recommended environmental envelope, new values marked in red, source [10].

The major advantages of expanding the environmental envelope are:

- Increased temperature and humidity operating deadband (range of values where no action is required by the system)
- Longer use of the economizers
- Higher thermodynamic efficiency due to smaller temperature differential in the cooling cycle

2.1.2.1 Other than HVAC

Energy efficiency controls and strategies at rack level that can be considered for non-HVAC related operations are:

- Virtualisation of the servers
- Identification and managing of unused servers

2.1.3 Emerging technology

Datacom equipment liquid cooling

An emerging technology with considerable potential for energy efficiency is the liquid cooling of datacom equipment. Differently from removal of heat from the rack by liquid cooling as analysed before, in this case it is the transport of heat from the actual source component (e.g. CPU) within the rack itself that is via liquid. This not a feature that can be added on to a rack by the data centre designer/managers; datacom equipment has to be manufactured with built-in liquid cooling. Such components are appearing in the market, especially for high density IT and communication equipment. Integration with the rack cooling solution can, also in this case, be with a closed or an open architecture as shown in Figure 11.

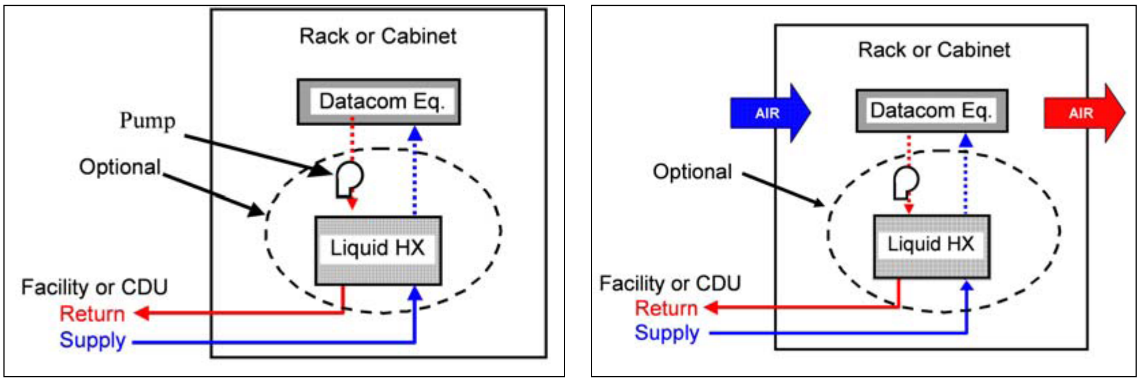

Closed architecture

Open architecture

Figure 11. Liquid cooling of datacom equipment, closed and open architecture, source [8].

2.2 In the data centre room(s)

2.2.1 Components & Systems

2.2.1.1 Air cooling

Air cooling systems are suitable for data centres with low (3-5 kW/rack) to middle (up to 15 kW/rack) heat load densities, depending on the air distribution configuration. Smaller facilities typically use CRAC units while larger facilities often use a centralised chilled water system with CRAH units in the computer rooms.

CRAC and CRAH units

Data centre rooms can be conditioned either by packaged Computer Room Air Conditioning units (CRACs) located directly in the data centre rooms, see Figure 12, or by a central station – often located in a separate technical room – connected to Computer Room Air Handling units (CRAHs) in the data centre rooms via a chilled water loop, see Figure 13.

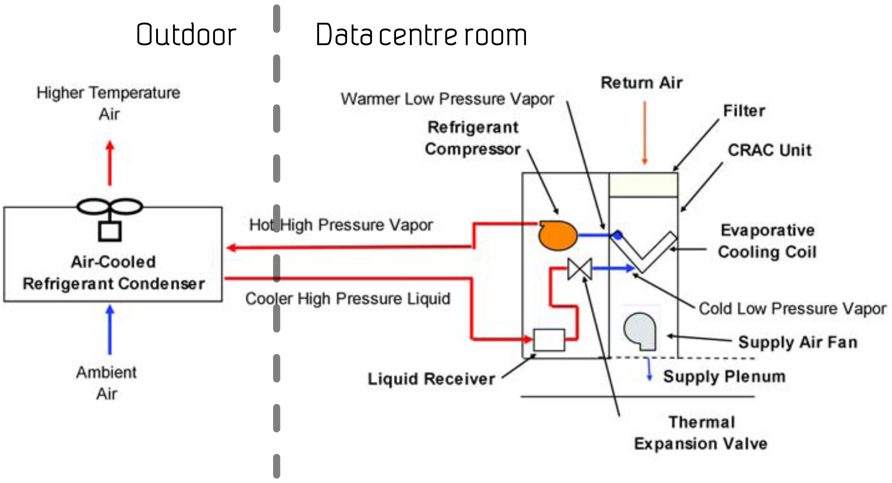

Figure 12. Schematic of HVAC system based on CRAC units, source [14].

CRACs are the most popular datacom cooling solution and are especially indicated for smaller data centres due to their simpler system architecture. CRAC units use an evaporative cooling coil with direct expansion of the refrigerant (Dx coils) to cool the computer room air. A compressor moves the refrigerant to the condenser – located on the outside of the building – for heat rejection, often simply by air-cooling with outdoor air and circulation fans. Hence, in CRAC systems there is a direct connection between the computer room and the outdoor with no need for separated technical rooms to host further equipment, as shown in Figure 12.

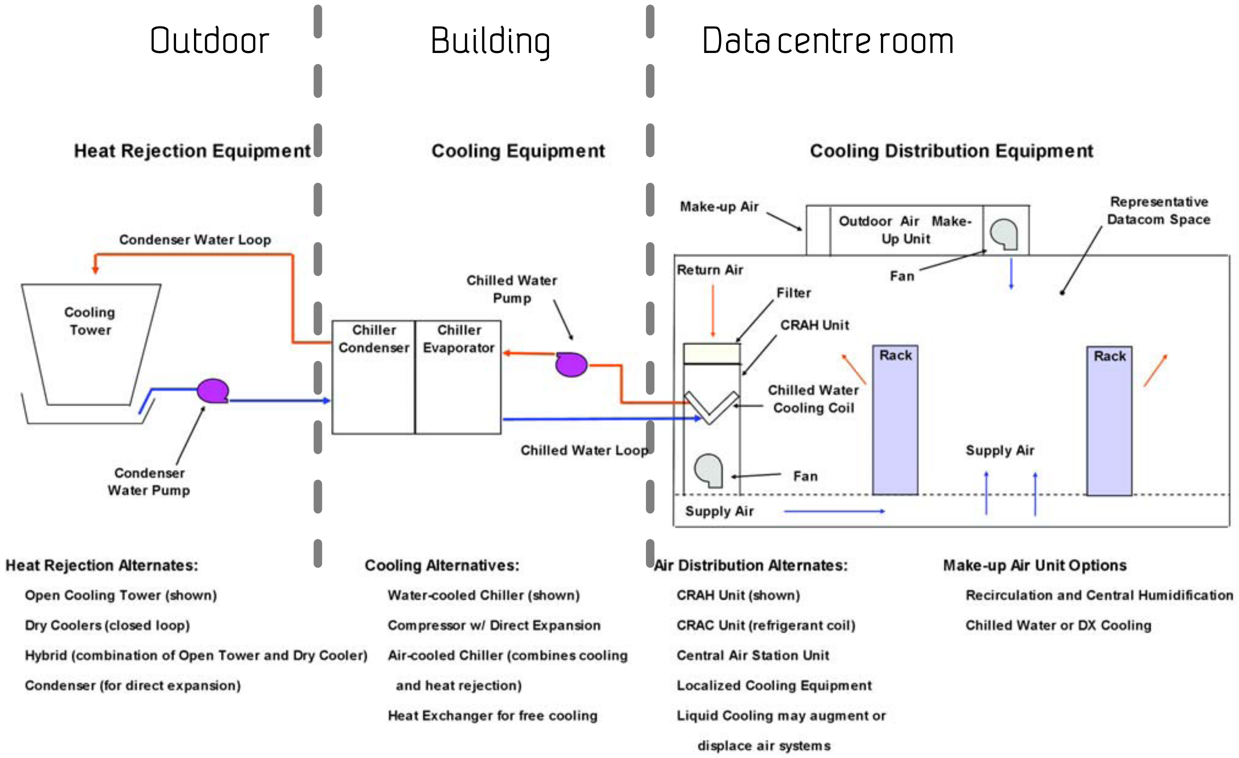

Figure 13. Schematic of HVAC system based on CRAH units and chilled water loop, with list of alternatives, source [14].

Larger datacom facilities often use a centralised refrigerant station to produce chilled water that is then sent via a chilled water loop to the CRAH units in the computer rooms. In this case it is necessary to dispose of additional space within the building where the chiller station is located, as shown in Figure 13.

Centralised air system

A centralized system offers many advantages over the traditional multiple distributed units. Centralized systems use larger motors and fans that tend to be more efficient, especially when operated at part-load. They are also well suited for Variable Air Volume operation, VAV, through the use of Variable Frequency Drives, VFD. A central system allows redundancy to be implemented in a manner that provides maximum reliability (in form of fans spinning reserve) and increases normal operating system efficiency. Implementation of an airside economizer system is simplified with a central air handler system. Reduced footprint, reduced maintenance and maintenance traffic in the data center are additional benefits [16].

Air distribution

Air distribution in data centres is based on the principle of the hot and cold aisles, as shown in Figure 14*a)*; the racks are aligned front-to-front, so that cold air is supplied in the cold aisle to the front of two rack rows and warm air is extracted from the hot aisle from the back of two rack rows. It is important to prevent air mixing between the cold and the hot aisles because that would be detrimental for energy efficiency. To avoid or minimise airflow “short circuits” internal to the racks it is advisable to use blanking panels to fill spaces between equipment. To avoid or minimise short circuits external to the racks it is advisable the use baffles for containment of the hot/cold aisles or to use overhead ducts that guarantee a better control of the airflow path and inlet conditions to the racks. Additionally, a common solution is then to have a raised floor, also named data floor, as plenum to distribute the cold air into the computer room by means of perforated floor tiles. The underfloor space is used also for cabling purposes so that it is important to pay attention that the cables do not create an obstacle to the air movement.

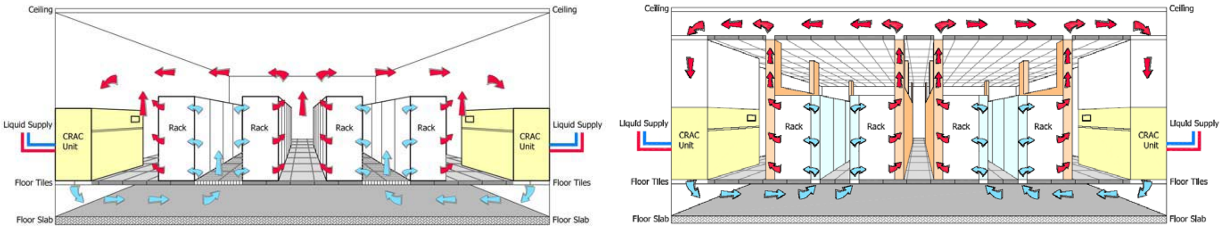

a) *b)*

Figure 14. Airflow distribution with *a)* hot/cold aisles and partially controlled path, *b)* fully controlled airflow path, source [4].

