ETSI TS 105 174-2-2 V1.1.1 (2009-10)

Technical Specification

**Access, Terminals, Transmission and Multiplexing (ATTM);
Broadband Deployment - Energy Efficiency
and Key Performance Indicators;
Part 2: Network sites;
Sub-part 2: Data centres**

Reference

DTS/ATTM-02005-1-3

Keywords

access, cable, optical, site engineering

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Copyright Notification

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Access, Terminals, Transmission and Multiplexing (ATTM).

The present document is part 2-2 of a multi-part deliverable. Full details of the entire series can be found in part 1 [13].

Introduction

The increasing interaction between the different elements of the Information Communication Technology (ICT) sector (hardware, middleware, software and services) supports the concept of convergence in which:

- multi-service packages can be delivered over a common infrastructure;
- a variety of infrastructures is able to deliver these packages;
- a single multi-service-package may be delivered over different infrastructures.

As a result of this convergence, the development of new services, applications and content has resulted in an increased demand for bandwidth, reliability, quality and performance, with a consequent increase in the demand for power which has implications for cost and, in some cases, availability. It is therefore important to maximize the energy efficiency of all the network elements necessary to deliver the required services.

New technologies and infrastructure strategies are expected to enable operators to decrease the energy consumption, for a given level of service, of their existing and future infrastructures thus decreasing their costs. This requires a common understanding among market participants that only standards can produce.

The present document is part 2-2 of a multi-part deliverable which has been produced by ETSI Technical Committee Access, Terminals and Transmission, Multiplexing (ATTM) in close collaboration with CENELEC via the Co-ordination Group on Installations and Cabling (CGIC). It offers a contribution to the required standardization process by establishing an initial basis for work on ICT networks and transmission engineering, with active collaboration from a number of other ETSI and CENELEC Technical Bodies. When complete, the multi-part deliverable will contain information that has been jointly evolved to present developments in installations and transmission implementation, and describing their progress towards energy efficiency in Next Generation Networks (NGN).

In order to monitor the implementation and operation of energy efficient broadband deployment, the present document also discusses Key Performance Indicators (KPI) for energy efficiency and focus on the possible consequences of standardization of installations, cabling techniques and equipment. In particular, the study will investigate possibilities and suggest solutions for development of processes for optimization in installation techniques and energy consumption.

1 Scope

The present document details measures which may be taken to improve the energy efficiency within operators sites and data centres for broadband deployment. Clauses 2 and 3 contain references, definitions and abbreviations which relate to this part; similar information will be included in the corresponding clauses of the other parts, thus ensuring that each document can be used on a "stand-alone" basis.

Within the present document:

- clause 4 introduces data centre concepts including those specifically related to network operators;
- clause 5 develops the concept of Key Performance Indicators (KPI), introduced in TS 105 174-1 [13], to enable consistent monitoring of energy efficiency;
- clause 6 details the approaches that may be employed to improve energy efficiency within the information technology infrastructure;
- clause 7 details the approaches that may be employed to improve energy efficiency within the environmental control systems;
- clause 8 details the approaches that may be employed to improve energy efficiency via the physical infrastructure of the buildings;
- clause 9 details the approaches that may be employed to improve energy efficiency within the power distribution system;
- clause 10 provides a summary of energy efficiency approaches within existing data centres;
- clause 11 provides a summary of energy efficiency approaches within new data centres and introduces wider issues concerning their location;
- clause 12 contains the conformance mechanisms of the present document;
- clause 13 contains the recommendations of the present document;
- clause 14 introduces future opportunities for improvements of energy efficiency;
- annex A provides indications of the first order effect of applying the approaches outlined in clauses 6, 7 and 9.

This will enable the proper implementation of services, applications and content on an energy efficient infrastructure, though it is not the goal of this multi-part deliverable to provide detailed standardized solutions for network architecture.

The present document focuses on energy efficiency. The CO_2 footprint is not taken in account in the present document.

Two separate aspects of energy efficiency are considered as shown in figure 1:

- actions to improve energy efficiency in existing data centres in the short or medium term;
- actions to improve energy efficiency in new data centres, in medium or long term.

The domains under study are:

- in the Information Technology (IT) infrastructure: all aspects of the technical infrastructure in the data centre, including servers, storage arrays, backup libraries and network equipment including routers, switches, etc.;
- in the IT operational strategy: all consolidation initiatives, such as virtualization, physical or logical consolidations, usage of specific software and processes;
- in the technical environment: all aspects concerning energy usage, cooling and, more generally, all disciplines involved in the technical environment of the data centre.

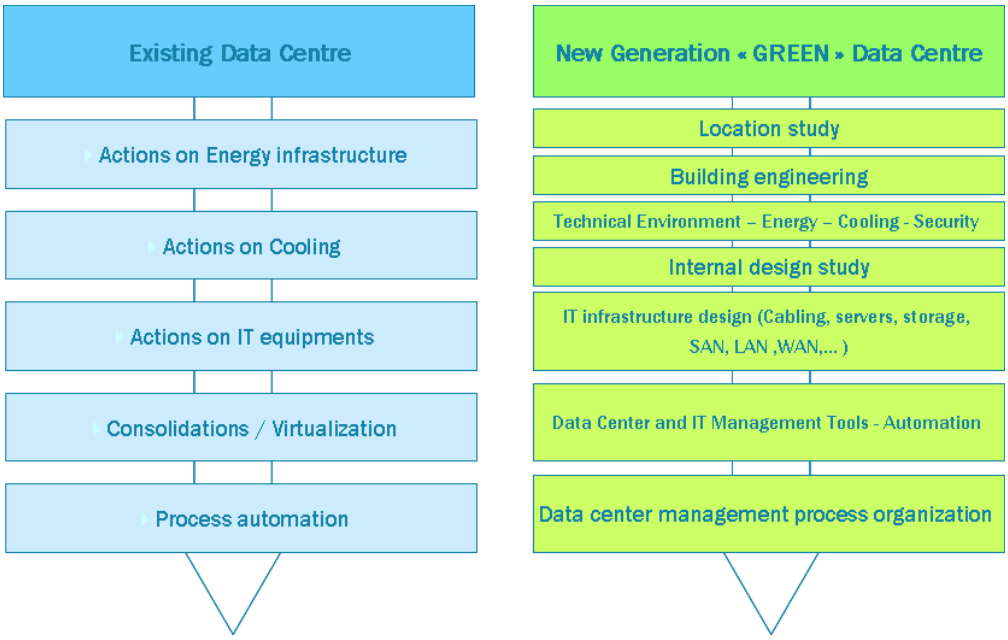

Figure 1: Aspects of data centres under consideration

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific.

- For a specific reference, subsequent revisions do not apply.
- Non-specific reference may be made only to a complete document or a part thereof and only in the following cases:
  - if it is accepted that it will be possible to use all future changes of the referenced document for the purposes of the referring document;
  - for informative references.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are indispensable for the application of the present document. For dated references, only the edition cited applies. For non-specific references, the latest edition of the referenced document (including any amendments) applies.

[1] ANSI/TIA-942: "Telecommunications Infrastructure Standard for Data Centres".

[2] Uptime Institute: "Tier Classifications Define Site Infrastructure Performance".

[3] Johannesburg: "Datacenter Dynamics Research Key Findings" August 2008.

[4] European Commission: "DG-JRC Code of Conduct on Data Centres Energy Efficiency".

[5] "Best Practices for the EU Code of Conduct on Data Centres".

[6] CENELEC EN 50173-2: "Information technology - Generic cabling systems - Part 2: Office premises".

[7] CENELEC EN 50173-5: "Information technology - Generic cabling systems - Part 5: Data centres".

[8] CENELEC EN 50174-1: "Information technology - Cabling installation - Part 1: Installation specification and quality assurance".

[9] CENELEC EN 50174-2: "Information technology - Cabling installation - Part 2: Installation planning and practices inside buildings".

[10] High performance buildings: "UPS report (Ecos Consulting-Epri Solutions)".

[11] ETSI EN 300 019-1-3: "Environmental Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment; Part 1-3: Classification of environmental conditions; Stationary use at weatherprotected locations".

[12] ETSI EN 300 132-3: "Environmental Engineering (EE); Power supply interface at the input to telecommunications equipment; Part 3: Operated by rectified current source, alternating current source or direct current source up to 400 V".

[13] ETSI TS 105 174-1: "Access, Terminals, Transmission and Multiplexing (ATTM); Broadband Deployment - Energy Efficiency and Key Performance Indicators; Part 1: Overview, common and generic aspects".

2.2 Informative references

The following referenced documents are not essential to the use of the present document but they assist the user with regard to a particular subject area. For non-specific references, the latest version of the referenced document (including any amendments) applies.

[i.1] ETSI TR 102 489: "Environmental Engineering (EE); European telecommunications standard for equipment practice; Thermal Management Guidance for equipment and its deployment".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

application: single program or a set of several programs executing a function or a service

availability: time or period during the application or the service has to be operational

NOTE: Availability is one of the criticality criteria.

blade server: server chassis housing multiple thin, modular electronic circuit boards, known as server blades

NOTE: Each blade is a server in its own right, often dedicated to a single application. The blades are literally servers on a card, containing processors, memory, integrated network controllers, an optional fibre channel host bus adaptor (HBA) and other input/output (IO) ports.

computer room: closed, secured and environmentally controlled room in which IT equipment is operating

criticality: level given to an application or service, linked to the impact for the enterprise in case of crash

NOTE: More the impact is strong, more the application or service is critical.

data centre: centralized repository for the storage, management, and dissemination of data and information organized around a particular body of knowledge or pertaining to a particular business

Data Centre Infrastructure Efficiency (DCIE): reciprocal of the PUE, that is "IT equipment power" divided by "total facility power", expressed as a percentage

NOTE: DCIE improves as it approaches 100 %.

Disaster Recovery Plan (DRP): all process (technical, organization, people) to launch in case of continuity disruption

disk array: cabinet containing physical disks

energy efficiency: search in existing DC, or for new future DC, of all tracks and actions allowing minimizing energy needs and costs

NOTE: Key drivers are Economic to decrease the energy bill by increasing the efficiency of all equipment and minimize power loss.

green data centre: in addition to energy efficiency, the "Green" approach will focus on carbon footprint

NOTE 1: Energy Efficiency is one way, to decrease CO_2 emissions, but it is not the only one.

NOTE 2: More "sustainable development" objective than economic, the key indicator is carbon footprint. Today, this concept is not still clearly defined, especially if we now that data centres are not directly producers of CO_2, but indirectly, due to their energy needs. If the source of power is becoming from renewable energies (hydraulic, solar, etc.) or nuclear (not so green for earth, but not producing CO_2) the carbon footprint of the datacenter is low. But if energy is becoming from coal, or fuel the CO_2 emissions are high.

information technology equipment: equipment such as computers, servers, mainframes, calculators and all storage devices as arrays, libraries, tape robots together with routers and switches within the local area networks

IT equipment power: total power needed for operate servers, racks, disk arrays, libraries, network telecommunications equipment (such as routers and switches), equipment used for monitoring the data centre (PC, laptops, terminals and workstations) and network telecommunications-specific equipment (such as DSLAM and BTS)

logical consolidation ratio: number of application instances per operating system image

logical server: one single instance of operating system

mainframe: high-performance computer used for large-scale computing purposes that require greater availability and security than a smaller-scale machine can offer

network telecommunications equipment: equipment providing direct connection to core and/or access networks including switches, DSLAM, BTS

operator site: premises accommodating network telecommunications equipment providing direct connection to the core and access networks and which may also accommodate information technology equipment

physical server: box containing supplies for energy, mother board, central processing unit, memory, slots

Power Usage Effectiveness (PUE): metric used to determine the energy efficiency of a data centre that is determined by "Total facility power" divided by "IT equipment power", expressed as a ratio (PUE is expressed as a ratio, with overall efficiency improving as the quotient decreases toward 1)

Recovery Point Objective (RPO): maximum allowed data loss

Recovery Time Objective (RTO): maximum authorized time during application or service can be stopped

server: computer program that provides services to other computer programs (and their users) in the same or other computers

total computing load: total computing power in the data centre, that can be evaluated by taking vendors specifications of computational power of each model of server multiplied by the number of servers (transactions per minute is one measure of total computing power)

total facility power: total power used by all power delivery components (such as uninterruptible power supplies, switches, power distribution units, batteries and transformers), cooling system components (such as chillers, computer room air conditioning units, pumps, fans, engines) and the non-technical energy (such as building lighting)

TPC Benchmark C (TPC-C): On-Line Transaction Processing (OLTP) benchmark measured in transactions per minute (TPMc)

NOTE 1: TPC-C is more complex than previous OLTP benchmarks such as TPC-A because of its multiple transaction types, more complex database and overall execution structure. TPC-C involves a mix of five concurrent transactions of different types and complexity either executed on-line or queued for deferred execution. The database comprises nine types of tables with a wide range of record and population sizes.

NOTE 2: TPC-C simulates a complete computing environment where a population of users executes transactions against a database. The benchmark is centred around the principal activities (transactions) of an order-entry environment. These transactions include entering and delivering orders, recording payments, checking the status of orders, and monitoring the level of stock at the warehouses. While the benchmark portrays the activity of a wholesale supplier, TPC-C is not limited to the activity of any particular business segment, but, rather represents any industry that manages, sells, or distributes a product or service.

utility computing: service provisioning model in which a service provider makes computing resources and infrastructure management available to the customer needs

NOTE: Like other types of "on-demand computing" (such as grid computing), the utility model seeks to maximize the efficient use of resources and/or minimize associated costs. This approach is becoming increasingly common in enterprise computing and is sometimes used for the consumer market as well, for internet service, web-site access, file sharing, and other applications.

Virtual Machine (VM): emulation of a physical server on a shared infrastructure

NOTE: Virtual machine embeds Operating System, specific softwares and application.

virtual server: "piece" of physical server dedicated to run a "virtual machine

virtualization: software that separates applications from the physical hardware on which they run, allowing a "piece" of physical server to support one application, instead of requiring a full server

virtualization ratio: number of Virtual Machines per server

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AC	Alternative Curent
ADSL	Asymetric Digital Suscriber Line
AS	Application Server
ATTM	Access Transmission Terminal and Multiplexing
B2B	Business To Business
B2C	Business To Customer
BTS	Base Transceiver Station
CPU	Central Processing Unit
CRAC	Computer Room Air Conditioning
CRIP	Club des Responsables d'Infrastructure et de Production
DC	Data Centre
DCIE	Data Centre Infrastructure Efficiency
DRP	Disaster Recovery Plan
DSLAM	Digital Subscriber Line Access Multiplexer
HBA	Host Bus Adaptor
HQE	Haute Qualité Energétique
HVDC	High Voltage Direct Current
ICT	Information Communication Technology
IEC	International Electrotechnical Commission
IO	Input Output
IS	Information Systems
ISP	Internet Service Provider
IT	Information Technology
ITIL	IT Information Library

KPI	Key Performance Indicator
LEED	Leadership in Energy and Environmental Design
M2M	Machine To Machine
MVS	Proprietary Operating System for IBM Mainframes servers
NGDC	New Generation Data Centre
NGN	Next Generation Network
OLTP	On-Line Transaction Processing
OS	Operating System
PDU	Power Distribution Unit
POD	Performance Optimzed Datacenter
PUE	Power Usage Effectiveness
RISC	Reduced Instruction Set Computer
RPO	Recovery Point Objective
RTO	Recovery Time Objective
SAN	Storage Area Network
SLA	Service Level Agreement
TCO	Total Cost of Ownership
TPC-C	TPC Benchmark C
TPM	Transaction Per Minute
TPMc	transaction per minute - count
TV	TeleVision
UPS	Uninterruptible Power Supply
VM	Virtual Machine
VMS	Proprietary Operating System for DEC Mainframe Servers
VOD	Video On Demand
VOIP	Voice Over IP
WAS	Web Access Server

4 Overview of data centres

4.1 Types of data centres

There are a number of different types of data centre:

- **a network data centre** has the primary purpose of the delivery and management of broadband services to the operator's customers. To enable their functionality, all network data centres must be connected to at least one core network operator site. For reasons of network resilience, data centres will invariably be connected to more than one operator site and to several other data centres. Data Centres may serve core networks operated by several network operators, thus enabling traffic between customers of different network operators;

- **an enterprise data centre** has similar functions and connectivity functions and connectivity to that of a network data centre but has the primary purpose of the delivery and management of services to its employees and customers;

- **a co-location data centre** is one in which multiple customers locate their own network, server and storage equipment and have the ability to interconnect to a variety of telecommunications and other network service providers. The support infrastructure of the building (such as power distribution, security, environmental control and housekeeping services) is provided as a service by the data centre operator;

- **a co-hosting data centre** is one in which multiple customers are provided with access to network, server and storage equipment on which they operate their own services/applications and have the ability to interconnect to a variety of telecommunications and other network service providers. Both the information technology equipment and the support infrastructure of the building are provided as a service by the data centre operator.

This clause will identify and explain the elements of the network sub-systems employed in broadband deployment.

4.2 Tiering of data centres

4.2.1 Tiers and criticality

Several levels of data centres have been defined, based on the criticality of the applications or the business processed which determine the global Recovery Time Objective (RTO). The lower the RTO, the more the data centre has to be supported by the use of redundant equipment in both the technical environment and IT infrastructure domains.

A number of schemes defining levels of data centres have been developed that are considered in the following clauses.

4.2.2 ANSI/TIA-942

ANSI/TIA-942 [1] defines requirements for reliability and availability of data centres, including the associated redundant support infrastructures, based on four "tiers". Network data centres are assumed to at least meet the requirements of Tier 3.