A typical solution with raised floor and uncontrolled return air path is suitable for heat loads up to 3 kW/rack, while adopting overhead ducts and/or aisle separation baffles may increase the capacity up to loads of 5 kW/rack [9]. Higher heat load densities can be covered when adopting a fully controlled air supply and exhaust path, as the one shown in Figure 14*b)*. This solution is effective when the air can effectively flow through the equipment in the rack, hence with racks equipped with their own ventilation fan. The maximum heat load with fully controlled air path is about 15 kW/rack [9].

2.2.1.2 Liquid cooling

As load densities raise air cooling becomes more and more inefficient (due to the large delta T in the thermodynamic cycle and due to air moving device performance) and subject to acoustic limitations. Liquid cooling is necessary when the heat load density increases above 15 kW/rack and is advisable already for loads above 5 kW/rack. The reason for liquid cooling being more effective than air cooling is easily understood by looking at heat transfer properties of liquids versus gases; a summary of indicative typical value ranges for convection heat transfer coefficients is given in Table 3. The heat transfer capability of liquids is two to three orders of magnitude greater than gases. As seen in §2.1.1 and 2.1.3, whether the cooling of electronics within the rack is air or liquid based, the extraction of heat from the rack is via liquid. There can be both closed and open architectures. A liquid cooling system with an open architecture has the advantage to utilize the room air volume as a thermal storage to ride through short power outages.

Table 3. Typical values of the convection heat transfer coefficient.

Process	**h [W/m^{2}K]**
Free convection	
Gases	2 – 25
Liquids	50 – 1000
Forced convection	
Gases	25 – 250
Liquids	100 – 20 000
Convection with phase change	
Boiling or condensation	2 500 – 100 000

When talking about liquid cooling in a data centre there are actually different systems to consider. With reference to Figure 15, the liquid cooling system, or cooling loop, may refer to:

- Datacom equipment cooling system; this was discussed in §2.1.3
- Technology cooling system: it includes cooling of the rack, as discussed in §2.1.1, and the liquid loop for further rejection of heat from the data centre room, as it is discussed in this section.
- The chilled water system and the condenser water system; these will be discussed in §2.3 and 2.4, respectively.

The datacom equipment system – inside the rack – may or may not be liquid based, as seen in §2.1. On the other hand chilled water and condenser water systems – inside and outside the building, respectively – may exist also in large facilities where the data centre rooms are actually cooled by air with CRAH units, as seen in §2.2.1.1. It is therefore the technology cooling system – within a data centre room – what is normally referred to when speaking of a liquid cooling system. The main characteristics of the technology cooling system are discussed in this section.

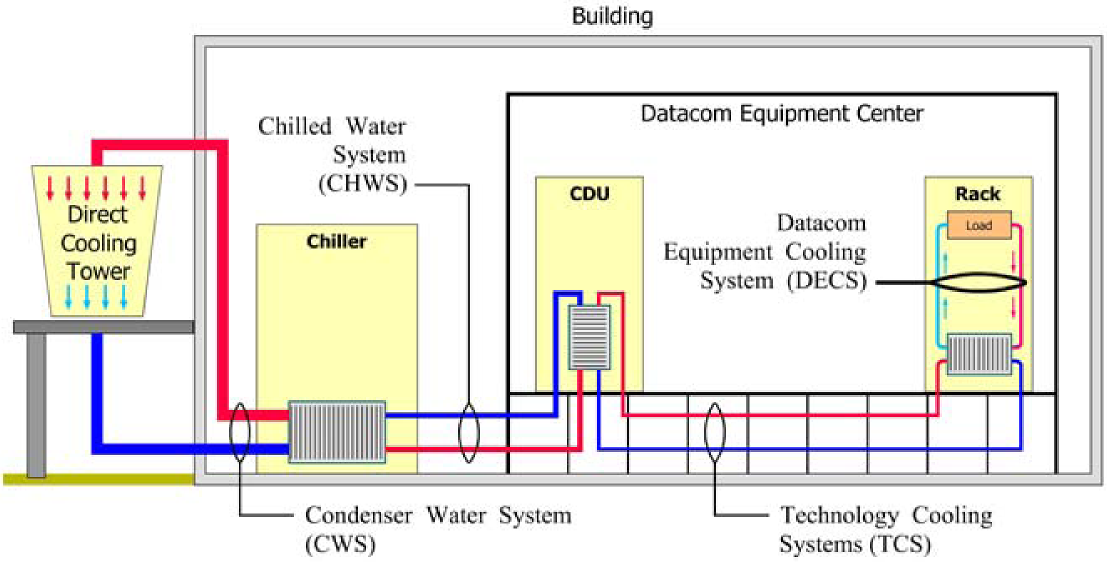

Figure 15. Liquid cooling systems/loops within a data center, source [8].

Coolant Distribution Unit, CDU

The CDU (Coolant Distribution Unit) is to the liquid cooling system what the PDU (Power Distribution Unit) is to the electric system in a data centre. The CDU allows separating the fluid in the data centre room cooling loop from the chilled water loop, and this has several advantages:

- control data centre room fluid temperature and avoid condensation at the cooling coils in the racks
- minimise filtration of building chilled water
- isolation between clusters in case of leak

Piping architecture

Chilled-water distribution systems should be designed to the same standards of quality, reliability, and flexibility as other computer room support systems. Many solutions exist in the design of the piping architecture [8], and they should be evaluated with regards to their characteristics in terms of

- Redundance
- Leakages
- Flexibility

Cooling fluids

There are mainly three types of cooling fluids used in data centres: water, dielectric fluids and refrigerants. The main characteristics of these fluids are summarised in Table 4.

Table 4. Cooling fluids properties, source [7].

Cooling fluid	Advantages	Disadvantages
Water	Low cost	Hazard in case of leakage
Dielectric fluids	Non-conductive --> no hazard in case of leaking	Fluid cost Worse thermal characteristics: larger pumps and piping
Refrigerants	Two-phases: leak as gas, no hazard Better thermal characteristics: smaller pumps and piping	Fluid cost, piping cost Presence of HFCs

2.2.1.3 Other than HVAC

Other non-HVAC components and systems to consider for energy efficiency at data centre room level are:

- Electric cabling (non-obstructive of the under floor air path)
- UPS, Uninterruptable Power Supply
- PDU, Power Distribution Unit

2.2.2 Controls & Strategies

Centralised humidity control

Data centres have small human occupancy and therefore there is little internal generation of humidity. The humidity load comes mainly from the outdoor air and this offers the opportunity to efficiently control humidity with a centralised outdoor air unit. The primary advantage of central humidity control is to eliminate adjacent CRAC units ”fighting”, the simultaneous side-by-side humidification and dehumidification that can occur with multiple CRACs/CRAHs providing humidity control to the same space. This is mainly due to the drift of humidity sensor readings in combination with too narrow humidity control ranges. Another advantage is that if dehumidification (and humidification) is handled centrally, the cooling coils “on the floor”of the datacom facility (e.g. the CDU coils) can run dry. This allows resetting the chilled-water set point to higher temperatures at part-load operation without increased relative humidity.

Part-load operation

HVAC components and systems should be selected and designed for efficient part-load as well as full-load operation, hence with VFD (Variable Frequency Driver) drivers for fans, pumps, compressors. As the load in a data centre is not fixed over time, flexibility is of paramount importance.

2.2.3 Emerging technology

Computational Fluid Dynamics, CFD

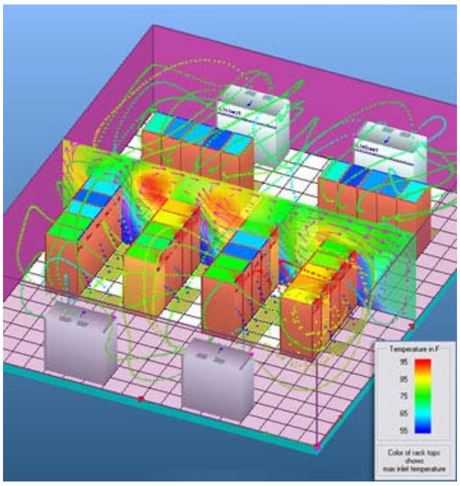

The most prominent advanced technology for air cooled data centres, even though not completely new in the field, is the use of CFD, Computational Fluid Dynamics, computer simulations for data room layout design. It allows to:

- Visualise airflow patterns and temperature distribution
- Identify potential weak areas of cooling capacity, hot-spots, mixing and short-circuits airflows
- Optimise data centre layout design

Figure 16. Graphical representation of temperature distribution and airflow path obtained with CFD.

2.3 Inside the Building

2.3.1 Components & Systems

Chillers

Chillers are typically one of the largest energy users in large data centres and therefore careful attention in the choice and design of the chiller system has big impact both in achieving energy efficiency and in reducing Total Cost of Ownership, TCO. There are various types of chillers, categorised according to their size and working principle. Smaller sizes are generally of the reciprocating type or scroll type; while larger chillers are screw or centrifugal. For these chillers the thermodynamic cycle is driven by a compressor and they are said to be “electric driven”. Another kind of chiller is the absorption chiller, and is discussed in §2.3.3, where the thermodynamic cycle is driven by an absorption and regeneration process, which in turn is driven by a high temperature heat source; these chillers are therefore said to be “heat driven”.

The thermodynamic efficiency of a chiller is sensitive to the temperature difference between the chilled water and the condenser water. The lower the difference, the higher the chiller efficiency. Important parameters affecting chillers efficiency are:

- Chilled water supply temperature; if there is centralised dehumidification the cooling coils can run dry and handle only the sensitive cooling load. Hence, the chilled water supply temperature can be raised, especially considering the ASHRAE 2008 expanded recommended environmental envelope [10].
- Chilled water differential temperature; chiller efficiency is quite insensible to evaporator differential temperature. Therefore it is usually advantageous to keep this differential higher in order to reduce chilled water flow rate, hence the pump energy consumption.
- Condenser water differential temperature; chiller efficiency decreases with increase in condenser differential temperature. Therefore a compromise should be found between chiller efficiency and condensed water pump energy consumption.
- Variable speed compressor driver; compressors with VFD (Variable Frequency Driver) have significantly higher efficiency at part-load and so they payback for their higher capital cost.

Economizers

An economizer is a system that reduces the need for mechanical cooling by utilization of outdoor air under certain conditions. Economizers can be both on the air-side and on the water-side.

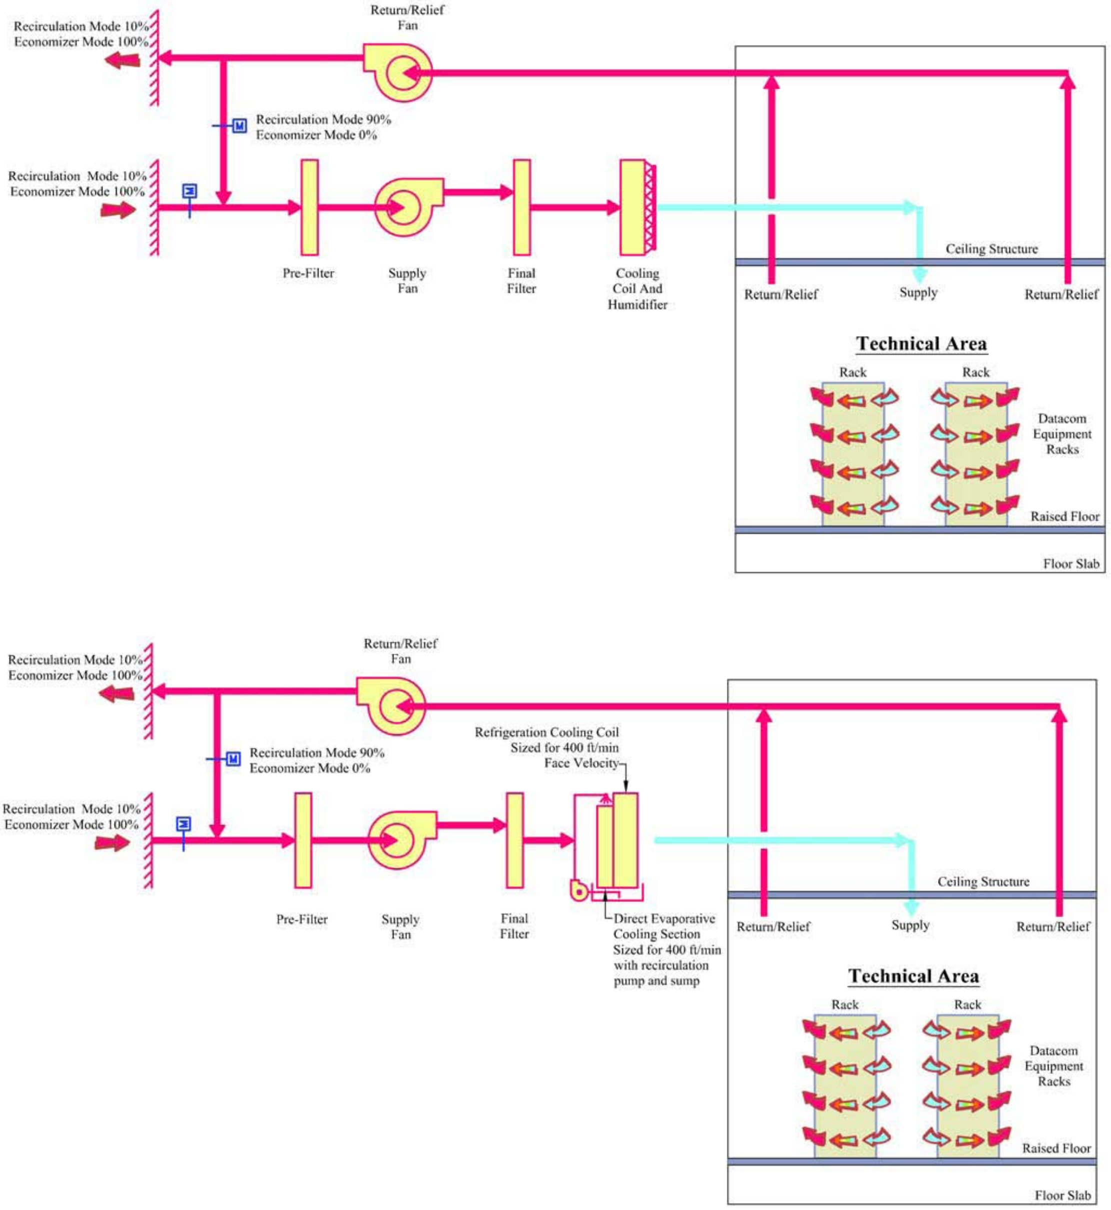

a)

b)

Figure 17. Schematic of *a)* air-side economizer in an air handling system, *b)* with evaporative coil for adiabatic cooling.

An air-side economizer, as the one shown in Figure 17*a)* and *b)*, is a system that has been designed to use outdoor air for cooling – partially or fully – when the outdoor air meets certain criteria; the control may be on dry bulb temperature or on total enthalpy. Air dampers are used to adjust the mixture of outdoor air and return air. In addition, an air-side economizer may include an evaporative coil for adiabatic cooling, see Figure 17*b)*. This device is mounted upstream the system cooling coil and can cool hot and dry air by spraying and vaporizing water in the outdoor air.

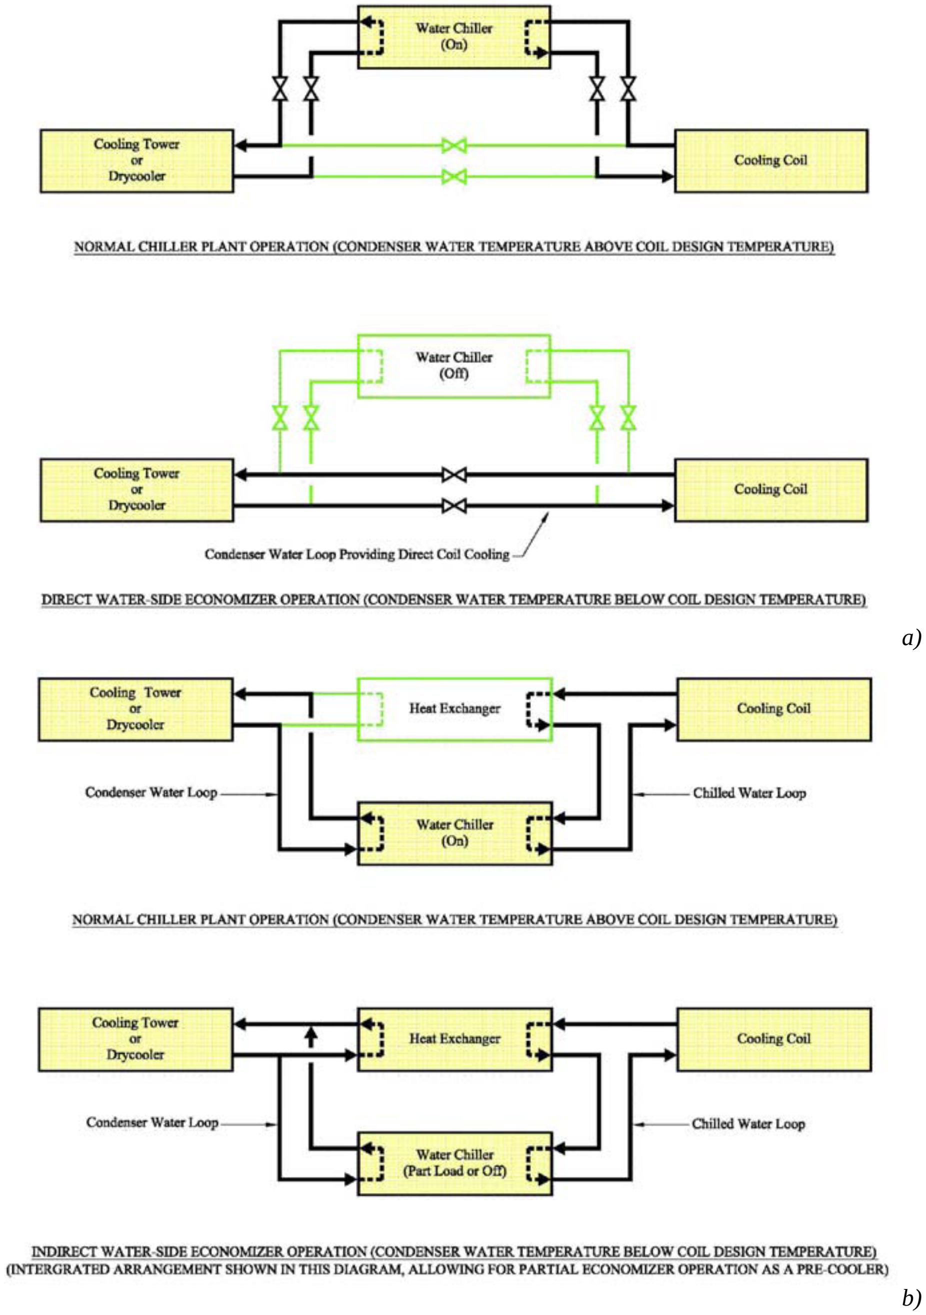

Figure 18. Schematic of water-side economizers *a)* direct, *b)* indirect.

Water-side economizers use cool outdoor air to generate condenser water that can partially or fully meet the data centre’s cooling requirements. There are two basic types of water-side economizers: direct and indirect, as shown in Figure 18*a)* and *b)*, respectively. In a direct system, the condenser water is circulated directly through the chilled-water circuit, hence in the cooling coils in the data centre rooms, while in an indirect system a heat exchanger is added to separate the condenser-water and chilled-water loops. A direct water-side economizer would be more efficient because it maximizes the number of annual economizer hours, but its operation can be more complex due to concerns associated with running condenser water directly through cooling coils (e.g. need for filtration). In the indirect approach, the cooling towers generate cold condenser water, which is then passed through a heat exchanger where it absorbs the heat from the chilled-water loop. When the condenser water is cold enough such that it can fully meet the cooling load, the water chillers can be shut off. Concerning the actual energy savings achievable and the impact on the TCO, the benefits of economizers strongly depend on climate. A summary is shown in Table 5.

Table 5. Air- and water-side economizer compared.

Economizer	**advantages**	**disadvantages**
Air-side	Higher annual utilisation Can be adiabatic (evaporative cooling)	De/humidification process efficiency is limiting factor
Water-side	Not constrained by de/humidification process	Lower annual utilisation (except dry climates)

2.3.2 Controls & Strategies

At building level as well as at data centre room level, the HVAC components and systems should be selected and designed for efficient part-load operation, e.g. with VFD (Variable Frequency Driver) drivers for fans, pumps, compressors. The single most important component for which efficiency should be maximised is the chiller.

Another fundamental point for energy efficiency is the control of economizers. These should be designed to allow for partial use, hence maximising their annual utilisation. Furthermore, it should be reminded that allowing more relaxed supply air temperature and humidity ranges in data centre rooms, as in [10], will increase cooling hours in economizer mode.

Analysis of the load and part-load for the full year are advisable in order to find the best control strategies for a variety of outdoor conditions. Computer simulation can be used to ensure high efficiency throughout the expected range of climatic conditions and part-load operations.