4.2.3 Uptime Institute

The Uptime Institute [2] defines an alternative system of "Tiers" based upon business objectives and acceptable downtime as shown in table 1. The Tier determines the redundancy of energy and cooling equipment as indicated in table 1 and shown in figure 2 and has some significant consequences on energy costs.

Table 1: Uptime Institute Tiers

Tier	Impact of failure	Design criteria	Downtime (maximum)
1	Internal company impact Mostly cash-based Limited on-line presence Low dependence on IT Downtime perceived as tolerable inconvenience	Single path for power and cooling distribution No redundant components	28,8 hours/year
2	Business critical applications Multiple servers Telephone system vital to business Dependent on e-mail	Single path for power and cooling distribution Redundant components	22,0 hours/year
3	World-wide presence Majority of revenue from on-line business VoIP telephone system High dependence on IT High cost of downtime	Multiple power and cooling distribution paths but only one path active Redundant components; concurrently maintainable	1,6 hours/year
4	Strategic or mission critical business Majority of revenue from electronic transactions Business model entirely dependent on IT	Multiple active power and cooling distribution paths; redundant components; fault	0,4 hours/year

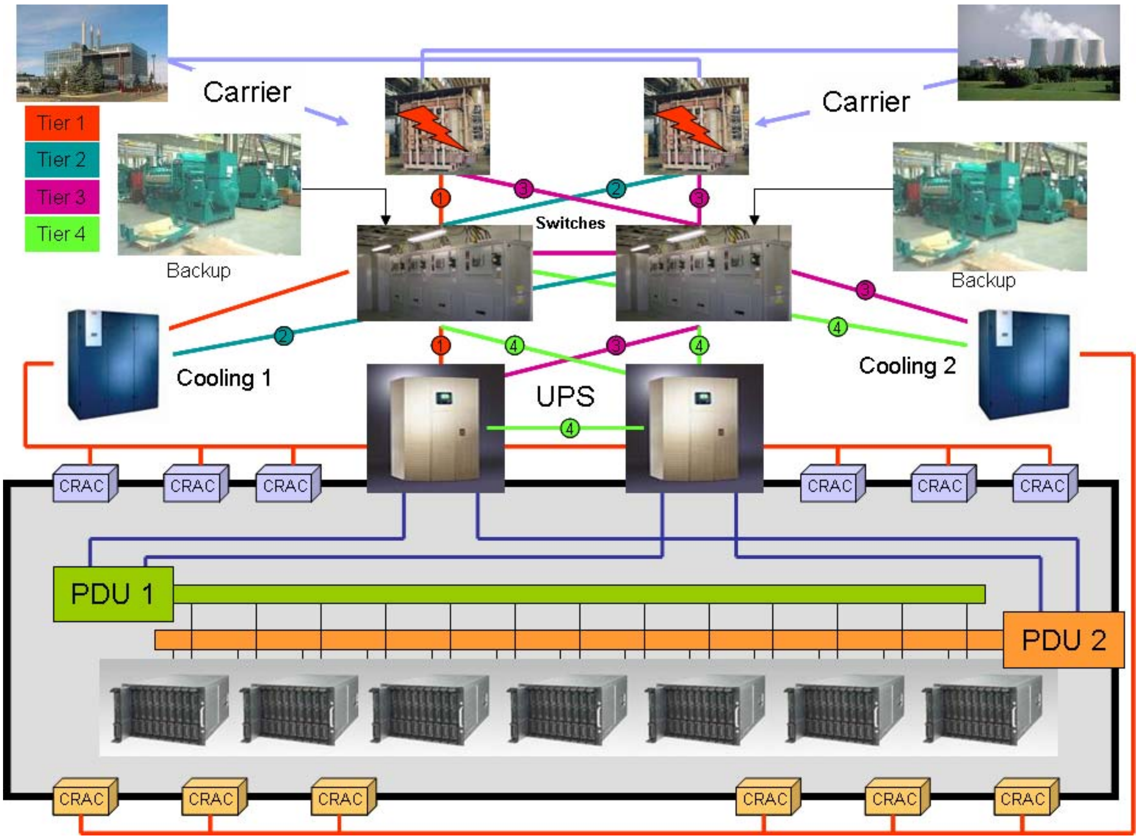

Figure 2: Uptime Institute Tier energy paths and redundancy scheme

4.2.4 Criticality levels

These levels as shown in table 2 are proposed by the Siska Hennessey Group and offers 10 levels of criticality, from 98 % estimated availability (175,2 hours of downtime/year) for a C1 level to C10 level (99,99999999 %, which corresponds to 0,0031 s of annual downtime). Levels above C4 are not considered to be achievable with available technologies.

For the purpose of comparison, the Uptime Institute Tiers of clause 4.2.3 lie between C2 and C4 criticality levels as proposed by the Siska Hennessey Group.

Table 2: Criticality levels proposed by Siska Hennessey Group

Tier	% availability	Annual downtime	Uptime Institute Tier	Status
C1	98	175,2 hours		Achievable
C2	99	87,6 hours	Tiers 1-2	Achievable
C3	99,9	8,76 hours	Tiers 3-3+	Achievable
C4	99,99	53 minutes	Tier 4	Achievable
C5	99,999	5,3 minutes		Not achievable (see note)
C6	99,9999	31 seconds		Not achievable (see note)
C7	99,99999	3,1 seconds		Not achievable (see note)
C8	99,999999	0,31 seconds		Not achievable (see note)
C9	99,9999999	0,031 seconds		Not achievable (see note)
C10	99.99999999	0,0031 seconds		Not achievable (see note)
NOTE: Not considered to be achievable with current technologies.				

4.2.5 Tiers and costs

The capital expenditure (Capex) and operational expenditure (Opex) of new data centres increase with the tier level. Figure 3 shows Capex and Opex (normalized to 100 for the Uptime Institute Tier 1) as a function of Uptime Institute tier.

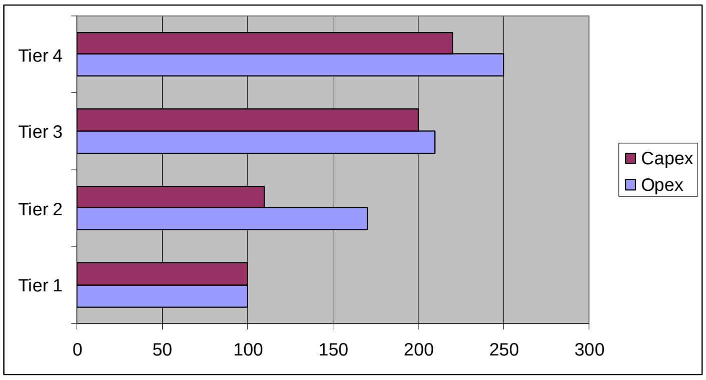

Figure 3: Uptime Institute Tiers and expenditure

Capex, which includes building, design, facilities (such energy, cooling and fire detection) shows the most significant increase between Tier 2 and Tier 3 due to creation of a fully redundant, concurrently maintainable facility including the power and cooling infrastructure.

The Uptime Institute assume that power requirement density (W/m^2) increases with each Tier level as shown in table 3 (normalized to 100 for the Uptime Institute Tier 1). Consequently, the Opex related to energy consumption would increase accordingly with some adjustment (for example, an additional 10 % to 20 %) for the inefficiencies in the redundant power and cooling equipment.

Table 3: Uptime Institute Tiers and power density ratios

Tier	Power density ratio
1	100
2	170
3	200
4	250

Opex also increases significantly between Tier 2 and 3 since the power distribution and cooling infrastructure runs at less than 50 % utilization, allowing for a system component failure without impacting service. There will be some additional Opex costs associated with the inefficiencies in the redundant power distribution and cooling plant, but the primary determinant of Opex cost will be the amount and density of the information technology deployed in the data centre.

4.3 Issues faced by data centres

4.3.1 General

Clause 4.3 reviews the situation in existing data centres and the issues that all enterprises, including network operators, are facing now or will face in the near future.

There are several types of data centre, from main strategic buildings, running enterprise, "mission critical" applications, for which maximum security and guarantee of service continuity is mandatory, to technical sites or computer rooms, for which the same level of integrity is not required (see table 4). This will have a direct and significant consequence on the energy costs, due to redundancy of the technical environmental and the IT infrastructure.

Table 4: Uptime Institute Tiers and mission criticality

Data Centre Type	Hosted applications	Disaster Recovery Plan (DRP)	Uptime Institute Tier
Strategic	Mission critical applications	Campus dual-site	3+/ 4
Secondary	Business critical/internal impact	Remote site	3
Local	Proximity equipment or controller	Equipment redundancy	1/2

4.3.2 Current issues

4.3.2.1 Overview

The information in table 5 summarizes the results obtained from an internal benchmarking exercise undertaken by the enterprise members of Club des Responsables d'Infrastructure et de Production (CRIP).

NOTE: CRIP is a French organization representing major companies such as in the banking, telecommunications, insurance, car manufacturing and general industrial sectors.

Table 5: Average values of actual situation (CRIP source)

Aspect of data centre design	No. of servers	
	100-200	**> 200**
Floor space for IT equipment (m^2)	139	2 405
Average total floor space of the data centre (m^2)	2 987	4 538
Average power consumption (W/m^2)	86	655
Autonomy following total failure of external electrical supply (days)	10	
	(% of those surveyed)	
Redundancy of electrical systems	67	94
Redundancy of cooling systems	91	91
Redundancy of telecommunications networks and rooms	58	90
Existence of local disaster recovery plan (DRP)	46	
Existence of campus mode DRP (dual sites within 10 km)	22	
Existence of metropolitan DRP (sites within 10 km to 100 km)	26	
Existence of continental DRP (sites separated by more than 100 km)	18	
Existence of effective DRP	51	
Supervision and control room backup	40	

4.3.2.2 Principle issues

An unsatisfactory situation exists in legacy data centres as a result of historical policies for the provision of servers, often with each application having its own dedicated physical server, sometimes as a result of running older operating systems not allowing virtualization features. As a result, these servers may have a Central Processing Unit (CPU) usage of only 10 % to 20 %, resulting in a very low energy efficiency. One of the reasons for this has been the lack of effective management of server capacity, another consequence of which is that servers may not be removed from service when they are no longer required. The overall result of this is that many data centres are now at their limits in terms of energy, cooling and floor space.

Research surveying benchmarked significant enterprises [3] indicates that the most significant concerns in existing data centres are:

- lack of energy;
- lack of cooling;
- absence of upgrade path due to new environmental legislation and other constraints;
- energy costs;
- new generation hardware, more efficient but creating areas with high energy density (from 0,7 kW/m^2 to 20 kW and potentially 40 kW per rack in 2010);
- average PUE (see clause 5.3.1) of 2,5 to 2,8;
- low usage of servers, especially Windows and Linux (10 % to 20 %).

Nevertheless, unless improvement in energy efficiency are implemented it is clear that energy costs, as a proportion of the Total Cost of Ownership (TCO), which increase from a typical level of 10 % to 50 %, meaning that energy costs (Opex) will exceed annual IT Capex.

4.3.2.3 Operator data centres

Historically, network data centres have often migrated into existing operator sites which are typically located in urban areas. The primary power supply to these locations was often not designed for the high levels of energy usage required by the technology now employed.

These buildings and their infrastructure were designed to accommodate network telecommunications equipment which had a power usage density several orders of magnitude lower than the modern information technology equipment that has replaced it. Modern building technology is capable of achieving far greater efficiency both in floor space utilization and energy usage; hence it is unlikely that the overall performance of legacy buildings could ever be made to approach that of purpose-built data centre complexes. It is, therefore, probably necessary to consider these as separate cases when comparing energy performance.

Additionally, these existing buildings often have a shortage of floor space that is difficult to increase due to commercial, building and planning constraints in urban areas; this, in turn, forces increased concentrations of processing capability. Legislative and environmental factors place severe constraints on the provision of the additional cooling equipment that becomes necessary.

As energy costs continue to rise and concerns regarding its availability increase, it will become even more necessary to employ new generation hardware with greater processing efficiency.

These factors require new strategies and practices to be employed in the design and operation of data centres, particularly in relation to energy efficiency.

4.3.3 Evolution and future trends

4.3.3.1 Power and cooling demands of IT equipment

The technology road-maps of the main IT equipment (servers and/or storage) vendors road-maps show that predicted power consumption values (in terms of kW/m^2) are not aligned with the capabilities of the majority of computer rooms within data centre facilities (both in terms of power provisioning and cooling capacity).

NOTE 1: The great majority of current data centres are not adapted to meet the energy needs required to host these new technologies (Source CRIP Members internal benchmark on Data centre trends - February 2008)

New technologies of servers (such as blades, chassis and Unix mainframes) have power requirements that exceed the capacity of computer rooms. Examples of such increases are from 0,5 kW/m^2 to 0,7 kW/m^2 to 4 kW/m^2 or 6 kW/m^2 in the two next years, and much more (10 kW/m^2) in the five next years.

NOTE 2: Gartner Group said in 2006: "50 % of existing data centres are becoming obsolete by the end of 2008 in terms of space floor, energy and cooling potential".

Such increases make the concept of kW/m^2 as a design criterion irrelevant and introduce a new evaluation method based on kW into a rack, forcing designers to consider high density areas within data centres with the necessary electrical power and cooling systems. These areas may be specially configured areas within a traditional computer room (see figure 17 and figure 18) allowing the remaining space to accommodate low-density equipment such as robotics, backup libraries, network switches, etc.

Alternatively, the required conditions may be created within closed and secured pre-racked environments (including electrical power management features, dedicated liquid cooling systems and containing all the necessary connections to external networks) which allows their use outside a computer room (see clause 7.2.2.1.5).

4.3.3.2 New projects

Some projects or services, especially in the telecommunications and Internet Service Provider (ISP) world, need huge computing power, linked to the natural growth of their activities or global consolidation initiatives. Additionally, new services (VOD, VOIP, TV ADSL, etc.), electronic exchanges with customers (B2C), suppliers or partners (B2B) and, in the near future, between machines (M2M) will impose new constraints in terms of connectivity, availability, security and random workload absorption.

An overall approach has to be taken to IT in the data centre, enabling sharing policies and being able to respond immediately to application needs, with the required level of service, in a cost effective manner.

4.3.3.3 Data centre consolidation programmes

A recent trend is the consolidation of the existing data centres of major organizations within fewer, major "Critical Data Centres". This has a potential impact on energy costs since "Critical Data Centres" are usually constructed at the Tier 3+ or Tier 4 level. However it is possible to have some areas or computer rooms of a low Tier level (1 - 2) in a data centre of a higher tier (3 - 4). If the data centre is of Tier 3 or 4, overall redundancy of power supplies will be implemented, but only the Tier 3 or 4 computer rooms will require fully redundant power and cooling systems and other equipment (applications) may be installed in computer rooms with a lower Tier level (as shown in figure 4).

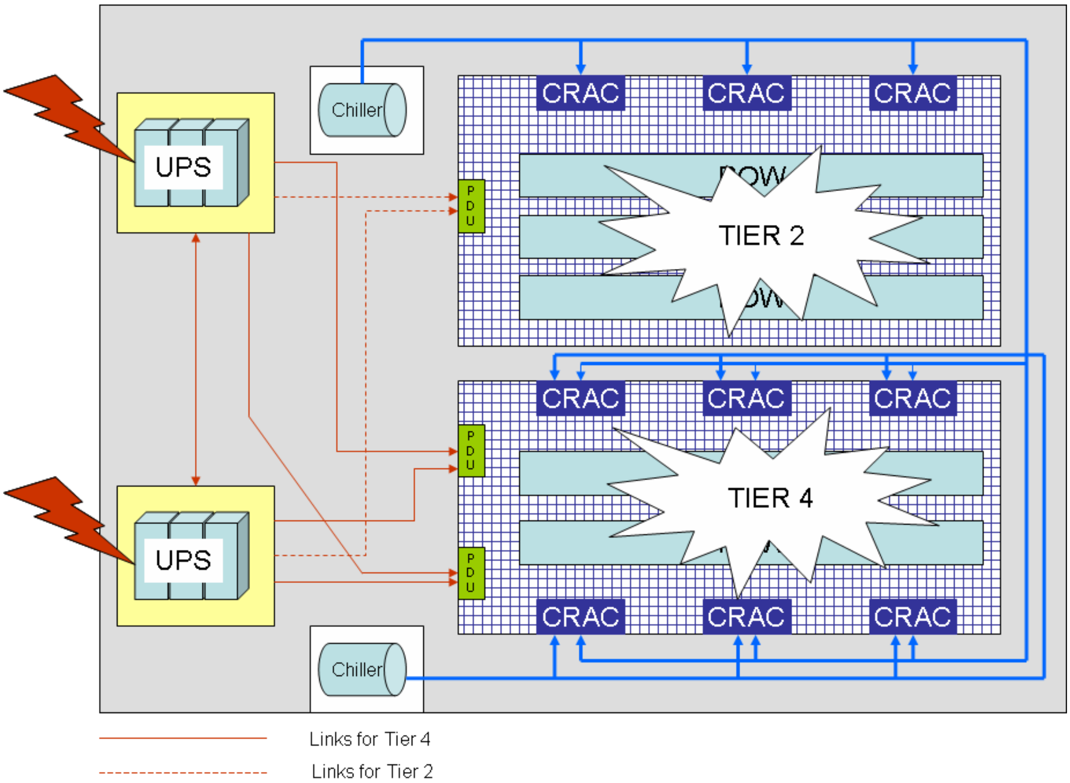

Figure 4: Mixed Tier-level data centre

This approach may reduce overall power consumption. In addition it may be possible to operate the different computer rooms at different temperatures (for example, information technology equipment may be segregated from network telecommunication equipment) allowing further reductions in power consumption (see clause 7.2.3).