2.3.3 Emerging technology

Absorption cooling

A technology that is mature but not extensively exploited in data centres is the absorption cooling. An absorption cooling cycle is similar to a conventional compressor cycle in that it has an evaporator (extracting heat from the indoor), a condenser (rejecting the heat outdoor) and an expansion valve (adjusting pressure and temperature between the two components), see Figure 19. What an absorption cycle differs from a compression cycle is the way the cycle is ‘closed’. In a compression cycle the refrigerant fluid that has vaporised in the evaporator is compressed by mechanical work before entering the condenser and so restarting the cycle. The mechanical work is performed by a compressor, which is in turn driven by electricity, hence the term “electric driven” or “electric compressor cycle” that is also used to address a compressor cycle. In the absorption cycle, instead, the vaporised refrigerant is absorbed by an absorbent in a spontaneous reaction that releases some heat – for air conditioning the most common is to have water as the refrigerant and lithium-bromide as the absorbent. The refrigerant-absorbent mixture is then pumped into the generator (or re-generator) where the two substances are separated again, in virtue of their different boiling points, by a heating process at high temperature (normally above 100ºC). Altogether, the absorption and regeneration process acts like a “thermal compressor”, see Figure 19*right)*, because the refrigerant is delivered to the condenser at similar pressure and temperature conditions as if it had passed through an electric compressor. In last analysis, what drives the cooling cycle is a source of heat at high temperature for the absorption-regeneration process, instead of electricity for the compressor. The efficiency, or the COP, of an absorption cooling cycle is always lower than that of a compression cycle. However, absorption cooling is worth considering if any of the following factors apply:

- Waste heat is available from other processes
- Low cost source of heat is available (e.g. geothermal)

- An existing site has an electrical load limit that would be expensive to upgrade
- A site is particularly sensitive to noise and/or vibration
- Solar energy can be harnessed

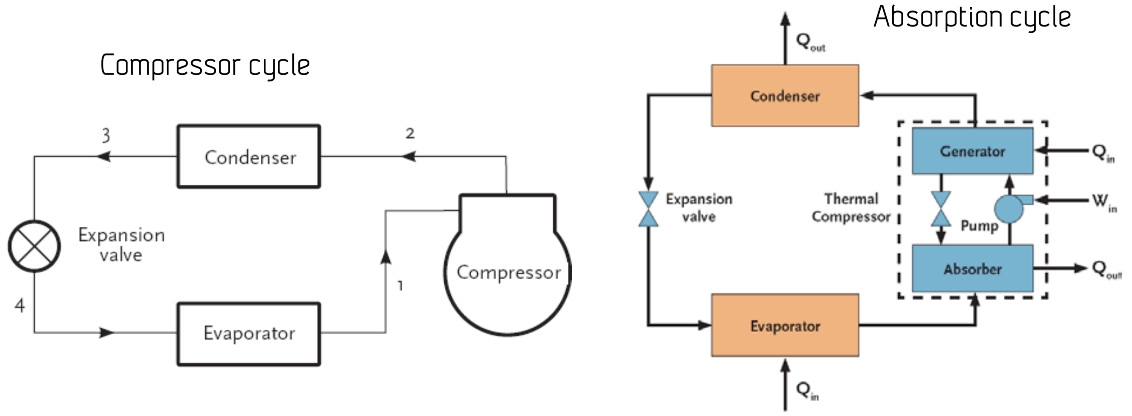

Figure 19. Cooling cycles: *left)* Compressor cycle, electric driven and *right)* absorption cycle, heat driven.

Thermal storage
A thermal storage system can be seen as the thermal equivalent of the electric back-up battery system. It may be used to provide uninterruptible cooling in case of short term chiller plant outage or simply to trim energy costs based on real-time pricing. Storage systems can be based on chilled water storage, ice storage or phase-change materials. Storage systems are space demanding if not designed to cover just short-term outages, especially if based on chilled water tanks, and may reduce the overall thermodynamic efficiency of the cooling process, especially if based on ice bank storage. Therefore a careful TCO analysis should be undertaken when considering and sizing a thermal storage system.

Other than HVAC
Among non-HVAC emerging technology for energy efficiency it is worth mentioning the DC power distribution system.

2.4 Outside the building

2.4.1 Components & Systems

Concerning the cooling system the main component to be found outside the building is the heat rejection device. Heat rejection devices are categorised in:

- Open cooling towers;
  Are more efficient and have higher capacity, hence are mostly used in large data centres with centralised systems. There is a need for water availability and there may be issues with protection against legionella or freezing.
- Dry coolers;
  Are less efficient and of smaller capacity, mostly used in small data centres with CRAC units.

2.4.2 Controls & Strategies

Also for heat rejection devices it is important to think about their efficiency at part load conditions (VFD drivers) and the effect they have on the efficiency of the overall cooling system. In particular, in the control of pumps or fans for heat rejection systems there is a trade off between higher delta temperature across the heat rejection device, ΔT, that would then require lower flows and reduce pump/fan consumption but would also lower the chiller efficiency, see also §2.3.1 on chillers. The control of the economizer according to outside air conditions involves also controlling of the heat rejection devices, see also §2.3.1 on economizers.

2.4.3 Emerging technology

Energy generation

Concerning HVAC installations, data centres that use absorption cooling machines may take advantage of a cogeneration system, such as Combined Heat and Power machines, CHP, and fuel cells. Such systems run on natural gas or biofuels and generate both electricity and heat at the same time. The electricity is directly used by the IT and auxiliary systems, while the waste heat is used to drive the absorption cooling machine.

Heat reuse

Inside a data centre there is rarely need for heating, hence the big amount of waste heat rejected at low temperature from a data centre may be reused outside the data centre itself. If the data centre is part of multi-use building or if it is part of a larger development, such as an office park, the waste heat may be transported and used in nearby offices for space heating and preheating of hot water. Reusing the heat outside the data centre does not increase the energy efficiency of the data centre itself but it helps reducing energy demand in the adjacent buildings.

Other than HVAC

Outside space of the data centre can be used to harvest renewable energy from the sun and wind and transform it into electricity using PV panels (Photo-Voltaic) or wind turbines, respectively. When a data centre becomes both a consumer and a producer of electricity it becomes also interesting to look into the possibilities offered by smart grid technology in order to minimise operating costs and optimise import and export with the grid.

3 Measuring and Reporting

3.1 Real-time measuring

In a report to the US congress the Environmental Protection Agency (EPA) analyses the power consumption of data centres and draws a number of possible scenarios for the future development [2]. The graph with projected aggregate US data centres energy demand is shown in Figure 20.

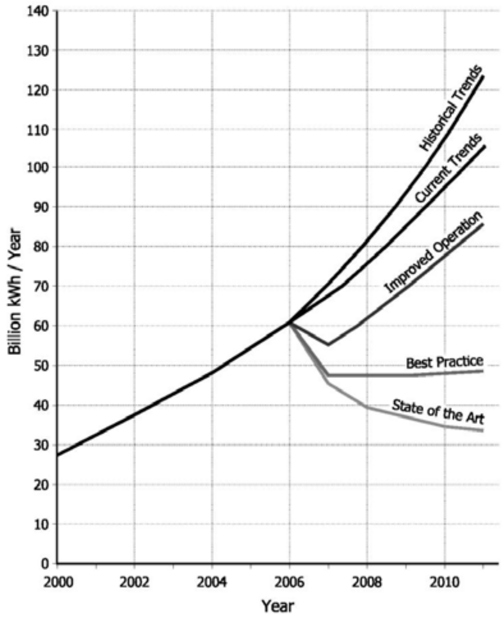

Figure 20. Projected data center energy use scenarios, source [2].

In [17] it is stated that one of the key ways in which the industry can achieve the state-of-the-art curve is via real-time energy efficiency, which is achievable only through the use of real-time energy consumption data using energy efficiency and productivity metrics. The reasons why real time energy efficiency can be so critical are derived from the previous chapters: data centres capacity is constantly increasing, the actual consumption depends on the equipment configuration and equipment is changed and upgraded quite frequently (3-years) causing an equivalent change in heat loads. Furthermore, the most efficient way to operate HVAC systems depends on the outdoor conditions, e.g. regulation of the economizer cycle. It should be noted that Real-time energy consumption measurements are only possible if all key subsystems are appropriately instrumented and properly communicating through use of data center level software, and that mixed-use facilities offer the greatest challenge in which to quantify real time energy consumption.

Real time energy measurements can apply also to the PUE indicator described in the next paragraph.

3.2 PUE (Power Usage Effectiveness) and PUE survey

Recognising the importance of establishing metrics for data centres efficiency, the Green Grid (a non-profit trade organisation of IT professionals) proposed in 2008 some indicators [18], whose measuring details were later analysed in more detail [19]. In the following years these indicators have been tested by members of the Green Grid and the most accepted one resulted being the PUE, Power Usage Effectiveness. As shown in Figure 21, the PUE measures the ratio between the total power use by the data centre, hence including also power needed to run the facility: switchgears, UPS, PDUs, HVAC systems, etc., and the power use devoted to IT equipment (meaning both pure IT and telecommunication equipment). Hence, for the PUE indicator: the lower the better, where a value of 1 represent an ideal target when all energy input to the data centre is used for IT purposes only.

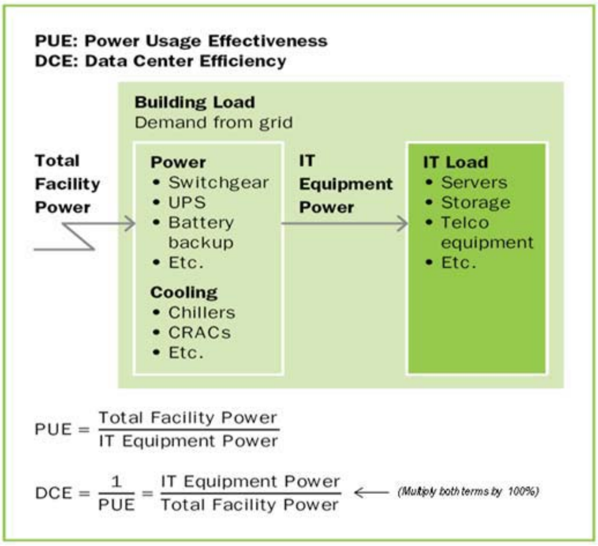

Figure 21. Data centre energy efficiency metrics, source [18].

In 2010 a group of data center industry leaders reached an agreement on the need for harmonised global metric for data center energy efficiency and formed a task force to work on the issue [20]. The agreement document states that the PUE is the preferred metrics for overall data centre facility energy efficiency. Besides, it is recognises the need to develop efficiency metrics also for measuring the actual IT-work and for renewable energy technologies and re-use of energy to reduce carbon emissions. For the time being, the formed task force published recommendations for measuring and reporting overall data center efficiency at dedicated data centers [21], sponsored by 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency's ENERGY STAR Program, United States Green Building Council and Uptime Institute, see Appendix A. Because in a data centre also other energy carriers than electricity may be employed, it is necessary to convert the energy uses of each specific carrier into "electricity equivalent" units. In [21] it is agreed to adopt source energy conversion factors derived from [22].

The fact that other energy carriers are weighted in terms of equivalent electricity units has an impact on the resulting PUE. Figure 22 shows the boundary for PUE calculation for two cases: with on-site generation of electricity from renewable sources – for example from solar or wind – and with the use of CHP (Combined Heat and Power) – the example is valid also for fuel cells. The use of renewable electricity does not affect the PUE calculation. Indeed, while on-site electricity generation may reduce the environmental impact and the CO_2 emissions associated with the operation of a data centre, it does not have any effect on its internal use of energy. The rationale is that the PUE shall measure how efficiently energy is used inside the data centre regardless of the energy supply. On the contrary, the use of CHP – assuming an absorption cooling system driven by the CHP waste heat – is seen as an efficiency measure rather than an alternative source of energy. Therefore the use of CHP should lower the PUE; this depends on the actual system. The overall efficiency of CHP machines or fuel cells (electrical efficiency + thermal efficiency) is normally around 80% or above. At the same time it should be considered that an absorption cooling system has a COP ranging from about 0.7 to 1.7 (with double or triple effect generators) compared to conventional compressor systems' COP of 3.5. However, the overall balance between the reduced purchase of electricity from the grid and the purchase of gas, would result in a reduced PUE when this is measured according to the conversion factors given in [21]. Furthermore, the fairly constant load of a data centre makes it a suitable field of application for CHP and fuel cells. Similarly, any reuse of the data centre waste heat within the same data centre infrastructure would reduce the PUE while a reuse in adjacent buildings would not have any impact on the PUE calculation.

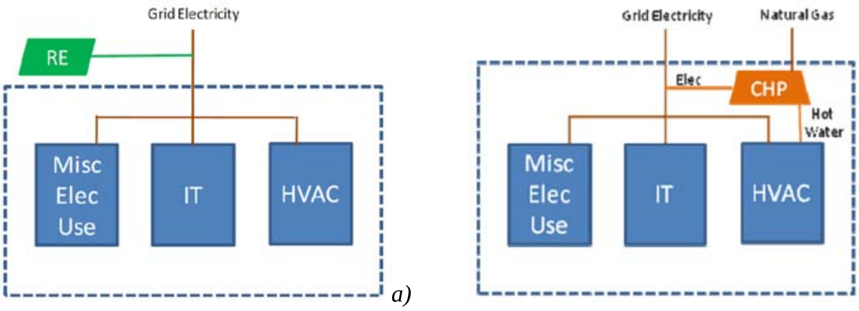

a) *b)*

Figure 22. Data centre PUE calculation boundary with *a)* renewable electricity on-site generation; *b)* CHP, Combined Heat and Power.

A survey performed by The Green Grid [23] shows that reported PUE falls roughly in the range between 1.5 and 3.0, with an average value of 2.03. These results are shown in Figure 23*a)* aggregated by PUE value distribution and in Figure 23*b)* by Tier level (where Tier from I to IV indicates the level of power supply redundancy – Tier I being the typical case for small facilities and Tier II being the most common case in large facilities). It appears, and it is not surprising, that large data centres are in average more efficient than small ones.

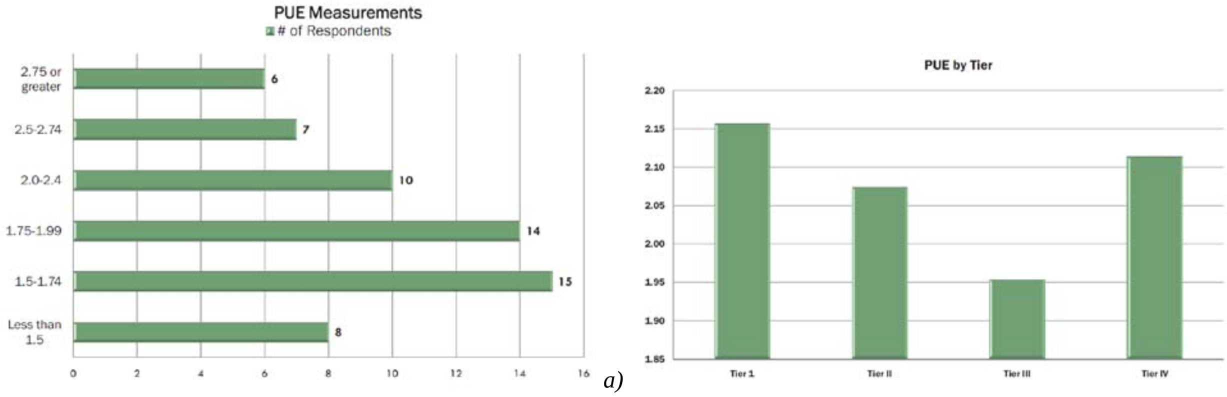

a) *b)*

Figure 23. PUE measurements *a)* by distribution; *b)* average by Tier level, source [23].

4 Data centres management

4.1 What do others do?

In one of its first activities The Green Grid conducted a survey in 2008 [24] in which participants, data center and IT facility managers, were asked what they experience being the main drivers and obstacles for undertaking energy efficiency measures; the results are summarised in Table 6. A similar survey was repeated in 2009 [23], and the results show, see Figure 24, that by far the three most important factors are: 1) rising energy costs, 2) need/desire to increase IT capacity and 3) corporate initiative to be "green".

Table 6. Drivers and obstacles for implementing energy efficiency measures, source [24].

Efficiency drivers	Obstacles
• Capacity Constraints • Cost Savings • Environmental Concerns • Impending Regulation	• Growth in IT demand • Lack of space in existing facilities • Risk aversion • Lack of funding for improvements • Shortage of expertise • Absent, weak or adverse incentives • Need for tools

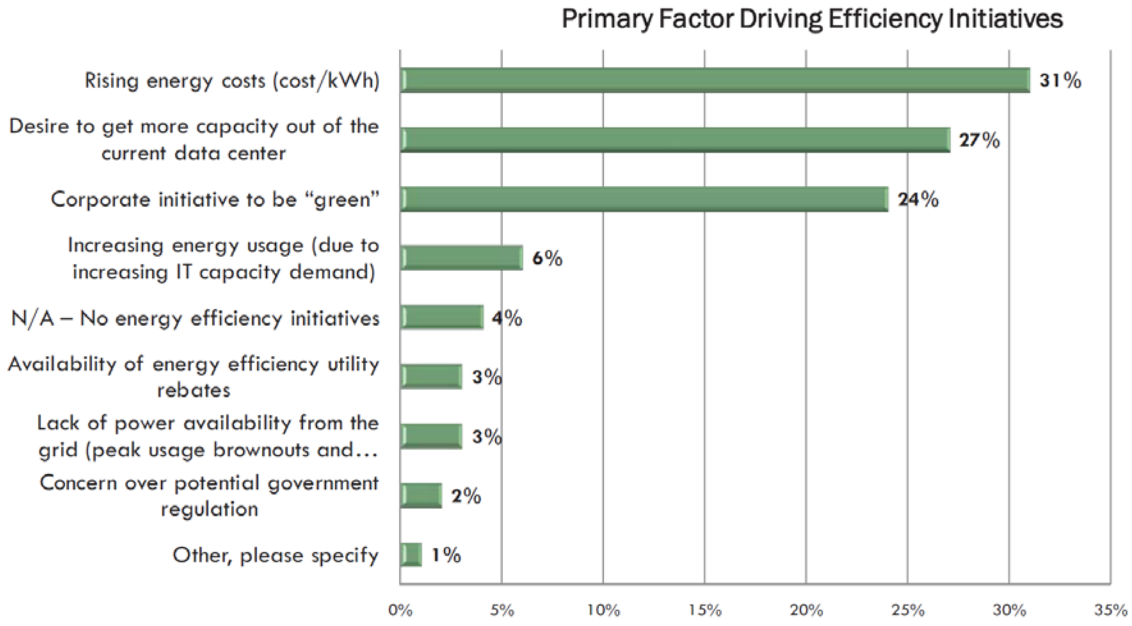

Figure 24. Primary factors driving energy efficiency initiatives, source [23].

In a further document, based on informal colloquia during a technical committee forum [25], the major findings were as follows:

- Efficiency continues to grow in importance
    - Being "green" can make also good financial sense
- Measuring and managing infrastructure equipment is too difficult
    - Reliability of measurements is not good; e.g. not possible to test if increasing the data center temperature cause servers to work harder

- While there are thousands of compute devices on the IT side, there is nothing similar on the infrastructure side
- Users are frustrated by the proprietary nature of today's infrastructure equipment and monitoring systems

With respect to the latter point, lack of comprehensive infrastructure monitoring system, the report states that: "The attendees of the session believe very strongly that this issue needs to be rectified for data centers to be able to operate as efficiently as they can and they felt that the Green Grid should take the lead in driving change in the industry. On their wish list for infrastructure monitoring is a system that is: vendor agnostic / brand agnostic, has an open interface, is factually scalable (to tens of thousands of data points), and is fast. While this is the baseline, there is also a desire for a system with the intelligence to analyze data with predictive assessment algorithms. A vendor that can do this well could create a competitive advantage. And finally, the ultimate benefit would be for this new infrastructure monitoring system to speak with existing IT systems, allowing for a common time-stamp and cause and effect analysis." This concern sounds as a major obstacle to the implementation of real time measuring and real time energy efficiency management as discussed in §3.1.

Concerning energy measurements, it appears somehow paradoxical that only about one third of the respondents has dedicated energy metering for the datacom equipment (compute and storage servers and telecom equipment) – that represents a data center core business. Metering of auxiliary electrical services and HVAC services such as cooling, instead, appear to be more common, see Figure 25.

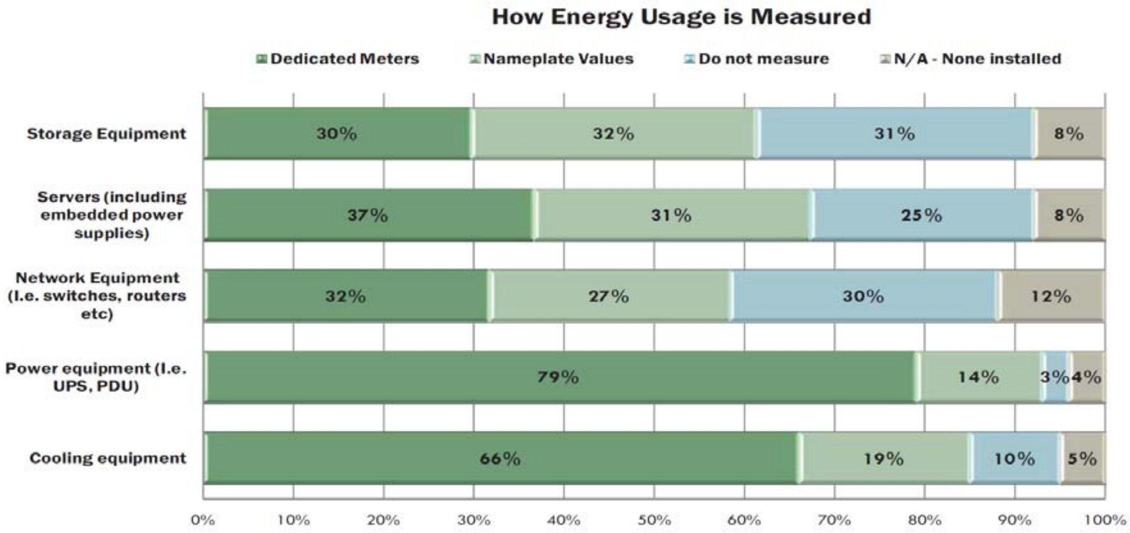

Figure 25. Energy measurements, source [23].