4.3.3.4 Environmental impacts

Regulations frequently require a prolonged case-by-case study by authorities which can introduce significant delays to the planning process. These constraints make it very difficult if not impossible to implement projects for the expansion of an existing data centre in towns which would increase heat dissipation or have significant energy requirements.

4.4 The new context

4.4.1 Energy consumption and energy efficiency

The model "one application, one physical architecture", each with its own servers, storage, is becoming obsolete and is being replaced by a new model based on shared IT components and mutualisation of technical infrastructure. In the near future, the data centre will be a true "production factory", with automation and industrialization and will become fully virtualized as described in clause 6.

New "energy efficient "equipment in all domains of IT, new cooling technologies (including solutions that only become practical as the power densities increase) and other solutions to manage power consumption more efficiently will all act to enable a reduction in energy costs for a given level of service. However, the greatest savings will only be obtained if other initiatives reflecting the "utility computing" concept are fully implemented (such as consolidation, automated provisioning of servers (network or storage) and recurrent operations).

The search for energy efficiency has to cover all disciplines and not only focus on the technical environment.

4.4.2 Factors impacting energy efficiency

The factors shown in table 6 contribute to poor energy efficiency and, consequently, high energy consumption.

Table 6: Principle factors leading to poor energy efficiency

Power distribution systems	Power distribution units and/or transformers operating well below their full load capacities.
	N+1 or 2N redundant designs, which result in under utilization of components.
	Decreased efficiency of Uninterruptible Power Supply (UPS) equipment when run at low loads.
Cooling systems	Air conditioners forced to consume extra power to drive air at high pressures over long distances.
	Pumps with flow rate automatically adjusted by valves (which reduces the pump efficiency).
	N+1 or 2N redundant designs, which result in under utilization of components.
Physical infrastructure	Under-floor "noodleware" that contribute to inefficiency by forcing cooling devices to work harder to accommodate existing load heat removal requirements, which can lead to temperature differences and high-heat load areas might receive inadequate cooling.
	Lack of internal computer room design.
IT infrastructure	Low usage of existing servers (10 % to 20 % CPU) specially in X86 world.
	Lack of capacity management process for technical environment and IT.
	Physical servers dedicated to applications.
	No consolidation, lack of sharing policy.
	Data redundancy, generating lot of storage capacity.
	IT equipment in active mode 24/24, 7/7 but only used at certain hours of the day and/or on certain days.
	Old generations of servers, with a low computing power/electrical consumption ratio.
	Lack of functional cartography, generating a lot of applications, data duplications, backups, and bit rate for exchanges, that contribute to dramatically increase the TCO (number of physical components, software fees, people, space floor, etc.) and energy consumption.

The actual situation is more generally the consequence of an important growth of the needs in term of IT equipment, due to the natural growth of the enterprises business and the creation of new services. Another factor increasing the IT load is that the functional cartography is very complex and generates a lot of applications.

Sometimes, the best way to make the maximum savings on energy is to decrease the number of applications and minimize data duplications.

4.4.3 On-going initiatives

Many organizations have introduced initiatives in response to the principle concerns identified within clause 4.3.2.2. These cover the following areas:

- IT infrastructure (discussed in detail in clause 6);
  - consolidation of their existing assets to decrease the number of physical components in computer rooms comprising:
    - storage consolidation - concentrating data on to shared arrays instead of using storage capacity dedicated to specific applications. This requires a development of a Storage Area Network (SAN) policy;

  - server consolidation initiatives are now in progress in many enterprises, particularly for servers using X86 technology, running Windows or Linux operating systems, due to their being more numerous than those using other operating systems (Unix, MVS, VMS). These servers are those from the Tier 1 (presentation) and Tier 2 (application servers) in the commonly used three tier software architecture model (outlined in clause 6.2). The consolidation of these servers is primarily implemented with:
    - new generation hardware such as blade or other racked servers to build a shared and common technical infrastructure;
    - virtualization software to move an application on to a "virtual server". Several applications can be consolidated in the same hardware, using virtual servers, with a 60 % to 80 % computational load, instead of only 10 % to 20 % for a non-virtualized server.

  These actions not only limit a dramatic increase of energy needs and costs but also reduce the time before the new generation data centres become operational. In the majority of cases, this consolidation is "physical" (with and without virtualization) using mainframes, racked servers, "pizza box" or "blade servers" technologies, due to their power capacity / space floor place ratio.

  - process automation for more precise provisioning of resources, introduction of new virtual servers and dynamically managing their workload and scalability;
  - capacity management for storage and servers, to have a better usage of existing resources;

- cooling systems (discussed in detail in clause 7):
  - increasing the ambient temperature;
  - free air cooling;
  - water cooling;
  - disabling air-humidifiers within areas that only contain equipment meeting the requirements of EN 300 019-1-3 [11];
- power distribution systems (discussed in detail in clause 9);
  - High Voltage Direct Current (HVDC) supplies for IT equipment (for new data centres).

5 Energy efficiency standards and metrics

5.1 Review of activities outside ETSI

5.1.1 EU Code of Conduct on Data Centres Energy Efficiency

The European Union Code of Conduct [4] provides the opportunity for operators of data centres to implement practices intended to reduce the energy consumption of their data centres. Details of the energy efficiency best practices [5] that are employed are detailed together with actual recorded energy consumption measured at specific points in the data centre.

5.2 Energy consumption in data centres

Of the total energy used in a data centre, the principal areas of consumption, shown schematically in figure 5, are:

- power distribution to information technology equipment and network telecommunications equipment in the computer room;
- environmental control (for example, cooling and humidity) applied to the computer room;
- lighting and equipment in offices associated with the data centre;

- lighting for the computer room.

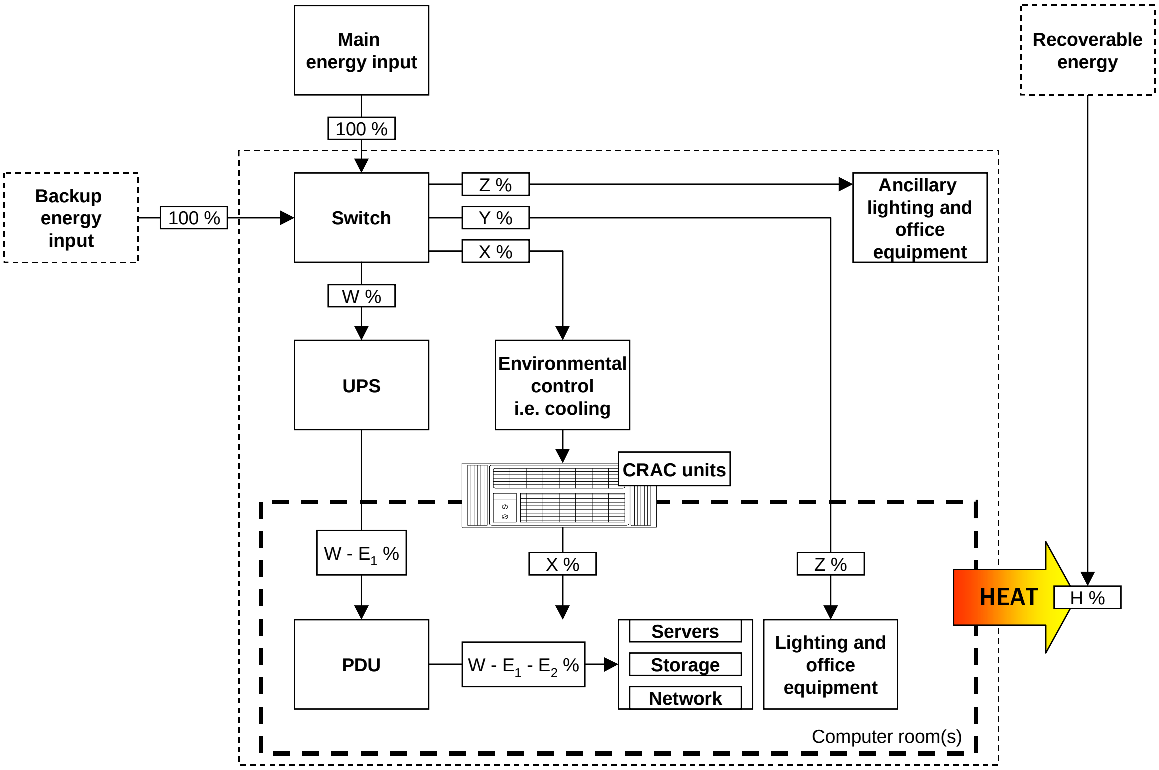

Figure 5: Schematic of energy consumption and wastage in data centres

With reference to figure 5,

- the proportion of the energy delivered to the information technology equipment is W - E1 - E2 % where:

W (%) = energy consumption at the input to the UPS;

E_1 (%) = energy wasted within/by of the UPS;

E_2 (%) = energy wasted within/by of the PDU.

- the energy consumed by the environmental control equipment X %.

Unless otherwise stated, the standard model used in the present document is follows:

- the % power required for IT equipment 45 (based on W = 60 and E1 + E2 = 15 in figure 5);
- the % power required for the cooling systems = 37 (that is X = 37 in figure 5);
- the % power required for building facilities = 3 (that is Y + Z = 3 in figure 5).

5.3 Energy Efficiency Key Performance Indicators (KPIs)

5.3.1 Power Usage Effectiveness (PUE) or Data Centre Infrastructure Efficiency (DCIE)

5.3.1.1 General

With reference to figure 5, PUE is defined as $100/(W - E_1 - E_2)$.

The European Union Code of Conduct [4] recognizes the DCIE (the inverse of the PUE expressed as a percentage as shown in table 7.

Table 7: PUE and DCIE conversion

PUE	DCIE
4,0	25 %
3,0	33,3 %
2,5	40 %
2,2	45 %
2,0	50 %
1,8	56 %
1,6	62,5 %
1,5	66,7 %
1,4	71,4 %
1,3	76,9 %

The standard model of clause 5.2 produces a PUE of 2.22 and a DCIE of 45 %.

Figure 6 indicates the PUE and DCIE figures for high efficiency, medium efficiency and low efficiency data centres.

In general, data centres built in the last five years have an operational PUE of between 2 and 2,5 (that is DCIE values of 40 % - 50 %). Some older data centres have operational PUE values greater than 3 (DCIE values less than 33,3 %).

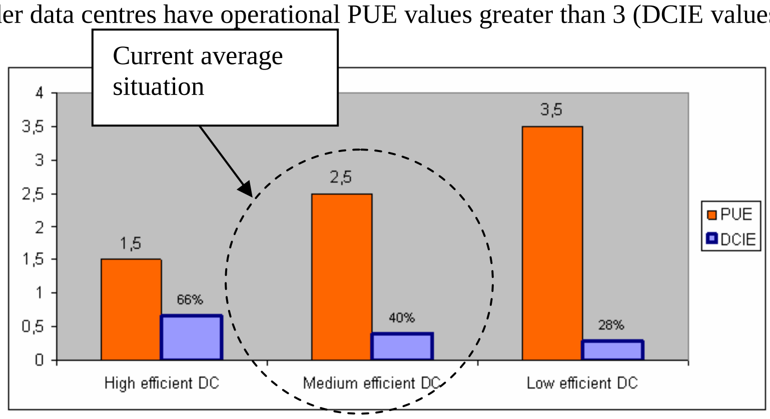

NOTE: The word "efficient" needs to be replaced with the word "efficiency".

Figure 6: Data centre energy efficiency spectrum and PUE/DCIE values

The latest Tier 3+ or 4 data centres may have PUE targets of as low as 1,4. Other data centres of Tier 2 that use free water or air cooling and equipped with the latest generations of UPS, pods (see clause 7.2.2.1.5), high-efficiency racks and servers are capable of localized PUE targets of as low as 1,2 in specific areas.

5.3.1.2 PUE for new data centres

For a data centre under construction, the KPIs of PUE (or DCIE) are appropriate (see clause 12.2).

A data centre under construction can be designed to have a specified PUE (or DCIE) and following construction the actual PUE (or DCIE) can be monitored against the design value. However, it should be recognized that any reduction of W (by means of improvement of the information technology or network telecommunications equipment or its usage, and by reduction of waste in the power distribution system, E_1 or E_2) without a equivalent reduction in the primary energy consumption parameters, X , Y and Z, will lead to an increase in PUE indicating a worsening in energy efficiency rather recognising a reduction in energy usage.

The design strategies for new buildings can be chosen in order to maximize their initial PUE and with the objective of maintaining that PUE during their subsequent growth, operation and evolution

As described in clause 7, the selection of data centre location, building engineering strategies and system selection can substantially reduce the energy consumption required for environmental control. As described in clause 8, effective planning of pathways and spaces in accordance with EN 50174-1 [8] and the data centre specific aspects of EN 50174-2 [9] can maximize the energy efficiency of environmental control systems.

As described in clause 6, internal design studies can influence the usage of equipment (servers, storage and networking). In addition, appropriate management tools used in conjunction with effective process organization and automation can significantly reduce the overall energy consumption of the information technology equipment.

The introduction of high efficiency power distribution components as described in clause 9 can reduce the value of W.

Using these approaches in combination, the typical design PUE of new buildings will reduce over the next few years from the current value of 2,0 to 1,6 (2010-2015) and, with the implementation of measures outlined in clause 14, can be expected to reach 1,3.

It is possible to specify even lower PUE values for dedicated areas within computer rooms but it is unlikely that, using the approaches detailed in the present document, the PUE values for the overall data centre will be much lower than the values shown. Table 8 gives several steps for a better efficiency. Some of them are directly achievable by the data centre owner, or the IT one, but some others are depending on technological evolutions.

Table 8: Progressive improvement in PUE

Year	Situation	PUE
2005-2008	Many stand alone servers Physical dedicated architecture per application No virtualization on Open systems Obsolete technologies and OS Manual provisioning of resources Lack of space, energy or cooling in computer rooms	2,5 - 2,0
2008-2010	Infrastructure policy (server farms, Unix mainframes) Storage consolidation - Virtualization - Tiering policy Physical server consolidation Virtualization technologies for servers and storage Automatic provisioning for server Operation process automation Capacity management for server and storage Need 2 kW/m^2 to 4 kW/m^2	1,8
2010-2015	Energy Efficient Data Centre building (LEED) Own production of renewable Energy (wind, solar, hydraulic, etc.) New generation cooling appliances - Free cooling Auto-cooled IT hardware Energy Efficient servers (blade farms, mainframes) Specific software for manage Energy capacity Software for thermal measurement and modelling in computer rooms Massive application consolidations - IS rationalization Full data centre processes automation	1,6

5.3.1.3 PUE in existing data centres

There is little opportunity within existing buildings to reduce energy consumption by modifications to building structures or to make substantial changes to the environmental control systems without incurring significant costs and operational disruption.

Small improvements may achieved by limited changes to environmental control systems and general improvements in energy efficiency in power distribution components and ancillary areas thereby reducing the values of E_1, E_2 and W in figure 4.

Instead, reduction of overall energy consumption may be achieved by, where possible, implementing effective planning of pathways and spaces in accordance with EN 50174-1 [8] and the data centre specific aspects of EN 50174-2 [9] in order to maximize the energy efficiency of environmental control systems.

Procurement of more energy efficient IT equipment and consolidation/virtualization/process automation initiatives represents the only realistic solution for a reduction in energy consumption in existing buildings for a given level of service.

However, these actions without a equivalent reduction in the primary energy consumption parameters, X, Y and Z as shown in figure 5, will lead to an increase in PUE indicating a worsening in energy efficiency rather than recognising a reduction in energy usage.

PUE is hard to evaluate without a set of tools generally missing in majority of legacy data centres. As mentioned above, some necessary primary actions on infrastructure such as virtualization will have a negative effect on PUE. PUE may be the most relevant indicator to give a global view of efficiency, but it is not the most appropriate KPI metric for measurement the improving energy efficiency in existing data centres.

Conformance to the present document (see clause 12.1) for existing data centres requires the use of a KPI other than PUE.

5.3.2 Other KPIs

5.3.2.1 Energy efficiency KPI

A number of these KPIs exist including:

- electrical power/space-floor ratio expressed in kW/m^2; typically, this KPI is used for legacy data centres and/or low density areas in computer rooms;
- the ratio of total energy consumption / total computational load;
- computing power/electrical power expressed as TPM-C/kW; this KPI gives a density of computing potential per kW and is useful for racks and/ or high density areas in computer rooms;
- total energy required for data centre per hour;
- total energy consumption per year.

5.3.2.2 Consolidation KPI

A number of KPIs exist for IT infrastructure consolidation including:

- number of physical servers;
- number of virtual servers;
- virtualization ratio;
- number of deployed operating system images (logical servers) - one physical server can contain several logical servers: this KPI is a principal measure of logical consolidation;
- logical consolidation ratio;
- number of applications; this is a KPI only for application consolidation and de-commissioning (relation has to be done between application, physical components, and technical environment needs saved with the consolidation);
- average computational load per family (RISC, X86) - this KPI evaluates the efficiency of servers, in terms of computational load with the objective having the highest usage of CPU power in order to minimize the number of servers;
- number of disk arrays.