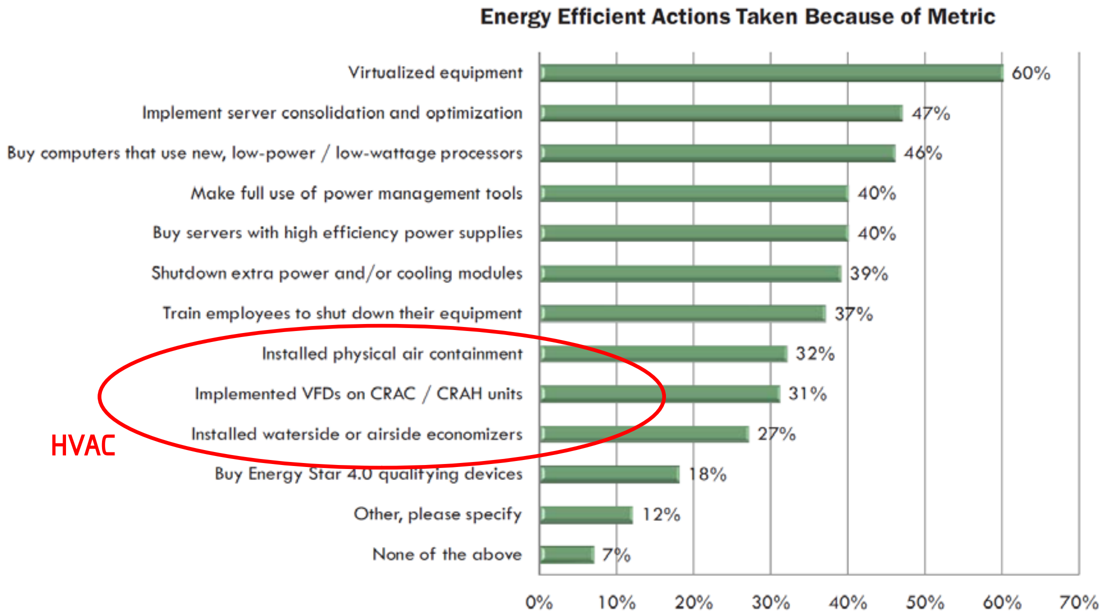

Figure 26. Actions prompted by efficiency measurements, source [23].

The same survey [23] also reports what energy efficiency actions were undertaken after that measurements were taken and in an attempt to improve a data center's PUE. Figure 26 shows that despite being HVAC measurements more spread than IT measurements, the actions taken to improve HVAC performance rank at the bottom of the list. There can be several reasons for that. First of all, IT energy efficiency measure such as virtualisation and consolidation of servers are virtually inexpensive and give short term results. Improvements on the HVAC system, on the other hand, are more difficult to implement, require more time, expertise and have an up-front cost. Another fundamental reason may be that data center managers have direct responsibility and knowledge on the IT side but lack competence or proper communication / coordination with the infrastructure side; often the IT and infrastructure personnel belong to different departments, and in some cases the HVAC services are even outsourced.

4.2 Energy efficiency checklist

The norm TIA-942 regulates in detail the data centre telecommunications space topology and cabling system. Unfortunately, there is not a similar norm regulating the HVAC system and related energy efficiency; the TIA-942 itself [12] does only include informative values on the environmental temperature and humidity, which are in turn taken from from the ASHRAE environmental guidelines for data centres [10]. Due to the variety and complexity of HVAC systems found in data centres, it is also not possible to summarise the related energy efficiency measures in a comprehensive and synthetic way. Guidelines for energy efficiency in selected data center subsystems exist in literature, as for example in [5], [8] and [13], and were considered in this sate of the art report. However, such guidelines are limited to a specific subject and are quite extensive.

In an attempt to synthetise the information discussed in the previous chapters, the energy efficiency checklist is proposed. The purpose of the checklist is solely to provide a quick overview of the most relevant HVAC system characteristics and related energy efficiency aspects. The list should by no means be regarded as exhaustive of the topics it mentions, and the literature in the references should be consulted for further insights.

References

[1] The Green Grid white paper #1 (2007) Guidelines for energy-efficient datacenters.

[2] Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431 (2007) U.S. *Environmental Protection Agency*, ENERGY STAR Program.

[3] Lawrence Berkeley National Lab. (2007) Benchmarking: data centers – Charts. Website http://hightech.lbl.gov/benchmarking-dc-charts.html, accessed 19/11/2010.

[4] ASHRAE Datacom series (2008) Book 2: Datacom Equipment Power Trends and Cooling Applications.

[5] ASHRAE Datacom series (2008) Book 1: Thermal Guidelines for Data Processing Environments.

[6] The Green Grid white paper #23 (2009) Proper sizing of IT power and cooling loads white paper.

[7] ASHRAE Datacom series (2008) Book 3: Design Consideration for Datacom Equipment Centers.

[8] ASHRAE Datacom series (2008) Book 4: Liquid Cooling Guidelines for Datacom Equipment Centers.

[9] Krav til ventilasjon og kjøling i IKT-rom (2008) UNINETT utkast, Trondheim, NO.

[10] 2008 ASHRAE Environmental Guidelines for Datacom Equipment – Expanding the Recommended Environmental Envelope.

[11] TIA-942 (2005) Telecommunications Infrastructure Standard for Data Centers, *Telecommunications Industry Association*.

[12] TIA-942-2 (2010) Telecommunications Infrastucture Standard for Data Centers Addendum 2 – Additional Media and Guidelines For Data Centers, *Telecommunications Industry Association*.

[13] The Green Grid white paper #11 (2008) Seven strategies to improve data center cooling efficiency.

[14] ASHRAE Datacom series (2008) Book 6: Best Practices for Datacom Facility Energy Efficiency.

[15] Fundamentals of Heat and Mass Transfer, Incropera, F. P. And De Witt, D. P. – 5th edition (2002), *John Wiley & Sons*, New York, US.

[16] High performance data centers (2006) A Design Guidelines Sourcebook, *Pacific Gas & Electric Company*.

[17] ASHRAE Datacom series (2008) Book 9: Real-Time Energy Consumption Measurements in Data Centers.

[18] The Green Grid white paper #6 (2008) Green grid data center power efficiency metrics: PUE and DCiE.

[19] The Green Grid white paper #22 (2009) Usage and public reporting guidelines for the green grid's infrastructure metrics (PUE/DCiE).

[20] Harmonizing Global Metrics for Data Center Energy Efficiency (2010) The United States of America, European Union and Japan Reach Agreement on Guiding Principles for Data Center Energy Efficiency Metrics.

[21] Recommendations for Measuring and Reporting Overall Data Center Efficiency: Version 1 – Measuring PUE at Dedicated Data Centers (2010) *Data Center Metrics Coordination Taskforce*, sponsored by 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency's ENERGY STAR Program, United States Green Building Council and Uptime Institute.

[22] Methodology for Incorporating Source Energy Use (2009) ENERGY STAR Performance Ratings, *Environmental Protection Agency*.

[23] The Green Grid white paper #26 (2009) Energy measurement survey results and analysis.

[24] The Green grid white paper #8 (2008): Data center baseline study report.

[25] The Green Grid (2010) Birds of a feather requirements gathering report.

Appendix A – Recommendation for Measuring and Reporting Overall Data Center Efficiency

Following document.

Recommendations for Measuring and Reporting Overall Data Center Efficiency

Version 1 – Measuring PUE at Dedicated Data Centers

15 July 2010

Table of Contents

15 July 2010

Recommendations for Measuring and Reporting Overall Data Center Efficiency

Version 1 – Measuring PUE at Dedicated Data Centers

1 Introduction

1.1 Purpose – Recommendations for Measuring and Publishing PUE at Dedicated Data Centers

The purpose of this document is to provide recommendations on measuring and publishing values for PUE at dedicated data center facilities. These recommendations represent the collective work of a task force representing 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency's ENERGY STAR Program, United States Green Building Council, and Uptime Institute. The recommendations are prepared in accordance with the guiding principles, presented in the following section. The task force continues to work towards a Version 2 document, which will provide corresponding guidance for mixed-use buildings.

1.2 Background – Guiding Principles for Data Center Efficiency Metrics

As business demands and energy costs for data centers rise, owners and operators have focused on the energy efficiency of the data center as a whole, frequently using energy efficiency metrics. However, the metrics are not always applied clearly and consistently. To address these inconsistencies, a group of leaders from across the industry met on January 13, 2010 to agree on data center energy efficiency measurements, metrics, and reporting conventions. Organizations represented included the 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now and Federal Energy Management Programs, U.S. Environmental Protection Agency's ENERGY STAR Program, U.S. Green Building Council, and Uptime Institute.

The following guiding principles were agreed to:

- Power Usage Effectiveness (PUE) using source energy consumption is the preferred energy efficiency metric for data centers. PUE is a measurement of the total energy of the data center divided by the IT energy consumption.
- When calculating PUE, IT energy consumption should, at a minimum, be measured at the output of the uninterruptible power supply (UPS). However, the industry should progressively improve measurement capabilities over time so that measurement of IT energy consumption directly at the IT load (i.e. servers) becomes the common practice.
- For a dedicated data center, the total energy in the PUE equation will include all energy sources at the point of utility handoff to the data center owner or operator. For a data center in a mixed-use building, the total energy will be all energy required to operate the data center, similar to a dedicated data center, and should include cooling, lighting, and support infrastructure for the data center operations.

This guidance is meant to help the industry have a common understanding of energy efficiency metrics that can generate dialogue to improve data center efficiencies and reduce energy consumption. Member organizations are committed to applying and promoting these guidelines to their programs.

A task force was created to further refine these metrics and to identify a roadmap for the future. The group also aspires to address IT productivity and carbon accounting in the future.

1.3 Scope of Recommendations

The purpose of this document is to provide recommendations from this task force on how to measure and publish overall data center infrastructure energy efficiency, based on the agreed upon guiding principles. This document does not address IT efficiency. It also does not directly address system-level metrics for cooling or heat rejection, air flow management, power distribution, lighting, etc.

The task force recognizes that many data centers operators may not currently have the capability to measure all energy consuming components within their facility accurately. The task force therefore recommends four (4) categories of measurement, which represent a subset of The Green Grid's (TGG) measurement methods[1]. These categories range from relatively simple measurements that provide a performance snapshot to more sophisticated measurement means that provide highly detailed performance data.