A number of KPIs exist for physical storage consolidation including:

- number of SAN ports;
- Tbytes per m^2;
- ratio of occupied storage space/m^2 - this KPI gives storage density;

- average array allocation ratio - this KPI is similar to average computational load, providing the ratio of space allocation in the array.

5.3.2.3 Data collection

The production of the KPIs in clause 5.3.2.2 requires the collection and aggregation of a wide range of data as shown in table 9.

Table 9: Date required for the production of consolidation KPIs

Vendors characteristics of all equipment in the data centre	Vendor
	Model, type
	Number of CPU , cores
	Computational power
	Electrical consumption (idle, full load)
	Heat dissipation
	Weight,
	Size
Inventory database from all equipment of the data centre	Per data centre
	Per computer room
	Per business process
	Per application
Database containing measurement values during business activities and during periods of non-activity	
Electrical consumption	Per data centre
	Per computer room
	Per business process
	Per application
	Per end-user
Computational needs	Per computer room
	Per business process
	Per application
	Per user
Cooling needs	Per computer room
	Per business process
	Per application
	Per user

6 Increasing the energy efficiency of IT infrastructures

6.1 General

Indications of the impact of some of the actions in this domain of energy efficiency are shown in annex A.

6.2 3 Tier Software Architecture Model

Figure 7 shows the 3 tier software architecture model adopted in data centre environments. This clause makes reference to the Tiers of this model - and these tiers should not be confused with those discussed in clause 4.

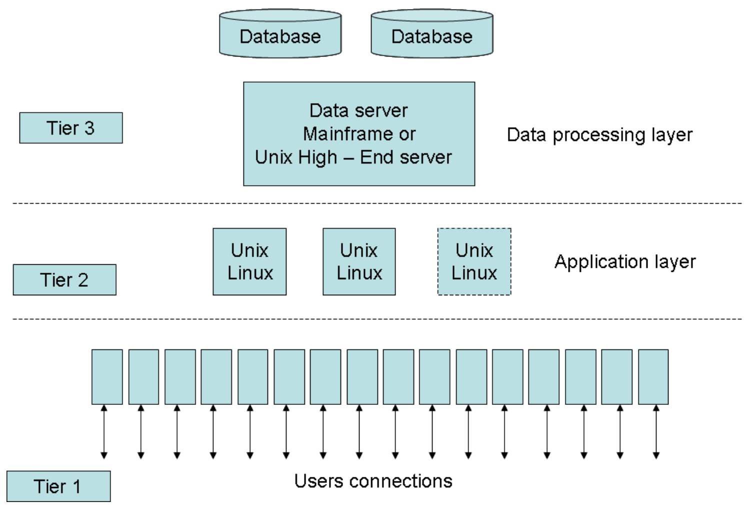

Figure 7: Typical architecture of 3 Tier model application

6.3 Energy efficiency solutions

6.3.1 Obsolete equipment

This involves the identification, turning-off and removal of all equipment without any activity such as old servers, modems and routers. This typically represents a small percentage of the installed equipment (possibly 5 %) but decommissioning of this equipment provides an immediate reduction in energy consumption without any reduction in service levels.

6.3.2 Replacement equipment

This involves the replacement of existing equipment from previous generations of technology with the most recent, more energy efficient IT equipment.

The choice of server equipment should be directed by their ability to run virtualized operating systems.

Blade server farms offer an excellent ratio of power consumption / computing power in a limited space.

High-end or mainframes are and, for the foreseeable future, will be necessary for the processing of large databases.

6.3.3 Power and capacity management

6.3.3.1 General

There are two separate rapid routes by which reductions of power consumption may be achieved by providing more efficient usage of existing resources within existing IT infrastructures without the need for changes to hardware. The routes are described as:

- power management (see clause 6.3.3.2);
- processing capacity management (see clause 6.3.3.3).

6.3.3.2 Power management

6.3.3.2.1 Activation of basic power management features

This involves the activation of any power management features within existing equipment.

The application of dynamic allocation of equipment resources (see clause 6.3.3.3.5) provides additional beneficial effects on power management.

6.3.3.2.2 Activation of "sleep" mode

This involves the activation of sleep mode (that is not a system shut-down of the equipment) during periods without application activity during certain periods during days, weeks or months and can be applied to a variety of equipment depending upon it role within the 3 tier software architecture model of clause 6.2.

It may be even be possible to consider a full system shut-down of certain pieces of equipment.

These solutions are applicable to all servers that are not used in continuous use (such as backup and development servers) and servers in Tiers 1 or 2 of the software architecture model of clause 6.2. Typically, these servers or applications respect a schedule of the type shown in table 10.

Table 10: Server schedules

Times	Activity	Tier 1	Tier 2	Tier 3
08h00 - 20h00	TP	User connections	Application	Database connections
20h00 - 08h00	Backup/batches	Any connection	Inactive	Backup/batch activity
Week-ends	None	Any connection	Inactive	Backup/batch activity

An example of the potential savings is shown below.

EXAMPLE:

An active, last generation X86 mono or bi-processor, server has a typical mean consumption of 240 W. The same server in "sleep" mode has a typical consumption of 80 W and 0 W when turned off. In table 11, 200 such servers were identified that could be turned-off or put in sleep mode:

- for 8 hours per day;
- during weekends and public holidays.

Table 11 shows that the potential energy savings from activating sleep-mode would be 31 % and from shut-down would be 46 %.

Table 11 - Example saving calculation

Times		Activity
No. of week-ends per year	W	52
No. of public holidays per year	P	9
No. of working week-days	D = 365-2W-P	252
Total energy-reduction hours	R = 8D + 24(2W+P)	4 730
Total energy consumption (without action)	8 760 x 240 W Per server	420 480 kWH
Energy reduction with sleep mode	160(8 760-R) per server	128 960 kWH (30,7 %)
Energy reduction with shutdown	240(8 760-R) per server	193 440 kWH (46 %)

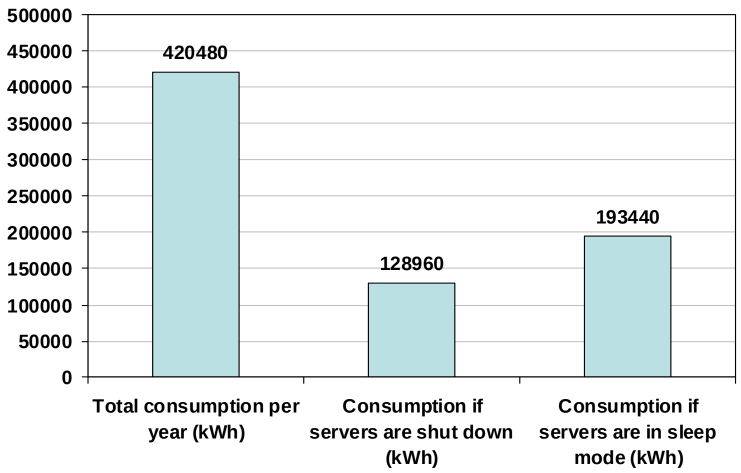

Figure 8: Example savings schematic

6.3.3.2.3 Reduction of energy consumption of environmental control equipment

It may also be possible to reduce the energy consumption for environmental control but the level of savings is depends upon the type of cooling employed. If it is not possible to dynamically adjust the cooling air-rate, any savings would be insignificant. However, if the cooling air-rate can be adjusted dynamically then the energy used to cool the servers could be reduced by up to 50 %.

6.3.3.3 Capacity management

6.3.3.3.1 General

Capacity management is the ongoing, operational, process of estimation and allocation of space, environmental needs, computer hardware, software and connection infrastructure resources to reflect the dynamic nature of data centre users or interactions. As shown in figure 9, capacity management addresses the following:

- is the data centre able to host new applications, services or support the growth of the activity?
- what is the capacity in terms of energy, space and cooling?
- what is the capacity in terms of storage, CPU, memory, I/O, ports, etc.?

Capacity management provides an exhaustive view of the real needs in terms of computational power and/or environmental capabilities by continuous management, measurement and monitoring of the servers and application activities.

The objective of capacity management is to ensure that new capacity is added just in time to meet the anticipated need but not so early that resources go unused for a long period. Successful capacity management, using analytical modelling tools (responding to "what will happen if" scenarios) implements trade-offs between present and future needs that prove to be the most cost-efficient overall. The emergence of new technologies together with changes to business strategies and forecasts change require capacity management to be under continual review.

Effective capacity management supports the use of products that are modular, scalable and also stable and predictable in terms of support and upgrades over the life of the product.

Capacity management has the following objectives:

- prediction and anticipation of future needs of the business due to both natural growth and new projects;
- implementation of actions on IT or environment to provide adequate resources;
- adjustment of infrastructure usage to the real needs of the business and prevent waste due to over-sizing of applications;

- determination of equipment usage;
- preparation of consolidation initiatives (see clause 6.3.4).

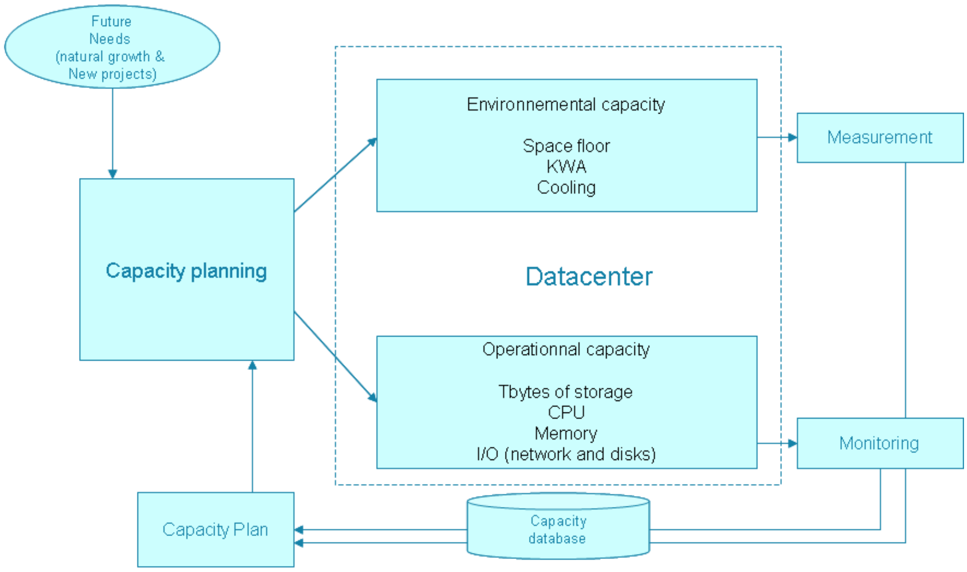

Figure 9: Capacity management

6.3.3.3.2 Environmental capacity management

This requires the measurement and subsequent management of electrical, cooling and space needs. In many cases this information is obtained manually, directly by the data centre personnel. However, the best method is to apply software solutions.

6.3.3.3.3 Storage

This involves the use of shared data storage, active data compression and data de-duplication in order to maximize the utilization of storage capacity. The implementation of thin provisioning for storage, allowing the right disk-space is critical to the management of storage capacity.

6.3.3.3.4 Servers

This involves the use of existing equipment when additional server capacity is required. This approach is a step towards the consolidation initiatives of clause 6.3.4.

6.3.3.3.5 On-demand scalability for on-line business

This requires the implementation of pre-packaged virtual environments, including all logical components necessary to run the application, and a "utility computing" tool to distribute them across the infrastructure taking account of, for example, the number of connections to the service (that is one new server provisioned all the 200 connections). A critical aspect is that the automated system has to be able to remove the additional capacity as soon as the number of connection decreases. Virtualization is the main key for this, as each new server installed is a Virtual Machine.

6.3.4 Consolidation initiatives

6.3.4.1 Consolidation of servers

The consolidation of processing within existing servers is the best way toward reduce energy costs for given level of service. The result of consolidation is a reduction in the number of servers which has a direct impact on the IT infrastructure power requirements which has a corresponding effect on reductions in requirements for cooling and floor space.

There are a number of types of consolidation that are covered in this clause:

- physical consolidation;
- virtualization;
- logical consolidation;
- application or Information Systems (IS) rationalization.

6.3.4.2 Physical consolidation

6.3.4.2.1 The process

Physical consolidation involves the gathering of stand-alone hardware within a physically more powerful container, as shown in figure 10 and can be achieved without using virtualization if the server technology allows partitioning features.

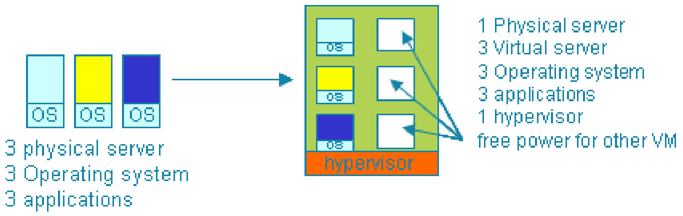

Figure 10: Physical consolidation - Virtualization

6.3.4.2.2 The effects

A physical consolidation programme has the following effects:

- reduction in the number of physical components (servers, storage arrays, robotics);
- savings on floor space, maintenance costs, cooling and power;
- Capex for new hardware.

6.3.4.3 Virtualization

6.3.4.3.1 The process

Virtualization is a method of physical consolidation and has to be implemented made on new generation servers and is a pre-requisite for a shared infrastructure policy.

Physical consolidation has no effect on the number of "logical" servers (that is reducing the number of hardware has no effect on the number of Operating System (OS) images and on software licences).

The effect on the global IT TCO is not significant.

6.3.4.3.2 The effects

Under certain conditions, virtualization can deliver energy reductions of 80 % on a selected set of servers.

6.3.4.3.3 Reduction of energy consumption of IT infrastructure

Table 12 provides a methodology to evaluate energy savings for a virtualized panel of servers. Other indirect savings could be also evaluated if the virtualization affects the cooling requirements in the computer room. In majority of cases, virtualization will not have a positive effect on the PUE of the data centre, but it contributes to a reduction in total energy consumption.

NOTE: Any savings will be ineffective unless the old servers are shut-down electrically.

Table 12: Virtualization savings profile

Number of servers to virtualize	X
Mean energy consumption per server (watt)	W1
Total consumption in Watt (X * W1)	C1
New server number (number of blades or boxes)	X / R
Average consumption per VM (including storage)	W2
Virtualization ratio	R
New consumption (X * W2)	C2
Savings on energy (watt) : C1-C2	S1
Savings on energy % (S1 / C1)	S2

The following figures are provided by a major telecommunications organization that has launched a virtualization project on X86 servers under production. The results are significant, since in addition to virtualization, logical consolidation (see clause 6.3.4.3.4) was also applied, by reducing the number of operating systems. Generally, mean values for energy savings using virtual environments are 40 % to 60 % of the energy consumed by the servers before virtualization.

In this example, the aim was to consolidate many physical servers of one multi-server application (170) into new generation technology servers, associated with a virtualization tool. The legacy servers were 2 CPU Intel X86 "racked" servers from previous generation, containing one image of an operating system and one application instance. The new servers were Intel X86 4 CPU 2 core blade servers, racked. A consolidation ratio of 12 (meaning that one new physical server contains at least 12 virtual servers) was achieved. Figure 11 shows the savings in electrical consumption that resulted.

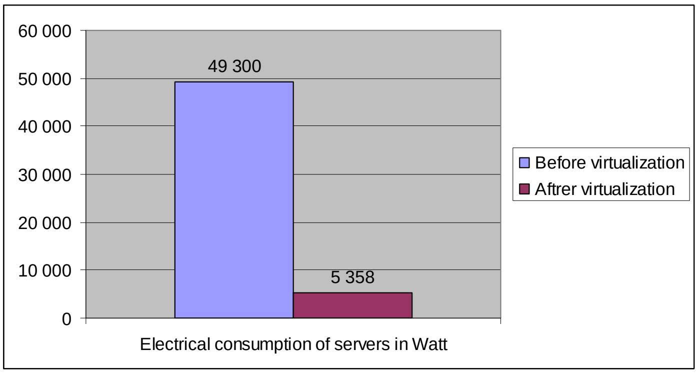

Figure 11: Electrical savings with virtualization

6.3.4.3.4 Reduction of energy consumption of environmental control equipment

See clause 6.3.4.3.4.

6.3.4.4 Logical consolidation

6.3.4.4.1 The process

Logical consolidation aims to decrease the number of logical servers (operating system images) as shown in figure 12 and so make some savings on licence fees in addition to the number of physical servers.

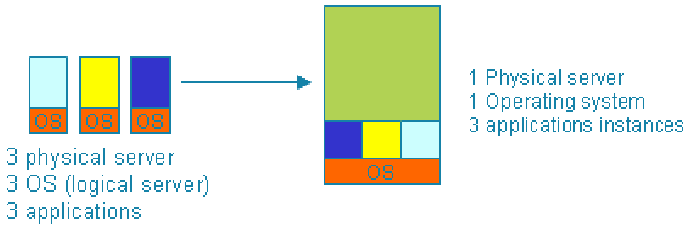

Figure 12: Logical consolidation

This is recommended for existing servers and can, without any Capex provide a better usage of legacy equipment. However, the latest generation of servers are recommended when creating a new data centre. Virtualization and logical consolidation can be implemented in tandem, providing a cumulative benefit.

There are three methods of achieving logical consolidation:

- Homogeneous: only servers containing same logical components of the same application (same web access servers (WAS), web server, etc.) are consolidated in the same OS.
- Semi-homogeneous: only servers containing same logical components, but from different applications (same WAS, web server, etc.) are consolidated in the same OS.
- Heterogeneous: different logical components, from different applications (this method is not the most popular and is implemented only in certain cases, with specific rules).

In all cases, logical consolidation has a pre-requisite of defining consolidation rules and criteria that address, amongst other issues:

- application sensitivity;
- application type;
- daily schedule;
- resource usage measurement (CPU, inputs/outputs (IO), memory, etc.);
- links with other applications.

6.3.4.4.2 The effects

- increased computational load of existing servers;
- reduction in the number of logical servers (OS instances) by strong sharing policy reducing software and maintenance fees;
- improved physical consolidation score, means less servers, and all related effects such as energy, cooling, etc.
- medium effect on global IT TCO, due to savings on licences.

6.3.4.5 Application consolidation

6.3.4.5.1 General

This is the highest level of consolidation initiatives. It addresses the problem, not from the infrastructure layer, but from the business-process layer. As shown in figure 13 it is a "top down" approach, in opposition with the "bottom-up" approach of physical or logical consolidation. This is not the easiest way, but it is the more profitable in terms of savings in all domains.

6.3.4.5.2 The process

There are several ways to undertake application consolidation, but in all cases, the process has to be led by business owners and developers.

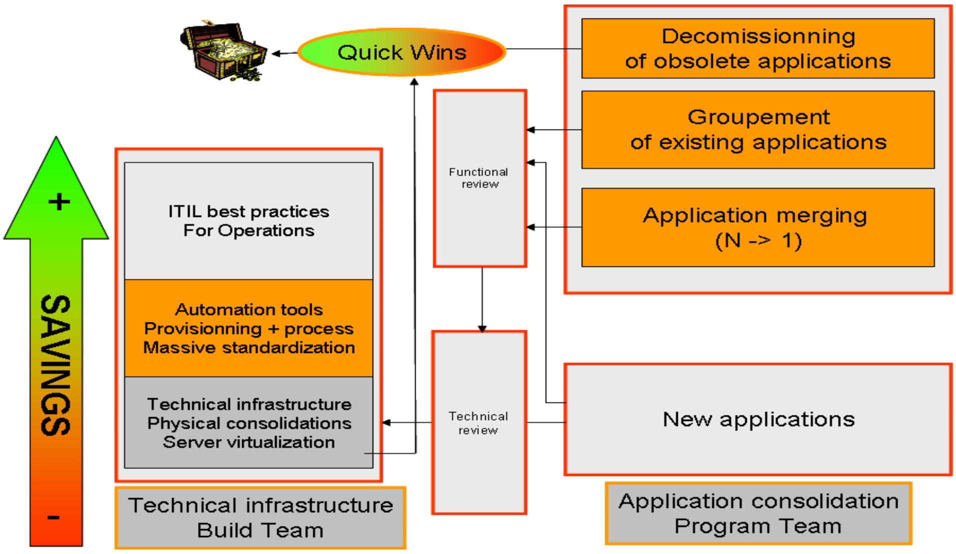

Figure 13: "Top Down" approach for application consolidation

In the majority of cases, application consolidation is not driven only by energy savings. Nevertheless, it is the best way to achieve maximum reduction in energy consumption since it addresses the primary causes of energy consumption inflation, etc. too many applications and too much complexity, redundancy, data duplication, etc.

Such a programme is more commonly implemented to positively impact TCO, improve quality of service, reduce "time to market" and enhance security.

6.3.4.5.3 The effects

- Build business process-oriented or service-oriented dedicated infrastructures, integrating scalability, availability, agility, etc.
- Can be made by consolidation of applications from the same vendor, from the same business owner or the same process on one dedicated infrastructure.
- Boost physical and logical consolidation scores, by decreasing complexity and redundancies.
- Develop consolidated and unique vision of applications relating to same process in order to propose one SLA per process.

7 Reducing the energy demand of environmental control systems

7.1 General

This clause describes the approaches that may be employed to reduce the energy demand of environmental control systems within the data centre which typically represents 35 - 40 % of total energy consumption in an Uptime Institute Tier 3 data centre.

Some of the approaches are applicable to legacy data centres while other is more likely to be applied in new data centres. Further details are provided in clauses 10 and 11.

The present document add to the information provided in TR 102 489 [i.1].

Indications of the impact of some of the actions in this domain of energy efficiency are shown in annex A.

7.2 Energy reduction solutions

7.2.1 Measurement of thermal behaviour

In order to enable reductions in energy usage it is important to determine the thermal patterns within the data centre by using available thermal measurement software tools. These may be deployed without significant Capex since the costs are restricted to the fees for the software and the Opex for the installation and customization.

7.2.2 Improvement of cooling efficiency

7.2.2.1 Zonal approaches to thermal isolation

7.2.2.1.1 General

The actual design in most of legacy data centres is that the Computer Room Air Conditioning (CRAC) units force cold air under a raised floor, into the cabinets, and draw up hot air coming from the top of the cabinets. However, if the hot air and cold air become mixed, bad aeraulic management results producing hot spots within the computer room. This can be shown by taking infra-red photographs of the room. Some simple and low cost actions can be taken to avoid this which generate some non-negligible savings on cooling and have a positive effect on the PUE.

7.2.2.1.2 Hot aisle and cold aisle segregation

This approach creates areas within the data centre that are designed as dedicated "hot aisles" and "cold aisles". Rows of cabinets are created in which the front of the cabinets face each other (cold aisles) and in the adjacent row, the rear of cabinets face each other (hot aisles). This segregation is shown in figure 14.

Cold air is drawn in through the front of the cabinets (cold aisles) and expressed through the rear of the cabinets into the hot aisles.

Infrastructure design, installation and operational procedures are critical to maximizing the benefit of basic hot aisle and cold aisle solutions (see clause 8).

7.2.2.1.3 Cold aisle covers

An improvement of the "hot aisle, cold aisle" concept is provided by placing covers over the cold aisles as shown in figure 15. This approach reduces the loss of cold air losses and is able to demonstrate a rapid reduction in energy usage.

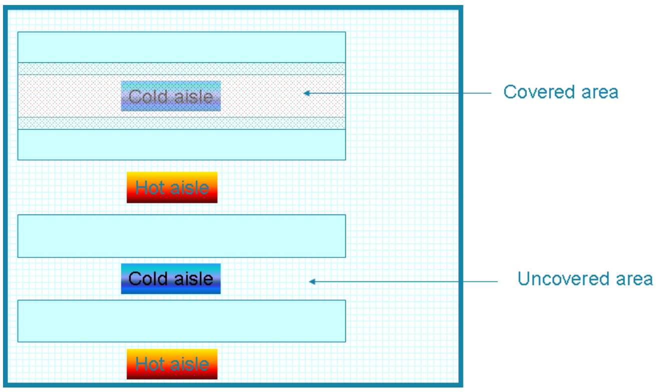

Figure 14: "Segregation of hot aisles and cold aisles

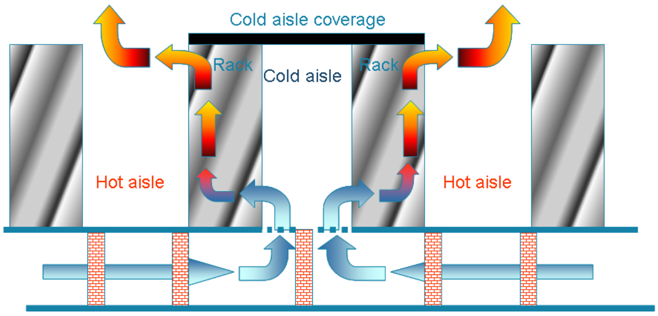

Figure 15: Cold aisle covers

7.2.2.1.4 Segregation using curtains

A low cost, simple and rapidly deployable solution to problems of air-flow employs plastic curtains hanging from the ceiling of the computer room to isolate hot areas from cold areas as shown in figure 16. This approach has the advantage that there is no impact on the installed infrastructure associated with fire detection and suppression (sprinklers).

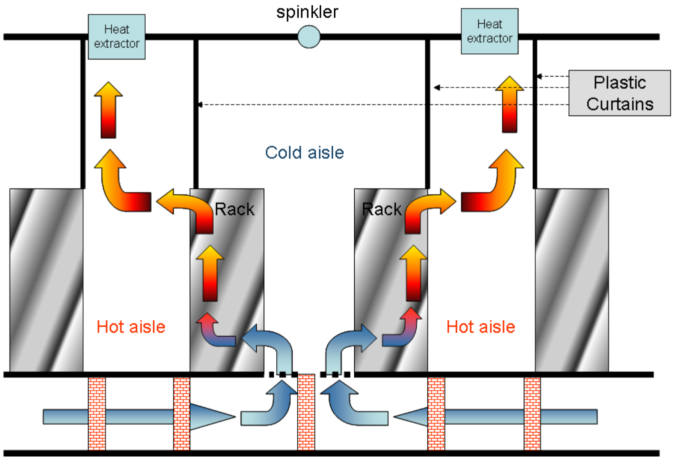

Figure 16: Segregation of hot/cold aisles using curtains

7.2.2.1.5 High density areas

Vendors are predicting trends for equipment with significantly increased the energy consumption density (kW/m^2). This is an issue for many of legacy data centres that are not designed to provide such high levels of cooling.

As an example, how could a computer room built with an average ratio of 0,7 kW/m^2 be adapted to new generation racks of blade servers, for which it is necessary to cool 20 kW over a 1 m^2 area.

If there are no restrictions on floor space, such as in a data centre following consolidation initiatives, this problem can be solved by not filling the cabinets fully, and locating the cabinets such that maximum efficiency of cooling is maintained without having to change existing air flows. However, where the use of floor space has to be optimized there will be commercial pressure to fully load cabinets. For this case, very high density areas are the appropriate answer, but these areas have to be considered separately from other "traditional" IT equipment areas and provided with their own energy and cooling needs.

The "pod" concept is supported by a number of vendors and extends the "hot aisle, cold aisle" concept by reducing the volume within the data centre that is subject to environmental control. Pod are enclosed and secured areas containing a specific set of cabinets/racks which are provided with the required level of environmental control (as shown in figure 17 and figure 18).

One of the advantage of such a solution is these areas can be installed outside computer rooms.

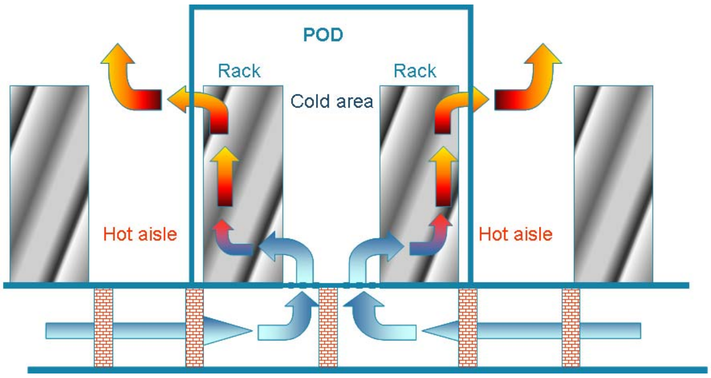

Figure 17: The "pod" concept

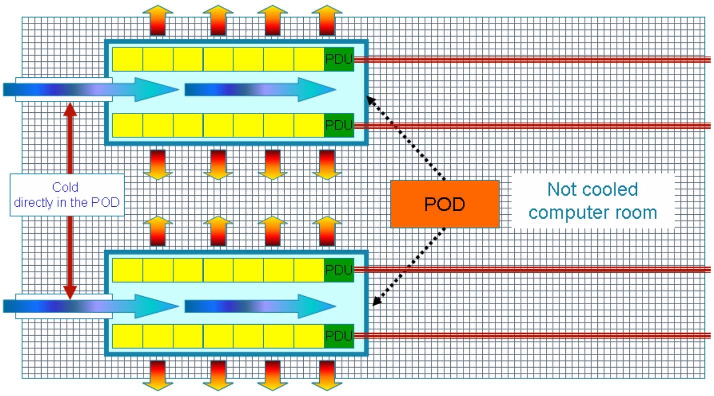

Figure 18: Pods within a computer room

7.2.2.2 Reduction of thermal waste in cabinets/racks

Any potential improvements in energy usage offered the approaches detailed in clause 7.2.2.1 by the approaches of may be impacted by the failure of installers and maintainers to fit/re-fit "blanking panels" in the front and rear of the racks within cabinets when equipment is not installed or has been removed - which creates significant losses of cold air coming into the rack and contributes to a inefficient management of cooling as indicated in figure 19.

This approach involves very little cost and is able to demonstrate a rapid reduction in energy usage.

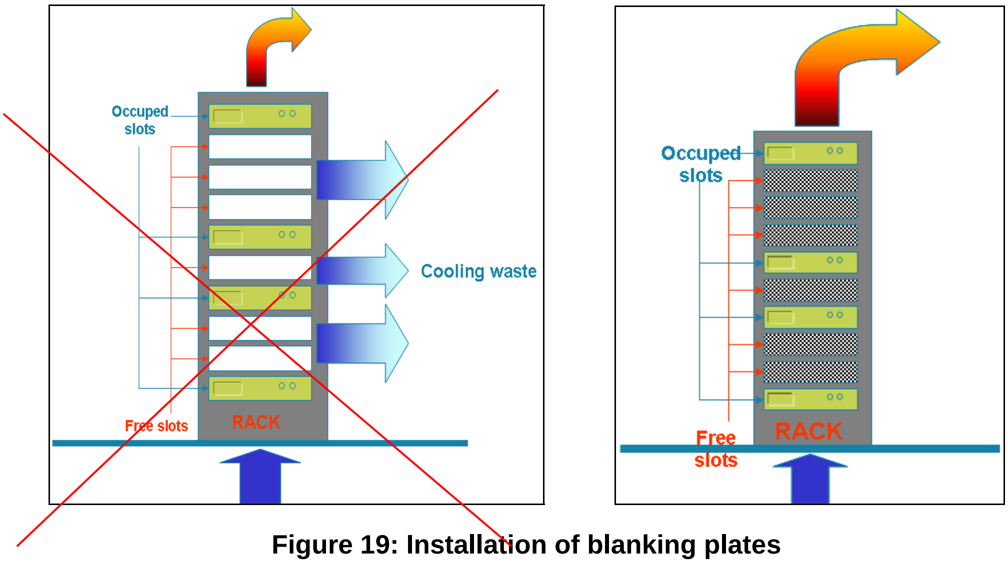

Figure 19: Installation of blanking plates

7.2.3 Modification of temperature and humidity

7.2.3.1 General

Increasing the temperature and adjusting humidity levels in computer rooms without violating vendors' specifications enables substantial reductions in energy usage associated with environmental control without significant Capex. The level of reduction depends upon some basic factors such as the size of the room, occupancy ratio.

EN 300 019 standards define the environmental classification for network telecommunications equipment. The European Code of Conduct [4] requests to vendors to consider changing the current temperature ranges applicable to information technology equipment to approach or adopt the same as those for network telecommunications equipments defined in EN 300 019-1-3 [11].

Experiments have been undertaken by some major telecommunications operators to determine the impact of increasing the average temperature in computer rooms without violating vendors' specifications.

7.2.3.2 Results of experimentation

The experiment described below was undertaken in a telecommunications operators data centre in Paris.

Details of experiment:

- The underground data centre comprised a 1 000 m^2 computer room operating non-critical systems (development, backup, etc.) within which the heat dissipation was in the range 300 W/m^2 to 1 500 W/m^2. The information technology equipment comprised servers, disk arrays, robotics and networking equipment and exhibited operational temperature ranges from 18 °C to 28 °C - 30 °C (dependent on vendor specifications) which is more restrictive than the EN 300 019-1-3 [11] applied to network telecommunications equipment.
- The computer room cooling was provided by air cooling (recycling mode) from seven dry cooler units (80 kW per unit) and the energy consumption for air cooling was between 40 % and 60 % of the total energy consumption
- Measurement instrumentation comprised 40 temperature and hygrometry sensors (measurement every 5 min) together with measurement of air cooling energy consumption.
- The first change in environmental conditions is defined in table 13. This resulted in reduced the energy consumption by 12 % without operational failure since 2007. This is shown in figure 20.

Table 13: Step 1 changes in environmental conditions

	Previous settings	*New settings*
Cooling unit start-up	*20-24 °C*	*22-26 °C*
Dehumidifier start-up	*50-65 %*	*60-75 %*
Humidifier start-up	*45-50 %*	*35-40 %*
Winter/summer switching point	*17°C*	*19°C*

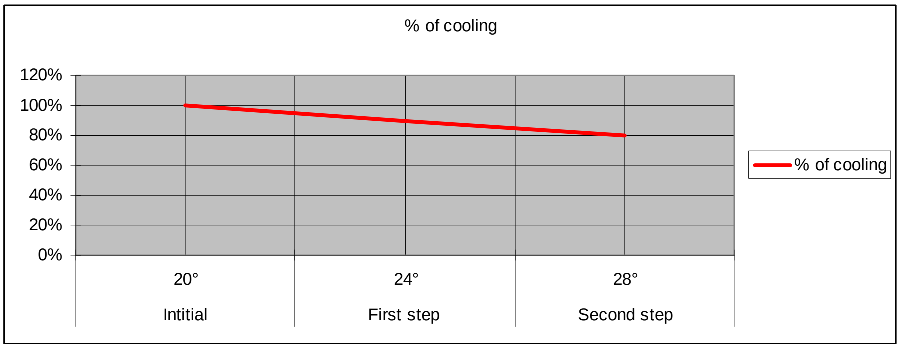

Figure 20: Computer room temperature and subsequent savings in energy usage

- The second change in environmental conditions is defined in table 14. This resulted in reduced the energy consumption by 20% without operational failure since 2007. This is also shown in figure 20.