The intent is to encourage operators with limited measurement equipment to participate while also defining a framework that allows operators to add additional measurement points to increase the accuracy of their measurement program. The goal is to recommend a consistent and repeatable measurement strategy that allows data center operators to monitor and improve the energy efficiency of their facility. A consistent measurement approach will also facilitate communication of PUE among data center owners and operators. It should be noted that caution must be exercised when an organization wishes to use PUE to compare different data centers, as it is necessary to first conduct appropriate data analyses to ensure that other factors such as levels of reliability and climate are not impacting the PUE.

At present the scope of these recommendations is limited to dedicated data center facilities. Version 2 of this document will include recommendations for data centers that are part of larger mixed-use facilities.

These recommendations can also be used in the planning phase for a new data center or a major renovation of an existing facility to assist in the placement and number of measurement points and monitoring equipment.

This document provides an overview of how to measure and report PUE per this task force's recommendations. This document is not intended to serve as a detailed technical reference. Further specific details and requirements on how to measure PUE are available in The Green Grid white papers (http://www.thegreengrid.org/).

2 PUE Metric Calculation in Dedicated Data Centers

Power Usage Effectiveness (PUE) is the recommended metric for characterizing and reporting overall data center infrastructure efficiency. The task force strongly recommends annual energy consumption (kWh) for all energy types as the unit of measure for PUE calculation. However an entry level measurement category has been included in the recommendations to allow operators that do not have consumption measurement capability to utilize demand based power readings.

PUE is defined by the following formula:

PUE = (Total data center energy consumption or power / IT energy consumption or power)

---

[1] The Green Grid nomenclature and supporting information: http://www.thegreengrid.org/en/Global/Content/white-papers/Usage%20and%20Public%20Reporting%20Guidelines%20for%20PUE%20DCiE

The task force recommends four (4) measurement categories for the measurement of PUE within a data center facility:

PUE Category 0

This is a demand based calculation representing the peak load during a 12-month measurement period. IT power is represented by the demand (kW) reading of the UPS system output (or sum of outputs if more than one UPS system is installed) as measured during peak IT equipment utilization. Total data center power is measured at the utility meter(s) and is typically reported as demand kW on the utility bill. As this is a snapshot measurement, the true impact of fluctuating IT or mechanical loads can be missed. However consistent measurement can still provide valuable data that can assist in managing energy efficiency. PUE category 0 may only be used for all-electric data centers i.e. it cannot be used for data centers that also use other types of energy (e.g. natural gas, district chilled water, etc.).

PUE Category 1

This is a consumption based calculation. The IT load is represented by a 12-month total kWh reading of the UPS system output (or sum of outputs if more than one UPS system is installed). This is a cumulative measurement and requires the use of kWh consumption meters at all measurement points. The total energy is typically obtained from the utility company bills by adding the 12 consecutive monthly kWh readings as well as annual natural gas or other fuel consumption (converted to kWh) if present. This measurement method captures the impact of fluctuating IT and cooling loads and therefore provides a more accurate overall performance picture then PUE Category 0.

PUE Category 2

This is a consumption based calculation. The IT load is represented by a 12-month total kWh reading taken at the output of the PDU's supporting IT loads (or sum of outputs if more than one PDU is installed). This is a cumulative measurement and requires the use of kWh consumption meters at all measurement points. The total energy is determined in the same way as Category 1. This measurement method provides additional accuracy of the IT load reading by removing the impact of losses associated with PDU transformers and static switches.

PUE Category 3

This is a consumption based calculation. The IT load is represented by a 12 month total kWh reading taken at the point of connection of the IT devices to the electrical system. This is a cumulative measurement and requires the use of kWh consumption meters at all measurement points. The total energy is determined in the same way as Category 1. This measurement method provides the highest level of accuracy for measurement of the IT load reading by removing all impact of losses associated with electrical distribution components and non-IT related devices, e.g., rack mounted fans, etc.

Table 1 summarizes the four categories.

Table 1: PUE measurement categories recommended by this task force.

	PUE Category 0*	**PUE Category 1**	**PUE Category 2**	**PUE Category 3**
IT energy measurement location	UPS output	UPS output	PDU output	Server input
Definition of IT energy	Peak IT electric demand	IT annual energy	IT annual energy	IT annual energy
Definition of Total energy	Peak Total electric demand	Total annual energy	Total annual energy	Total annual energy

*For PUE Category 0 the measurements are electric demand (kW).

2.1 Definitions

Dedicated data center: In this context, a dedicated data center is a facility in which all the spaces and supporting infrastructure (HVAC, lighting, electrical) are directly associated with the operation of the data center. Dedicated data centers are stand alone buildings whose exclusive purpose is IT operations and their support functions.

UPS Output is defined as the output of the UPS that serves IT equipment loads. UPS output does not include efficiency losses from the UPS system but does include losses from downstream electrical distribution components such as PDU's. It may include non-IT ancillary devices installed in IT racks such as fans. If there is non-IT equipment supported by the UPS system, (e.g., CRAC, CRAH, In-row coolers, etc.) it must be metered and subtracted from the UPS Output reading (the metering approach should be consistent with the metering required for the PUE category e.g. continuous consumption metering for PUE categories 1,2 and 3).

PDU Output is defined as the output of the PDU that serves IT loads. PDU output does not include efficiency losses of any transformation that occurs within the PDU but may include downstream non-IT ancillary devices installed in IT racks such as fans. If there is non-IT equipment supported by the PDU system, (e.g., CRAC, CRAH, In-row coolers, etc.) it must be metered and subtracted from the PDU Output reading (the metering approach should be consistent with the metering required for the PUE category e.g. continuous consumption metering for PUE categories 1,2 and 3).

Server input is defined as the IT load as measured at the point of connection of the IT device to the electrical power system, e.g., the power receptacle(s). Server input captures the actual power load of the IT device exclusive of any power distribution losses and non-IT loads such as rack mounted fans.

Peak IT electric demand (kW) is the annual peak electric demand for the IT load.

Peak Total electric demand (kW) is the annual peak electric demand for the entire data center including IT and supporting infrastructure measured at the point of utility handoff.

IT annual energy is calculated as annual electrical consumption of the IT. IT includes all IT equipment, servers, networking and storage as well as telecom equipment typically installed in telecom data centers.

Total annual energy is calculated as the weighted sum of the annual energy consumption for all energy types serving the data center at the point of utility handoff. This includes electricity, natural gas, fuel oil, and district utilities such as supplied chilled water or condenser water. All energy types must be converted into the same units before they are summed. Total annual energy must include supporting infrastructure (see definition below). Section 2.2 describes how to weight different energy types. (Note that weighting only applies for data centers that have energy types in addition to electricity.)

Supporting infrastructure includes the following:

- Power systems: Transfer switch, UPS, DC batteries/rectifiers (non UPS – telco nodes), generator, transformer (step down), power distribution unit (PDU), rack distribution unit (RDU), breaker panels, distribution wiring, lighting.
- HVAC systems: Cooling towers, condenser water pumps, chillers, chilled water pumps, computer room air conditioners (CRAC's), computer room air handlers (CRAH's), dry cooler, supply fans, return fans, air economizer, water-side economizer, humidifier, in-row, in-rack, & in-chassis cooling solutions.
- Physical security: Fire suppression, water detection, physical security servers/devices.
- Building management systems: Server/devices used to control management of data center, probes/sensors.

2.2 Weighting of Energy Types Based on Source Energy

For data centers that have electricity as well as other energy types, the different energy types must be weighted according to their source energy. Source energy represents the total amount of raw fuel that is required to operate the building. It incorporates all transmission, delivery, and production losses, thereby enabling a complete assessment of energy efficiency in a building. Table 2 provides the weighting factors for each energy type, normalized to electricity. These weighting factors are based on national average source factors used by EPA in their building energy benchmarking.

Weighted energy for each energy type = (Annual energy use X source energy weighting factor).

Note that all energy types must be converted into the same units before they are summed[2]. For example, if electricity is in kWh and natural gas is in kBtu, both must be converted to a common unit.

Table 2: Source energy weighting factors[3]

Energy Type	*Weighting Factor*
Electricity	1.0
Natural gas	0.31
Fuel oil	0.30
Other fuels	0.30
District chilled water	0.31
District hot water	0.40
District steam	0.43

Appendix A has an example of PUE calculation using source energy Weighting Factors.

---

[2] For conversion factors, see: http://www.energystar.gov/ia/business/tools_resources/target_finder/help/Energy_Units_Conversion_Table.htm

[3] The EPA source energy factor methodology is available at: http://www.energystar.gov/ia/business/evaluate_performance/site_source.pdf.

The weighting factors presented in this recommendation document are obtained by dividing each EPA source factor by the reference source factor for electricity, 3.34.

2.3 Renewable Energy Sources

Electricity from renewable energy (RE) sources should be included in the total energy or power and assigned the same source factor as grid electricity. In other words, RE sources are outside the PUE calculation boundary (figure 1). Therefore, the installation of RE sources does not change the PUE i.e. an all-grid building and all-RE building would have the same PUE. The rationale for this is that the purpose of the PUE is to evaluate how efficiently energy is used in the data center, regardless of the energy supply. Renewable energy in this context includes solar and wind power.

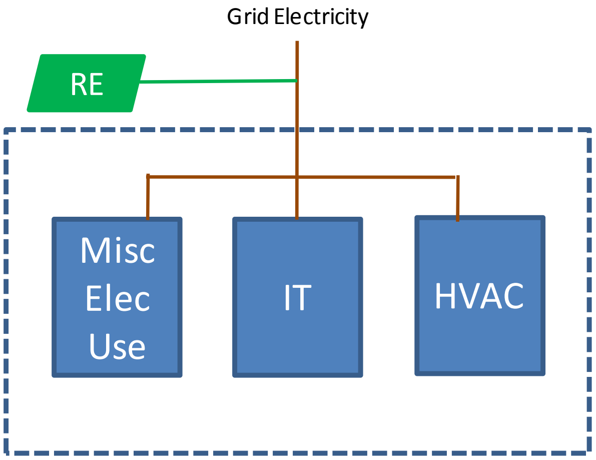

Figure 1: Data center PUE calculation boundary with renewable energy

2.4 Combined Heat and Power Plants

If the data center facility has a dedicated combined heat and power (CHP) plant, the inputs to the CHP should be included in the total energy or power (figure 2) and assigned the same source factors as those shown in table 1. The outputs from the CHP should not be included in the total energy or power. The use of CHP (like the use of more efficient cooling or other equipment) should lower the PUE. The rationale for this is that CHP is considered to be an efficiency measure rather than an alternative source of energy.

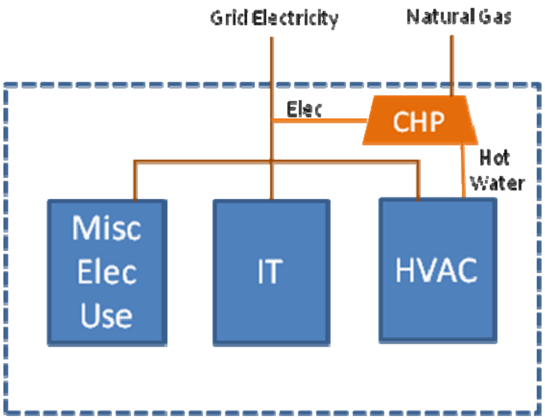

Figure 2: Data center PUE calculation boundary with CHP.