Table 14: Step 2 changes in environmental conditions

	Previous settings	***First step***	***New settings (2)***
Cooling unit start-up	*20-24 °C*	***22-26 °C***	***24-28 °C***
Dehumidifier start-up	*50-65 %*	***60-75 %***	***60-75 %***
Humidifier start-up	*45-50 %*	***35-40 %***	***30-35 %***
Winter/summer switching point	*17°C*	***19°C***	***19°C***

7.2.3.3 Time before "system-shutdown"

Many types of information technology equipment undergo automatic shutdown when temperatures exceed the vendors' maximum operating temperature specification (typically 30 °C or 32 °C). The available time to repair and/or restart cooling before automatic shut-down occurs is a major operational concern.

The time to repair is defined as the interval between a total failure of the air conditioning system (with no redundancy) and the time at which the temperature in the room reaches a maximal functional limit. It depends upon a number of variables, the most important of which are:

- the area and volume of the space/room being cooled;
- the contents of the space/room;
- the quantity and type of information technology equipment (racks, servers, etc.) and the global electrical consumption;
- the type, number and capacity of the CRAC units;
- air conditioning system redundancy;
- the operating temperature of the room.

In circumstances where the total cooling demand of the equipment is very high the time to repair can be short and is further reduced if the higher operating temperatures are applied.

Figure 21 shows the results of experimental work (based on computation and observation of real events) in a computer room with the following characteristics:

- thermal load: 362 kW;
- computer room area: 1 080 m^2 (thermal load per unit area: 335 W/m^2);
- computer room volume: 3 240 m^3 (thermal load per unit volume: 112 W/m^3);
- cooling system: an air-conditioning system working in pure recycling mode (no free cooling).
  Seven dry cooler units (each with a cooling power of 80 kW) are used, associated with seven air treatment units. Cold air is blown through perforated tiles on the floor and the hot air outputs through the ceiling.

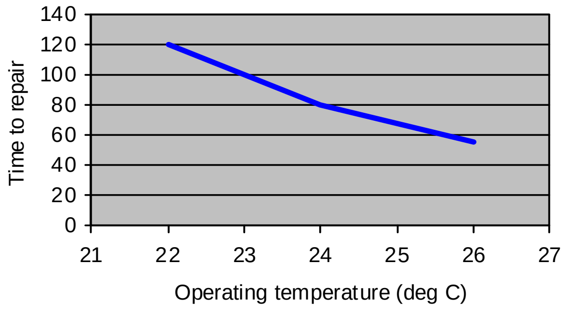

Figure 21: Example showing effect of operating temperature on "time to repair"

Figure 21 clearly shows that increasing the operating temperature significantly reduces the time before system-shutdown. This is an obvious concern for IT Managers and Operation Managers and may lead to this approach to energy usage reduction being ignored. However, a variety of approaches can be applied to provide lower, but significant, savings including:

- segregation of strategic "mission critical" business from other less critical activities and only apply increased operating temperature and humidity to those less critical areas;
- segregation of network telecommunications equipment from the information technology equipment and only apply the increased temperature and humidity to the network telecommunications equipment.

NOTE: Such segregation would not be required and greater savings would be possible if information technology equipment vendors change their operating specifications to those of network telecommunications equipment as requested by the European Code of Conduct [4].

7.2.3.4 Restrictions on implementation

In an Uptime Institute Tier 4 data centre, the risk posed by a cooling system failure is reduced by the presence of redundant environmental control systems and/or power distribution equipment. Any associated risks of operating the information technology and network telecommunications equipment are also minimized.

However, in data centre meeting the requirements of lower Uptime Institute Tiers, those risks and their resultant impact on the business conducted from and by the data centre need to be analysed.

7.2.4 Alternative cooling mechanisms

7.2.4.1 Free cooling

If the opportunity exists and outdoor conditions permit, the re-introduction of air-side economizers (free cooling) or water-side economizers (cooling tower) approaches may be considered to reduce energy usage without heavy Capex.

Free cooling uses external air or water temperature conditions for cooling rooms by introducing fresh air to the computer room equipment using traditional methods. This can produce significant reductions in energy usage and has a direct effect on the improvement of PUE (since cooling may represent 35 per cent to 45 per cent of total energy consumption.

There are two principle approaches to "free cooling":

- "free air-cooling": based on the air temperatures outside the building containing the data centre;
- the use of existing cold water sources, as sea, lake, river or other.

The ideal climatic combination at the intake and exhaust of the equipment lies in the range 22,5 °C to 22,7 °C and 45 % to 50 % relative humidity without irregular deviations in either parameter. If the suppliers of information technology equipment were to adopt the operational environmental specification applied to network telecommunications equipment (see EN 300 019-1-3 [11]) the period of time that free cooling could be applied and the resulting saving (see table 16) would generally increase.

NOTE: Optimum environmental control will be achieved by monitoring the environmental conditions should take place as close as possible to the intake and exhaust at the rack level.

The effectiveness, and therefore the use, of free cooling are determined by two principle climatic factors external to the data centre premises:

- mean air temperatures throughout the year (see figure 22 shows the mean temperatures across the world);
- mean relative humidity throughout the year.

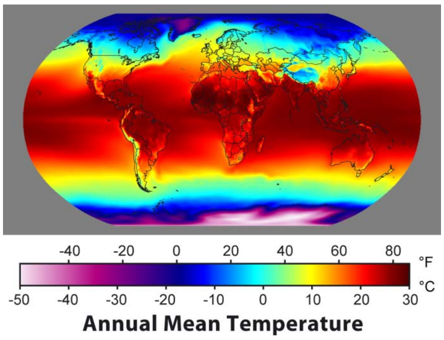

Figure 22: Annual mean temperature across the world

Total reduction in energy usage is directly linked to the external climatic conditions. The potential reduction is proportional to the total period during which the external environmental conditions match the ideal combination. Table 15 provides a methodology to evaluate energy savings for the use of free cooling.

NOTE: 100 % of matching climatic conditions do not correspond to 100 % reduction of energy consumption since the remainder of the cooling infrastructure such as chillers, CRACs, pumps and fans are still required and consume energy (only the refrigerant batteries are inactive)."F" in table 15 is typically 0,6 (as is used in figure 16 and figure 23).

Table 16 and figure 23 show the result for a series of examples using this methodology.

Table 15: Free cooling savings profile

Total energy consumption of environmental control system	W
Total number of hours per year	8 760
Number of hours of matching conditions during year	H
Energy consumption fraction of environmental control system made inactive during periods of matching conditions	F
% of potential free cooling per year	P
Savings	R
% of potential free cooling	P=(H/8760)
Savings on cooling	R = P x W x F

Table 16: Example of free cooling savings

No of hours of potential free cooling (hours[days]/year)	% of free cooling per year (H/8760)	Cooling load savings (P x W)	Total energy savings (R)
1 000 [42]	11,42 %	4,22 %	2,53 %
2 000 [83]	22,83 %	8,45 %	5,07 %
3 000 [125]	34,25 %	12,67 %	7,60 %
4 000 [167]	45,66 %	16,89 %	10,13 %
5 000 [208]	57,08 %	21,12 %	12,67 %
6 000 [250]	68,49 %	25,34 %	15,20 %
7 000 [292]	79,91 %	29,57 %	17,74 %
8 000 [333]	91,32 %	33,79 %	20,27 %
8 760 [365]	100,00 %	37,00 %	22,20 %

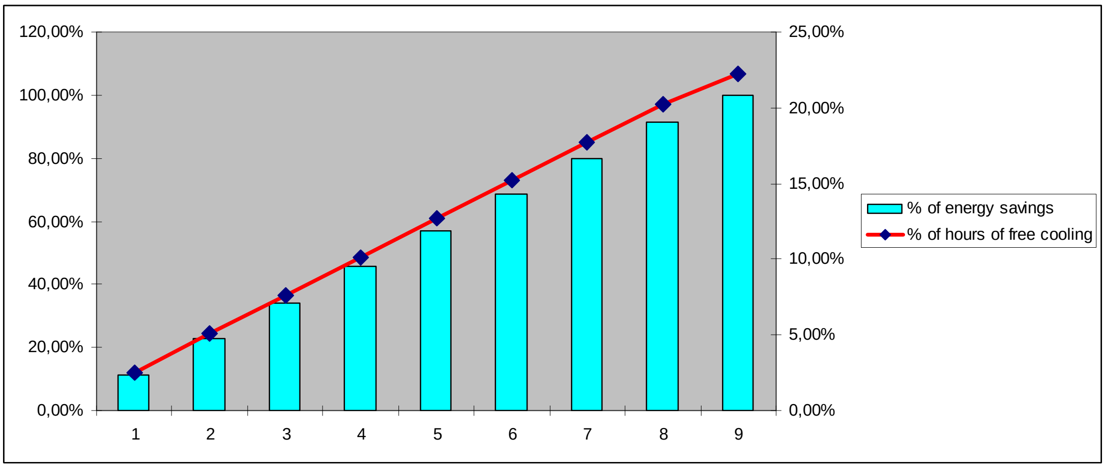

Figure 23: Free cooling savings

7.2.4.2 Direct liquid cooling

This topic is included here for completeness but it is for future study.

7.2.4.3 Emerging technology (auto cooled chassis or chip-level cooling)

7.2.4.3.1 Cooling-on-the-chip

This next generation technology aims to directly apply the cooling to the semiconductor packages within equipment such as servers. Hardware manufacturers are working on these technologies which are predicted to provide significant improvements in energy efficiency.

7.2.4.3.2 Auto-cooled chassis

Chassis or racks already exist which contain integrated cooling systems (air or liquid).

7.2.5 Enhancements of cooling systems

The following enhancements of cooling systems involve high Capex and Opex and represent a significant risk in terms of business continuity and quality of service during the implementation phase:

a) installation of high-efficiency variable-speed air-handler fans and chilled water pumps;

b) optimization of data centre airflow configuration;

c) installation of high-efficiency CRAC units;

d) sizing/re-sizing of cooling systems and the configuration of redundancy to maximize efficiency;

e) increasing the temperature difference between chilled water supply and return in order to allow a reduction in chilled water flow rates.

8 Infrastructure requirements to optimize energy efficiency

The implementation of effective planning of pathways and spaces in accordance with EN 50174-1 [8] and the data centre specific aspects of EN 50174-2 [9] is intended to maximize the energy efficiency of environmental control systems. The key aspects are:

- the use of generic cabling leading to structured growth and changes via distributors as opposed to ad-hoc direct point-to-point cabling within the computer room (except in very specific localized areas);
- the use of hot aisle and cold aisle or POD approaches together with detailed requirements and recommendations to ensure that cooling air reaches the intended location and exhausts are not obstructed.

9 Improvement of energy efficiency of power distribution systems

9.1 General

This clause describes the approaches that may be employed to increase the energy efficiency of the power distribution systems within the data centre.

Some of the approaches are applicable to legacy data centres while others are more likely to be applied in new data centres. Further details are provided in clauses 10 and 11.

Figure 24 shows the power distribution systems within the data centre from the main electrical arrival to, and within, the network telecommunications and information technology equipment. Each component in the power distribution system has a specified and measurable energy efficiency.

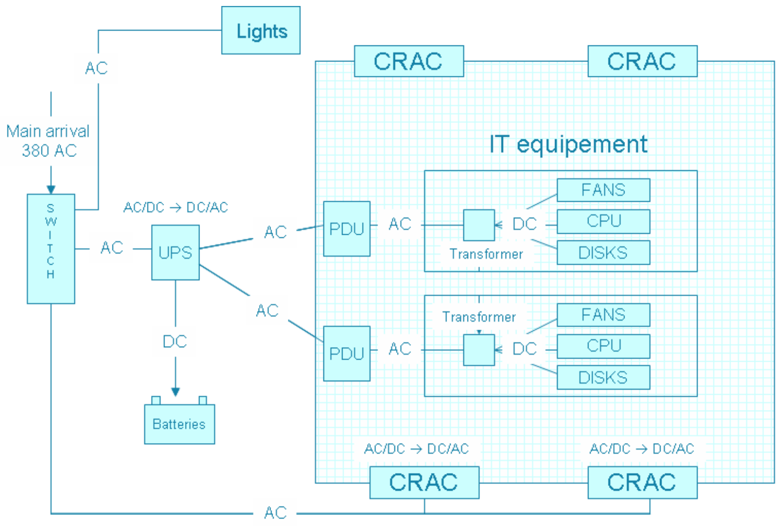

Figure 24: Power distribution system

Actions to increase the efficiency of the power distribution system in an existing data centre will, except for actions such as the deployment specific software solutions for energy management, require significant Capex investments and is not without risk for business continuity. Indications of the impact of some of the actions in this domain of energy efficiency are shown in annex A.

9.2 Uninterruptible Power Supplies (UPS)

9.2.1 Efficiency

The most popular UPS technologies including "line interactive " and "double conversion" (see figure 25). UPS technologies such as "delta conversion" are recognized but are proprietary technology from one specific vendor.

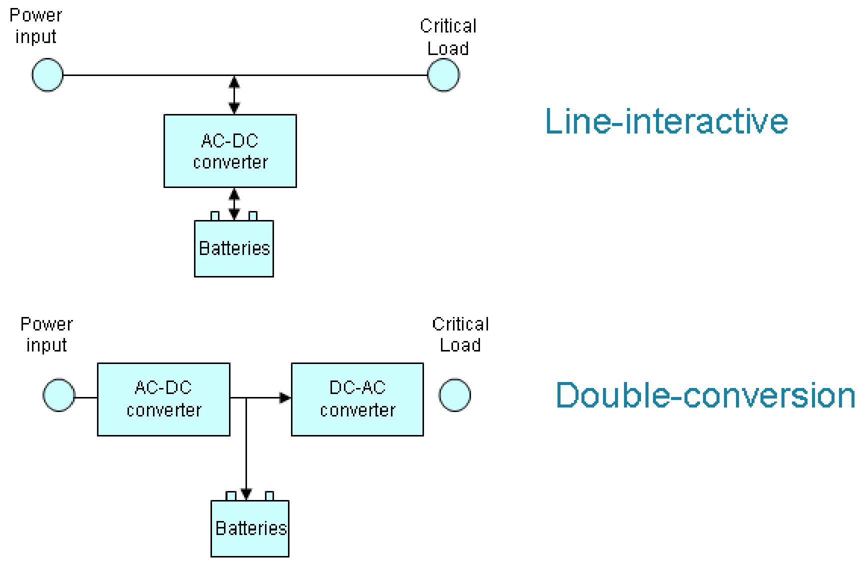

Figure 25: Multi-vendor UPS technologies

The IEC defines the efficiency of a UPS as "the ratio of output power to input power under defined operating condition".

As shown in figure 26:

- a UPS is generally more efficient when used at full load;
- line-interactive technology is more efficient than the double conversion for reasons linked to input load.

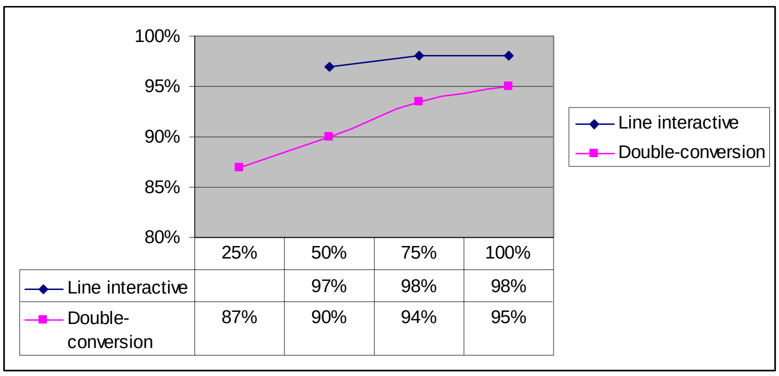

Figure 26: UPS efficiency

However, the data in figure 26 is based on an average of specified performance values taken from main UPS vendors and was reported in a report from Ecos Consulting-Epri Solutions [10]. The efficiency of the different technologies varies significantly. For example:

- at 50 % load, the most efficient technology is the "line-interactive" (with a 98 % efficiency);
- "double-conversion" technology is the least efficient;

- in general, UPS running in excess of 50 % load have good efficiency, but there is a dramatic decrease under 40 % (see figure 26).

In many legacy data centres, UPS systems are installed with the maximum capacity anticipated for the future needs. It is not uncommon for this capacity to never be fully used. In addition, redundancy requirements also promote the operation of UPS systems below their full capacity. In the case of 2N redundant configuration (independent of the UPS technology), it is always necessary to manage power so that systems are loaded beyond 50 % of capacity. Consequently every "2N" mode operates at less than maximum efficiency.

Modular UPS systems (see clause 9.2.2) allow the capacity to be mapped to the demand thereby improving energy efficiency.

Both UPS types can be mixed, in either the same or in different zones of the data centre. Some use a traditional UPS as their main source, but use smaller, modular systems as the second source for their most critical hardware to give "2N" redundancy without incurring that cost for the entire data centre.

9.2.2 Modular UPS

In legacy data centres, UPS commonly use double conversion technology, as shown in figure 25. This means two conversions (AC-DC and DC-AC), generating consequent losses of energy.

Recently, vendors have begun to offer smaller modules (from 10 kVA to 50 kVA) to build "modular" UPS systems. The main advantage of the modular UPS approach is the ability to grow capacity taking account of real needs (with an initial sizing). These modules are "hot pluggable" and can be removed for exchange or maintenance. Modular systems are also generally designed to accept one more module than required for their rated capacity, making them "N+1 ready" at lower cost than large traditional UPS system.

Used correctly, modular UPS systems run at higher efficiency since they can be used close to maximum rated capacity (see figure 26).

The main concern regarding modular UPS is reliability since the probability of failure increases with the increased number components employed.

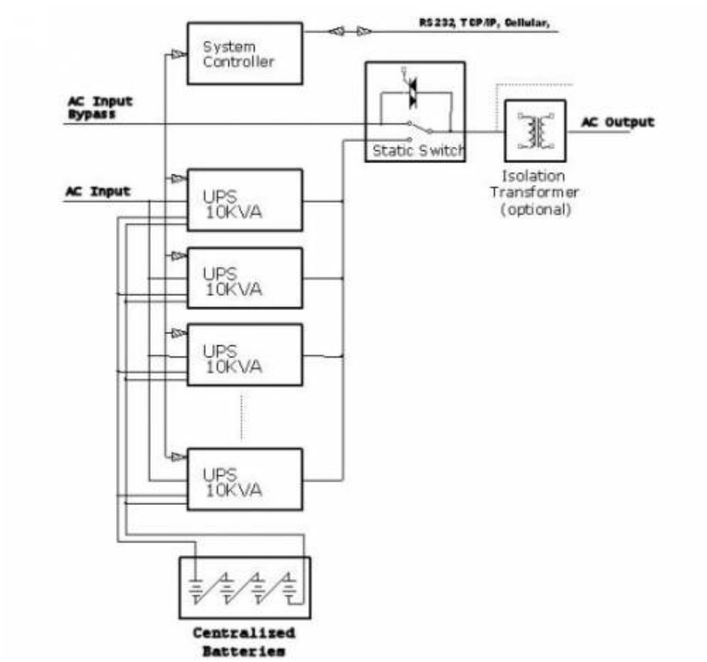

Figure 27: Example of modular line interactive UPS technology

9.3 Energy efficiency improvement solutions

9.3.1 Measurement of energy efficiency of existing equipment

This involves the review of the existing power distribution equipment within the chain in terms of its energy efficiency. This may be undertaken without significant Capex.

9.3.2 Energy capacity management

This involves the use of electrical capacity management tools in order to deliver a more efficient usage of consumption. This requires significant Capex (for the monitoring equipment) and Opex (installation and software).

The implementation of electrical capacity management can generate some immediate reduction in terms of consumption, without any work on the power distribution components, and without any risks for business continuity.

9.3.3 Review of policy

9.3.3.1 General

The implementation of improvements in energy efficiency by changing the distribution policy (such as AC or HVDC) have an impact on the equipment used and can represent significant Capex and a risk to business continuity. Such changes cannot be recommended in existing data centres.

9.3.3.2 HVDC versus AC

ETSI has published EN 300 132-3 [12] which discusses 400 V AC and DC power distribution. Some trials of HVDC have been undertaken on an international level by a number of organizations such as power distribution manufacturers, hardware vendors and also by laboratories, universities. These trials clearly demonstrate an interest to investigate more deeply and provide full 400 V DC directly into the IT equipment.

Many documents, studies, and standards suggest that direct DC can generate some savings from 5 % to 15 % depending on several conditions.

Technically, servers from main vendors could accept direct 400 V DC and work on this continues.

One of the main problem concerns safety rules when working in a HVDC environment (which is quite unusual in the IT world) making maintenance and technical interventions on a DC supply chain much more costly than on AC, due to lack of competencies on the market.

The present document judges that it is too early to take a definitive position on the usage of HVDC and that the work undertaken, results from trials and the announcements of main actors (vendors, manufacturers) will be reviewed before further recommendations are made.

NOTE: Several standardization institutes are working relevant areas such as HVDC 400 V plugs, connectors and safety breakers.

9.3.4 High efficiency distribution equipment

This involves the installation of individual power distribution equipment with improved energy efficiency specifications. Typical examples include:

- high-efficiency power distribution units;
- high-efficiency motors in fans and pumps;
- UPS units that exhibit improved efficiency over the over full range of load;
- rotary-based UPS units;
- correctly sized power distribution and conversion to optimize efficiency.

These solutions represent significant Capex and a risk to business continuity.

9.3.5 Backup power

The use of on-site generated energy (more often with fuel generator sets) as back-up to the main electrical power source does not improve energy efficiency. It is intended to be deployed in crisis situations where the aim is to get back to normal operation as quickly as possible.

Note that backup power infrastructures shall be designed to provide 100 % of the energy needs for a specified and not insignificant period (e.g. 72 hours minimum).

10 Energy efficiency within existing data centres

Table 17 shows the short and medium term actions that may be employed to improve energy efficiency within existing data centres by reference to the clause numbers of the present document.

Table 17: Short and medium term actions within existing data centres

Technology area	Short term actions	Medium term actions
IT infrastructure	6.3.1 6.3.2 6.3.3.2 6.3.3.3 6.3.4	
Environmental control systems	7.2.1 7.2.2.1 7.2.2.2 7.2.3 7.2.4.1	7.2.5
Physical infrastructure	8	
Power distribution systems	9.3.2	9.3.3 9.3.4

11 Energy efficiency within new data centres

11.1 General

This clause details the discipline that requires consideration during the design, installation and operation of a data centre in order to maximize savings on energy usage and reducing its "Carbon Footprint".

Many of the solutions detailed in clauses 6, 7, 8 and 9 are applicable during the design of a new data centre. However, it is clear that the design and operation of the most energy efficient data centres requires a multi-disciplinary approach with each discipline contributing in its own domain to make energy usage more efficient - rather than one specific initiative concerning a single domain.

Every aspect of the data centre requires detailed consideration, ranging from its location via the consolidation actions employed to the tools and processes adopted to monitor the operational aspects of the data centre. The process is described schematically in figure 28.

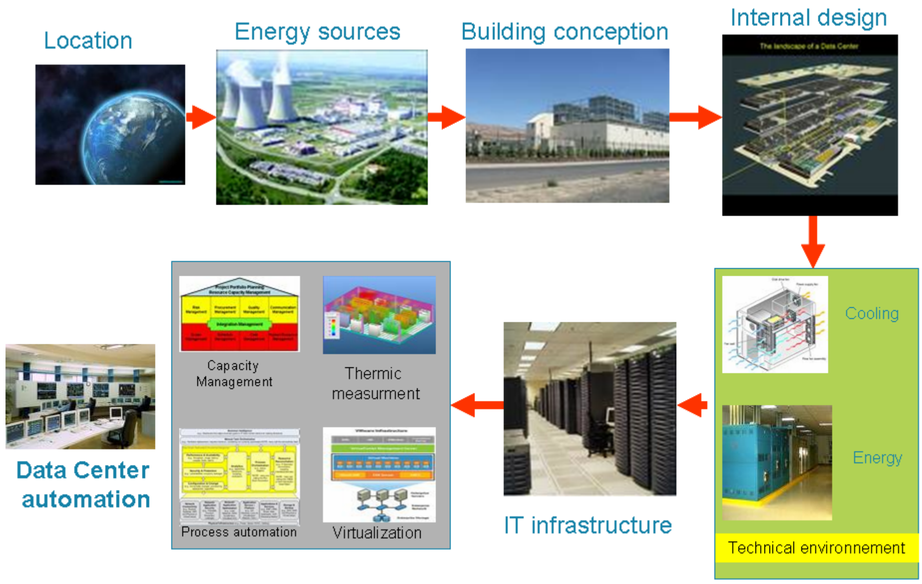

Figure 28: The path to energy efficiency within new data centres

11.2 Multi-disciplinary approach

11.2.1 Location

The selection of a location is primarily affected by the availability, and quality, of power supplies.

National or local legislation, regulation and constraints concerning industrial buildings influence the selection of location of new data centres. However, facilities offered by local authorities such as tax reductions, subventions, green bonus also have an effect.

The selection of location is fundamental in relation to the "green" credentials based upon the following criteria and impacts:

- the availability of "green" energy sources such as solar, wind or hydraulic;

NOTE 1: Solar or wind energy sources do not guarantee the continuity required for a data centre, but can represent a non-negligible part of energy when appropriate conditions are respected.

- the nature of energy present (renewable or not), its costs, its quality and the availability of backup supplies;
- the type of materials you will use for the building (walls, roof, windows, etc.);
- proximity of airport, ports, railway stations, highways - which has consequences on the shipping costs and on CO^2 emissions associated with transport.

The selection of a location is important for energy efficiency since climatic conditions will influence "cooling" approaches (which will both reduce the need for cooling).

NOTE 2: In case of free-water cooling, proximity of a river, lake, sea, or other source of cold water will be a main factor. In such a case, floods or Tsunamis risks have to be taken in account.

11.2.2 Energy sources

11.2.2.1 Main energy

In case of an Uptime Institute Tier 4 data centre, two separate sources of energy, coming from two remote plants or two different sources of energy are required and this dramatically limits the range of potential locations around the world that can support such data centres.

11.2.2.2 Backup energy

Backup energy has to be dimensioned to cover 100 % of data centres needs for a determined period and provision has to be made using "on site" sources with generators or other renewable energy solutions.

11.2.3 Building conception

The specification of a building housing a data centre is directly impacted by its location. The orientation of the building together with the thermal insulation properties of the building walls, roof, windows, have a significant effect on the power demands associated with environmental control of the building interior.

European standards for energy efficient buildings are in preparation. National approaches exist for energy efficient buildings such as HQE (*Haute Qualité Energétique*) for France and LEED (*Leadership in Energy and Environmental Design)* founded in the USA.

11.2.4 Internal design

In order to maximize the overall energy efficiency, the internal structure and configuration of the data centre should be analyzed to ensure that the appropriate power and environmental control resources are applied in different areas. A minimum set of identifiable spaces within data centres are:

- Computer rooms (principally accommodating information technology equipment).
- Network telecommunications rooms (accommodating network telecommunications equipment).
- Technical rooms (accommodating UPS, batteries, etc.).

Each of these, and any other, spaces shall be assessed in term of their demands for electrical power and cooling based upon their size and operating environmental conditions. Computer rooms also should be assessed in terms of the need for, and structural implications of, the constructions of raised floors.

The internal design process should also cover the impact on cooling system introduced by the necessary utilities and services in the spaces. To support this activity, the design of data centre cabling is specified in EN 50173-5 [7] and EN 50173-2 [6] and the planning and installation of the cabling is specified in EN 50174-1 [8] and EN 50174-2 [9].

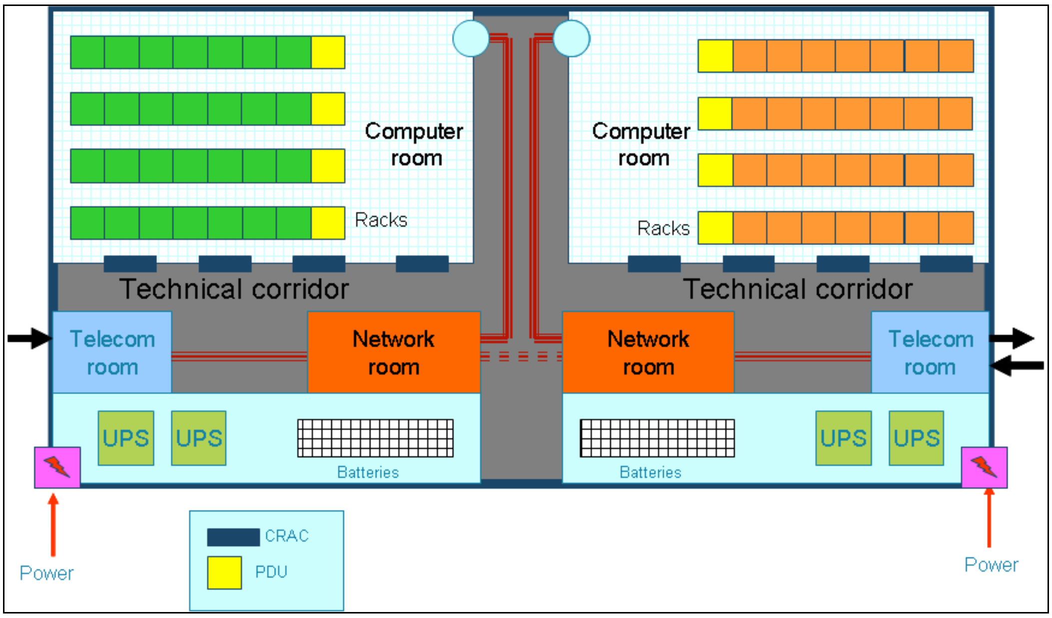

Figure 29: Example schematic of an Uptime Institute Tier 3 data centre design

11.2.5 Energy and cooling

The introduction of high efficiency power distribution systems and components as described in clauses 9.3.3 and 9.3.4 has a direct effect on external power demand in support of the information technology and network telecommunication equipment but does not have an automatic effect on PUE unless accompanied by improvements in the power demand of the environmental control systems.

The use of free cooling as described in clause 7.2.4.1 represents the primary method of improving PUE in new data centres. The reduction of energy involved in environmental control by means of increasing operating temperature and humidity as described in clause 7.2.3 also has a direct effect in PUE.

11.2.6 IT infrastructure

The IT infrastructure of new data centres has to:

- be designed for maximum agility, scalability, and a common shared architecture;
- adopt the most efficient hardware technologies proposed by the vendors (requiring vendors to prove the real efficiency of their technologies, and their impact on TCO).

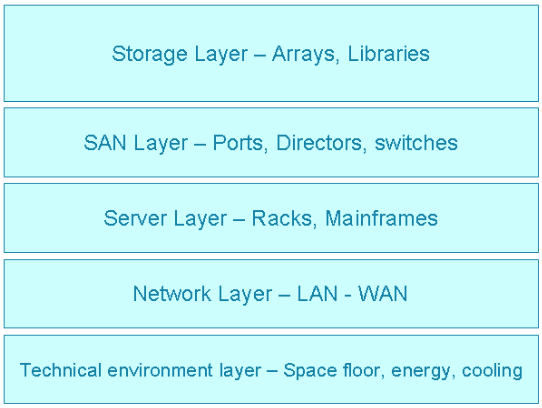

Figure 30: IT infrastructure layers

Independent of the Uptime Institute Tier level a new data centre will comprise a combination of:

- a rack-based architecture using very powerful serves will need high-density areas using, as a minimum, covered cold aisles with raised floors(as described in clause 7.2.2.1.3) or pods (as described in clause 7.2.2.1.5) in perhaps otherwise uncooled rooms and which are exhausted outside the building). This means the size of the computer room is not a main criteria, but the greatest concern is to carry the right resources to the right place. For these areas, the most appropriate indicator is the ratio of total power provided to the rack / total computational power (that is X kW → Y TPM-C, described in clause 5.3.2.1);
- legacy servers or Unix mainframes, which have the same problems as racks, consuming a lot of energy in a small area, and with an important heat dissipation;
- other kind of equipment such as robotics, storage, network telecommunications equipment, or non racked servers that may be accommodated in traditional computer rooms, for which the right ratio is the kW/m^2 (see clause 5.3.2.1). It should be noted that storage arrays, due to increased density, have energy needs approaching those of servers. We calculate an average need of 300 W per Terabyte. (For example, last generation of disks arrays can have more than 100 Tb capacity on a few square metres) .A future trend for next generation storage will be "Diskless storage" but the actual costs of these solutions to compare with disks technology make it not attractive for Companies.

11.2.7 Software

Infrastructure management software brings some benefits in the way to manage the energy needs more efficiency, and have a clear view of the potential of power in a computer room, in a rack, in a server.

Same if some thermic measurement tool is implemented in the Computer room. This can visualize the hot zones, in the room or in a rack, and propose some modelization for a better design of the computer room. This contributes to the improvement of the efficiency of energy usage for cooling.

11.2.8 Processes

Two major types of actions are to be launched.

First, a global consolidation programme to reduce the number of technical and logical components, which are the main cause of the energy expenses (see topic on consolidation methods),

Second time, a set of automation tools for operations in the data centre. The NGDC has to be seen as a "Service Factory", providing the right service, with the right resources, at the right time. That means all operational steps have to be automated - from the provisioning of resources to the way applications are managed.

The ITIL (IT Information Library) best practices are key for a good management of the IT.

Involved ITIL process in energy efficiency:

- Asset management to know what is existing in the IT;
- Capacity Management measure how what exists is used and size more precisely application needs;
- Service Level Management to provide the right level of service for an application.

12 Conformance

12.1 Existing data centres

To conform to the present document the following assessment needs to be undertaken:

- A: Awareness of total energy consumption of the designated data centre.
- B: The total computational load.
- KPI = B/A.

To achieve minimal conformance over a period of time the KPI has to show an improvement of 15 % compared to the original KPI.

To achieve basic conformance over a period of time the KPI has to show an improvement of 25 % compared to the original KPI.

To achieve enhanced conformance over a period of time the KPI has to show an improvement of 50 % compared to the original KPI.

12.2 New data centres

To conform to the present document:

- the target PUE shall be calculated in accordance with clause 5.3.1.
- equipment and processes shall be established to enable measurement of the electrical consumption at different points to determine PUE (i.e. to determine the comparative values of clause 5.3.1).

The initial PUE is unlikely to meet the target PUE but may be expected to reduce and based on the assumptions of the plan shall meet the target PUE when the total computational load matches that of the target. Unless there are changes in application or performance of the data centres that justify a re-calculated target value, the actual PUE of the data centre based on total computational load shall continue to meet or exceed the target value.