2.5 Reused Data Center Energy

The concept of reusing waste energy (generally in the form of heat) is receiving a lot of study as part of an overall effort to improve energy efficiency. The value of this energy may or may not affect PUE, depending on the location of the re-use.

- If the energy is reused within the datacenter, the value of the reused energy will already be contained within the PUE calculation and needs no special consideration. An example will be reusing waste heat to warm a battery space in a cool climate. The PUE would generally be lower because the HVAC system does not have to add as much energy to heat that battery space, thereby reducing the PUE when energy is measured and PUE calculated.
- If the energy is reused outside the data center (e.g. re-using heat to warm a lab that is not part of the data center), the PUE of the data center is not affected. While the effort to conserve energy is laudable, it cannot be accounted for in the calculation of the data center PUE due to the strict definition of PUE used in these recommendations. Currently there are on-going industry efforts to define a metric that could be used to account for this beneficial use, but it is specifically excluded from PUE.

3 Publishing Format and Related Resources

When publishing PUE, the category must clearly be indicated using a subscript e.g. PUE_0, PUE_1, PUE_2, PUE_3. A PUE reported without the subscript is not considered to be in compliance with these recommendations.

Note that PUE_0 may only be used for all-electric data centers i.e. it cannot be used for data centers that also use other types of energy (e.g. natural gas, district chilled water, etc.). An all-electric data center can utilize any of the PUE categories.

When publishing PUE calculated in accordance with these guidelines, users may state that "PUE was calculated in accordance with the recommendations of the Data Center Metrics Coordination Taskforce, sponsored by 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency's ENREGY STAR Program, United States Green Building Council, and Uptime Institute."

3.1 Related Resources and Tools

The Green Grid (TGG) white paper #22[4] provides a range of options for calculating PUE, as well as an annotation system to indicate how PUE was calculated. The recommendations of this task force are essentially a subset of the options in TGG framework, as shown in Table 2, and are therefore consistent with TGG framework.

Table 2: Mapping of PUE categories to TGG framework

PUE category	*TGG annotation*
PUE_0	$PUE_{L1,Y--}$
PUE_1	$PUE_{L1,YC}$
PUE_2	$PUE_{L2,YC}$
PUE_3	$PUE_{L3,YC}$

---

[4] The Green Grid nomenclature and supporting information: http://www.thegreengrid.org/en/Global/Content/white-papers/Usage%20and%20Public%20Reporting%20Guidelines%20for%20PUE%20DCiE

The U.S. EPA Energy Star Portfolio Manager tool utilizes a PUE definition that is consistent with PUE_1.

The U.S. DOE's DC Pro tool uses PUE definitions that are consistent with PUE_2.

The forthcoming LEED Data Centers standard has information related to the amount and type of monitoring equipment that are either prerequisites or for achieving additional points in the monitoring and verification credit category. The PUE definition used is consistent with $PUE_{2.}$

4 Next Steps

The task force will work to refine these recommendations so that they maybe applied at mixed use buildings, which will be presented in Version 2. In the interim it is recommended that dedicated data centers begin to adopt these recommendations when publishing PUE values. When publishing PUE calculated in accordance with these guidelines, users may state that "PUE was calculated in accordance with the recommendations of the Data Center Metrics Coordination Taskforce, sponsored by 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency's ENERGY STAR Program, United States Green Building Council, and Uptime Institute."

Appendix A. Example PUE Calculation Using Source Energy Weighting Factors

PUE should be based upon Source Energy as described in Section 2.2 of this document as well as the EPA source energy factor methodology as described in:
http://www.energystar.gov/ia/business/evaluate_performance/site_source.pdf.

Table 2 from the main document is reproduced here for convenience

Energy Type	*Weighting Factor*
Electricity	1.0
Natural gas	0.31
Fuel Oil	0.30
Other fuels	0.30
District Chilled water	0.31
District hot water	0.40
District steam	0.43

Each component of the PUE calculation needs to be multiplied by the appropriate Weighting Factor.

Recall the definition of PUE: Total Energy divided by IT Energy.

$$PUE = \frac{TotalEnergy}{ITEnergy}$$

Total energy should include all fuel sources. For an example data center that purchases electricity and district chilled water, the equation can be expressed as:

$$PUE = \frac{Electricty + DistrictChilled\ Water}{IT\ Energy}$$

To compute a PUE in accordance with these recommendations it is critical to include all fuels. For illustrative purposes, consider two fictitious data centers: Case A is an all electric data center, while Case B purchases chilled water from a local utility. The following tables show the input fuels, the end uses, and the approximate energy use for each end use. In order to compute the PUE it is not necessary to have the energy for each end use separately. The total energy is simply summed across all energy inputs to the data center (at the point of utility hand-off). The IT energy is the only end use that must be sub-metered within the building.

Case A – All Electric Data Center

Energy Input to Data Center	*End-use*	*Energy Use*
Electricity (1,705,000 kWh total)	IT Load	1,000,000 kWh
	Power delivery loss	250,000 kWh
	Lighting	50,000 kWh
	Cooling	400,000 kWh
	Other	5,000 kWh

The PUE for Case A would be computed as:

$$PUE = \frac{1{,}705{,}000 * (1.0)}{1{,}000{,}000 * (1.0)} = 1.70$$

Case B – Data Center with District Chilled Water

Energy Input to Data Center	*End-use*	*Energy Use*
Electricity (1,305,000 kWh total)	IT	1,000,000 kWh
	Power system loss	250,000 kWh
	Lighting	50,000 kWh
	Other	5,000 kWh
District Chilled water (1,300,000 kWh total)	Cooling	1,300,000 kWh (4.44 M BTUs)

In Case B, district chilled water is used to provide cooling (remove excess heat). The PUE for Case B would be computed as:

$$PUE = \frac{1{,}305{,}000 * (1.0) + 1{,}300{,}000 * 0.31}{1{,}000{,}000 * (1.0)} = 1.70$$

It is informative to consider that the PUE for each case is 1.70. The implication of this is that the efficiency of the cooling system in the all electric data center is identical to the efficiency with which the District Chilled Water is produced. If the Case A's cooling system was more efficient, then it would have required less than the 400,000 kWh in the example and the PUE would be lower. If Case A's cooling system is less efficient then it would have used more than the 400,000 kWh to cool the data center (including distribution losses and lighting) and the PUE would be higher.

Appendix B – Checklist for energy efficiency considerations

Planning and operation	
Inside racks and cabinets	
• Components & Systems	
o *Heat load*	Consider power trend evolution, recognise that datacom equipment loads will change over the next 10-15 years.
o *Proper sizing*	Use equipment thermal report, vendor estimation tools or real-time measurements. Never use equipment nameplate values or simple de-rating methods.
o *Air and liquid cooling*	Rack liquid cooling advisable for heat density > 5 kW/rack and necessary for > 15kW/rack. Evaluate closed or open architecture.
• Controls & Strategies	
o *Recommended environmental envelope*	Colder does not mean better! Use ASHRAE 2008 expanded recommended environmental envelope.
o *Other than HVAC*	Virtualisation of servers. Identify and manage unused servers.
• Emerging technology	
o *Datacom equipment liquid cooling*	Liquid cooling of datacom equipment (from manufacturer).
In the data centre room(s)	
• Components & Systems	
o *Air cooling*	Small facilities typically use CRACs in the rooms: simplest system. Larger facilities use centralised chilled water system with CRAHs in the rooms, or completely centralised air system with ducts for air distribution to the rooms. • Load < 3 kW/rack Air distribution in the rooms with hot/cold aisles and raised floor. Use blanketing panels to minimise air mixing inside racks. • 3 < Load < 5 kW/rack Use overhead ducts and/or aisle separation bafles to minimise air mixing outside the racks. • Load < 15 kW/rack Use fully controlled airflow path to/from the racks, and use racks equipped with own ventilation fan that allow effective airflow cooling.
o *Liquid cooling*	Consider using a Coolant Distribution Unit, CDU, to separate liquid cooling water loop from the chilled-water loop (avoid condensation in the room, minimise filtration, isolation of leakages). Realise a piping architecture in consideration of redundance, leakages and flexibility. Choose appropriate cooling fluid for the application.
o *Other than HVAC*	Arrange electric cabling to be non-obstructive of the under floor air path. Use efficient UPS, Uninterruptable Power Supply. Use efficient PDU, Power Distribution Unit.
• Controls & Strategies	
o *Centralised humidity control*	Eliminate adjacent CRACs/CRAHs ”fighting” – simultaneous humidification adn dehumidifying. Cooling coils in the rooms ”run dry” – allow for chilled

	water T reset at light loads
o *Part-load operation*	HVAC components and systems selected and designed for efficient part-load operation, e.g. VFD (Variable Frequency Driver) drivers for fans, pumps, compressors.
• Emerging technology	
o *Computational Fluid Dynamics, CFD*	Use CFD to optimise the layout.
Inside the Building	
• Components & Systems	
o *Chillers*	The single most energy consuming component, careful attention in selecting chiller and its design conditions – balance between low ΔT (higher chiller efficiency) and high flow (increased pump energy).
o *Economizers*	Air-side and water-side, designed also for partial use to increase annual utilisation – that means water-side must have a heat exchanger.
• Controls & Strategies	Develop an efficient part-load operation sequence to maximise chillers efficiency. Maximise annual economizer utilisation - higher supply air T in data centre room allow increased cooling hours in economizer mode. Computer simulation to ensure high efficiency throughout the expected range of climatic conditions and part-load operation.
• Emerging technology	
o *Absorption cooling*	The cooling cycle is heat driven (with high T source) instead of electric driven.
o *Thermal storage*	For uninterruptible cooling and stabilisation of the cooling system.
o *Other than HVAC*	DC power distribution.
Outside the building	
• Components & Systems	Small capacity: dry-coolers. High capacity: open cooling towers (need water availability).
• Controls & Strategies	Overall efficiency of the condenser-water loop to be evaluated in coordination with chiller and pumps.
• Emerging technology	
o *Energy generation*	Use of cogeneration with Combined Heat and Power machines, CHP, or fuel cells. This fits well with the constant load of a data centre and the utilisation of heat to run an absorption cooling system.
o *Heat reuse*	Mainly in adjacent buildings.
o *Other than HVAC*	On-site electricity generation from solar PV and wind. Use of smart-grid (where available) functionality to minimise operating costs.
Measuring and reporting	
• Real-time measuring	Real-time measurements are necessary to achieve real-time energy efficiency at state-of-the-art data centres.
• Power Usage Effectiveness, PUE	Measure and report PUE according recommendations of the Data Center Metrics Coordination Taskforce, see Appendix A or reference [21].

SINTEF

Technology for a better society

www.sintef.no