New data centres should be designed with a low PUE target, which means every aspect of the design has to be studied, from the location where the data centre is to be built, to the way the applications will be operated in the data centre. This also covers the future plans of the vendors to improve the energy efficiency of the equipment, such as direct cooling of the heat sources at chip level. Assessment of energy needs in terms of energy per square metre will no longer be appropriate, due to the localized needs of some individual racks (perhaps 40 kW to 100 kW in 2010). Some specific closed high power-density areas (pod) should be built inside the computer rooms.

In parallel a programme of information system (IS) rationalization should be launched, to minimize the number of applications, operating systems and IS complexity, since these are frequently the major causes of the excessive energy consumption. Such work on functional complexity is the most beneficial approach to maximising energy efficiency.

13 Recommendations

13.1 Existing data centres

13.1.1 General

There are three mains approaches concerning energy management in an existing data centre:

- reduction of PUE (increase efficiency);
- reduction of energy consumption (cost effective);
- optimum usage of existing resources (environment and/or technical).

For each of approaches some specific actions should be initiated. It should be noted that some actions will have positive effects via multiple approaches.

13.1.2 Reduction of PUE

In addition to conforming to clause 12.1, it is recommended that equipment and processes be established to enable measurement of the electrical consumption at different points to determine PUE (i.e. to determine the comparative values of clause 5.3.1).

All actions related to reducing power consumption of cooling systems infrastructures (chillers, fans, engines, pumps) and increasing he energy efficiency of the power distribution system (UPS, PDU, Transformers), will have a positive effect on the PUE. It should be noted that most of these actions are costly, and represent a risk for business continuity.

13.1.3 Reduction of energy consumption

In addition to conforming to clause 12.1, it should be noted that actions to reduce the energy consumption of the information technology equipment (such as power management, capacity management and consolidation initiatives) will have a positive effect on energy consumption and energy costs, but could have a negative effect on PUE. These actions can be balanced by the actions in computer rooms as such as aisle segregation, higher operating temperatures and free cooling.

13.1.4 Optimum usage of existing resources

In addition to conforming to clause 12.1, actions should be undertaken to optimize the usage of existing resources including capacity management, virtualization and logical consolidations in order to increase computational load of servers.

13.1.5 High density areas

High-density areas should not be created that would impact air-flow and the effective cooling of existing areas. In enterprise data centres, it is recommended not to load racks fully and to balance the use high- and low-density equipments in rows.

13.2 New data centres

13.2.1 General

For a Tier 1 or 2 data centre, the target PUE should be in the range 1,3 and 1,5 (DCIE 66,7 % to 76,9 %).

For a Tier 3 or 4 data centre, the target PUE should be in the range 1,6 and 2,0 (DCIE 50 % to 62,5 %).

NOTE: The target ranges reflects the opportunity for free-cooling as shown in table 16 and the higher values of PUE for higher Tier data centres reflects the need for redundancy of cooling and power distribution systems.

13.2.2 Location study

This should be the main factor in the planning of future data centres.

Following a comprehensive study covering risks (including natural, political, economical etc.) the choice of the location of the data centre will have a determinant aspect on all costs (Capex and Opex) as follows:

- locations with mean annual temperatures able to provide the maximum free cooling potential can generate up to 60 % of savings on cooling energy;
- the requirements for redundant infrastructures (separate power supplies using separate paths) in higher Tier level data centres place constraints on possible location.

13.2.3 Data centre construction

13.2.3.1 External

The nature of the building components has an important impact on energy consumption. The building should neither as a "pressure cooker" or a "Thermos flask". The building should be constructed according to "green" standards such as HQE and LEED.

13.2.3.2 Internal

Data centres should be designed with different areas which will not only provide improved energy efficiency but also minimize Capex and Opex, by considering:

- high density areas, treated separately from other areas, in small specific rooms;
- traditional computer rooms, with a kW/m^2 ratio;
- segregation of mission-critical from non-strategic business equipment;
- several environmental conditions regarding the criticality of the equipment and/or the free cooling capabilities during the year.

This is shown schematically in figure 31.

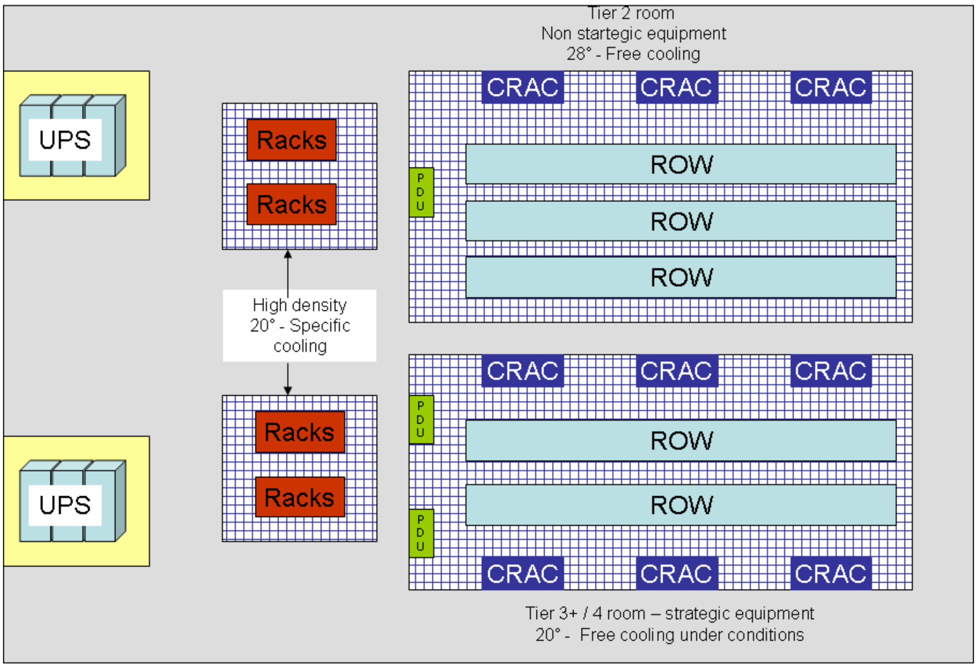

Figure 31: Environmental conditions

13.2.4 Cooling

13.2.4.1 systems

There should be a mix of traditional and free cooling systems. Traditional cooling should be applied:

- where areas needing heavy cooling, such as high density areas, rooms with strategic mission-critical IT equipments;
- when temperatures are not compatible with free cooling.

In all cases, the cooling equipments should be "energy efficient" requiring the latest generation of fans, pumps, engines.

13.2.4.2 Temperature control

Mean temperature in computer rooms should reflect the criticality of the equipment or the application.

It is not recommended to maintain low temperatures (18 °C to 20 °C) in rooms containing non-critical equipment or equipment that is specified to operate in accordance with EN 300 019-1-3 [11].

The introduction of an ETSI requirement for information technology equipment within network data centres to conform to EN 300 019-1-3 [11] is critical to allowing substantial reductions in energy consumption of cooling systems.

13.2.5 IT Infrastructure

13.2.5.1 Architecture and policy

Policy should be focused on shared components of the IT infrastructure.

As described in figure 32 the architecture should be based on blade farms for X86 servers, and partitionable high-end Unix servers or traditional mainframes for large database processing. For storage, a common shared infrastructure should be built on a SAN architecture. A data-tiering policy should be adopted with appropriate physical support (Low-cost storage for non critical data, high-end disk arrays for critical).

13.2.5.2 Automation and capacity management

As described in figure 33 all operation processes should be automated using specific tools including:

- supervision;
- backup;
- high-availability - disaster recovery;
- inventory - asset management;
- scheduling;
- monitoring (measurement);
- capacity management;
- utility computing.

13.2.6 Organization - processes

It is recommended that the data centre is viewed as a component of a "Service", for which it provides technical environment and a set of functionalities, specific for operating and monitoring those Services, as shown in figure 32. Note that the technical infrastructure components and the logical layers are existing only because the Service Delivery stands on it.

It is recommended that the data centre is considered as a factory, producing some Services, with an engagement of service level as shown in figure 33.

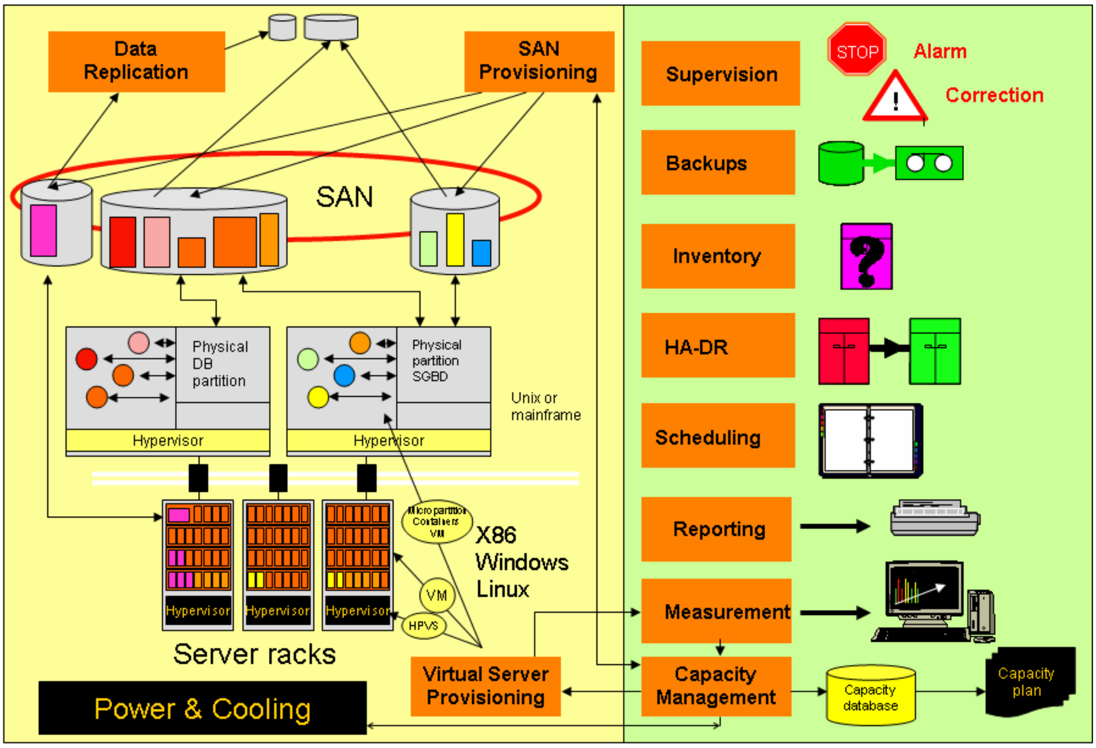

Figure 32: The future data centre internal operation architecture

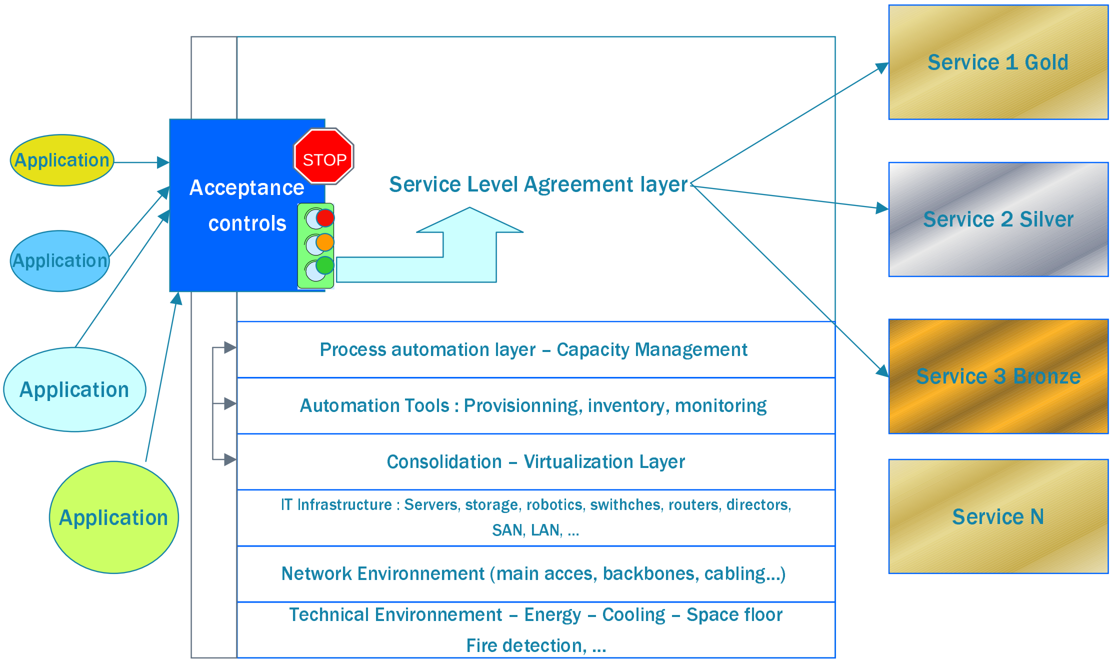

Figure 33: The data centre as a service provider

14 Future opportunities

14.1 General

Vendors in all areas are considering new solutions to meet changing needs for energy consumption, cooling and computing power and are proposing future equipment which integrates these demands.

14.2 Energy

- Produce your own renewable on-site energy, as much as possible.
- Re-use wasted heat.
- HVDC if significant energy consumption decrease is proved and tested.
- New generation, higher efficiency, UPS.

14.3 Cooling

Free air or water cooling opportunities are of major concern for new data centres and will have a great impact on the PUE.

Liquid or gas cooling directly in the servers or auto-cooled chassis or frames (previously used for mainframe computing components) are to be studied.

14.4 Energy efficient IT hardware

Follow the technical road-map of vendors, ensure in all sourcing or purchasing process of the real efficiency of equipment:

- servers: new standards concerning temperature range and component efficiency (transformers, direct DC in the server, cooling on the chip, efficient fans, dynamic power management, etc.);

- storage arrays, libraries, etc.;
- switches, routers, directors, ports, etc.

14.5 Software, tools and re-engineering

Data centres have to be equipped with a set of tools and gauges in order to be able to easily measure efficiency, and launch a set of initiatives concerning existing infrastructures including:

- measurement tools for monitoring IT activity (CPU, Memory, I/O, etc.);
- capacity management for IT infrastructure (the right resource at the right moment);
- utility computing (automation for processes) - inventory, provisioning;
- computing workload management;
- virtualization - grid;
- consolidation initiatives.

14.6 Consolidation initiatives

Think consolidation for servers, storage, and prepare common shared infrastructure to host applications as shown in figure 32.

Annex A (informative): Indications of the effect of energy efficiency actions

Table A.1 provides indications of reduction of energy consumption using some the actions detailed in clauses 6, 7 and 9.

The reductions shown are first order effects only, that is they do not show associated changes in one domain resulting from an action in another domain.

Table A.1: Indicative reductions in energy consumption

Actions		**Typical savings (%)**		**Nature of action**
Description	**Clause references**	**Within the domain**	**Overall**	
Domain: IT infrastructures				
Removal of obsolete equipment	6.2.1, 6.2.2	8	3,6 (see note 1)	Short term
Enablement of power management	6.2.3.2.1	10	4,5 (see note 1)	Short term
Activation of sleep mode	6.2.3.2.2	30	13,5 (see note 1)	Short term
Capacity management	6.2.3.3	8	3,6 (see note 1)	Short term
Virtualization	6.2.4.3	40	18,0 (see note 1)	Short term
Logical consolidation	6.2.4.4	45	20,3 (see note 1)	Short term
Functional consolidation	6.2.4.3, 6.2.4.4 and 6.2.4.5	60	27	Long term
Domain: Environmental control systems				
Improvement of cooling efficiency	7.2.2.1.2 7.2.2.1.3 7.2.2.1.4	30	11,1 (see note 2)	Medium term
High density areas	7.2.2.1.5	25	9,3 (see note 2)	Medium term
Modification of temperature and humidity	7.2.3	40	14,8 (see note 2)	Short term
Free cooling - air-side	7.2.4.1	50	18,5 (see note 2)	Medium term
Free cooling - water-side	7.2.4.1	30	11,1 (see note 2)	Medium term
Cooling-on-the-chip and auto-cooled chassis	7.2.4.3	80	29,6 (see note 2)	Long term
Increasing the temperature difference between chilled water supply and return	7.2.5	3	1,1 (see note 2)	Short term
High-efficiency variable-speed air-handler fans and chilled water pumps	7.2.5	3	1,1 (see note 2)	Medium term
Sizing/re-sizing of cooling systems and the configuration of redundancy	7.2.5	10	3,8 (see note 2)	Short term
Domain: Power distribution systems				
HVDC	9.3.3	7	1,1 (see note 3)	Long term
High-efficiency power distribution units	9.3.4	5	0,8 (see note 3)	Long term
High-efficiency motors in fans and pumps	9.3.4	2	0,3 (see note 3)	Medium term
Improved efficiency UPS units	9.3.4	10	1,5 (see note 3)	Medium term
Rotary-based UPS units	9.3.3	8	1,2 (see note 3)	Medium term
Capacity management	9.3.4	3	0,5 (see note 3)	Short term
NOTE 1: The overall saving shown assumes that IT infrastructure initially constitutes 37 % of overall energy consumption. NOTE 2: The overall saving shown assumes that environmental control systems initially constitute 45 % of overall energy consumption. NOTE 3: The overall saving shown assumes that power distribution systems initially constitute 15 % of overall energy consumption.				

History

Document history		
V1.1.1	October 2009	Publication