

2011 Best Practices

for the EU Code of Conduct on Data Centres

1 Document Information

1.1 Version History

Version ¹	Description	Version Updates	Date
3.0.1	2011 Release fix	Minor correction in section 2.6 table	01 March 2011
3.0.0	2011 Release	None	28 February 2011
2.1.1	Draft of 2011 Update	Incorporated feedback from Mark Acton, Dominique Roche, amended 4.1.9	February 2011
2.1.0	Draft of 2011 update	Incorporated feedback from Magnus Herrlin, Intel, Steve Strutt, TelecityGroup, Roger Schmidt (ASHRAE)	16 December 2010
2.0.0	2010 Release	Consolidated final comments	19 November 2009
1.1.6	Draft of 2009 Update	Updated with feedback after 2009 Ispra meeting	30 October 2009
1.1.1	Draft of 2009 Update	Updated with feedback from 2009 Ispra meeting	15 October 2009
1.1.0	Draft of 2009 Update	Incorporated changes from HP (EYP), added practice numbers, added "Practices to become minimum expected" section. Updated multiple practices see change tracking.	16 September 2009
1.0.0	First Release	Changes accepted	23 October 2008
0.2.9	Final Release Candidate	Incorporated comments in response to final review version 0.2.8	19 August 2008
0.2.8	Updated Review Copy	Incorporated further comments on storage, rewrite of Introduction	28 July 2008
0.2.7	Updated Review Copy	Incorporated further comments from Harkeeret and others	22 July 2008
0.2.6	Updated Review Copy	Restructured the document as suggested by and incorporated comments from Sophia and Robert of EYP	07 July 2008
0.2.5	Updated Review Copy	Incorporated further comments	06 July 2008
0.2.2	Updated Review Copy	Incorporated further comments received since April CoC meeting	21 April 2008
0.2.1	Updated Review Copy	Incorporated comments received during April London CoC meeting	15 April 2008
0.1.3	Updated Review Copy	Incorporated comments from Jan Viegand of DEST	02 April 2008
0.1.2	Updated Review Copy	Chloride Power, Keysource, Stulz comments incorporated	02 April 2008
0.1.1	Updated Review Copy	Telecom Italia, EEIOCG Energy Efficiency InterOperator Common Group and Dell comments and additions incorporated	01 April 2008
0.1.0	CoC Review Copy	Review pass through document	26 March 2008

¹ Version numbering standard, integer number for release version, first decimal point for major revisions, second decimal point for minor revisions

0.0.2	Addition of new content	Restructured Cooling Section, added IT sections	22 March 2008
0.0.1	First Version for Internal Review	First version, no updates	18 March 2008

Version History**1.2 Release History**

Version	Description	Authoriser	Date
3.0.0	2011 Release	Liam Newcombe	28 February 2011
2.0.0	2010 Release	Liam Newcombe	19 November 2009
1.0.0	First Release	Liam Newcombe	23 October 2008

Release History

2 Introduction

This document is a companion to the EU Code of Conduct on Data Centres v2.0. This document provides the full list of identified best practices for data centre operators as referenced in the Code of Conduct.

2.1 Role of Best Practices

This Best Practice supplement to the Code of Conduct is provided as an education and reference document as part of the Code of Conduct to assist data centre operators in identifying and implementing measures to improve the energy efficiency of their data centres. A broad group of expert reviewers from operators, vendors, consultants, academics, professional and national bodies have contributed to and reviewed the best practices.

This best practice supplement is a full list of the identified and recognised data centre energy efficiency best practices within the Code of Conduct. The best practice list provides a common terminology and frame of reference for describing an energy efficiency practice, to assist Participants and Endorsers in avoiding doubt or confusion over terminology. Customers or suppliers of IT services may also find it useful to request or provide a list of Code of Conduct practices implemented in a data centre to assist in procurement of services that meet their environmental or sustainability standards.

2.2 Expected Minimum Practices

To help ensure that Participants to the Code of Conduct are recognised as having committed to a useful and substantial level of energy saving effort, a subset of the best practices are identified in this document as being the expected minimum level of energy saving activity for Participant status.

The less disruptive or intrusive of the practices are identified as being applied to the existing data centre and IT equipment, retrospectively where necessary. It is accepted that a number of the practices identified as expected are inappropriate or present an unnecessary burden when applied to an existing running data centre. These practices are identified as being expected either when new IT equipment or software is sourced and deployed or during a retrofit of the facility. These practices provide substantial benefits and are intended to achieve efficiency improvements through the natural churn of equipment and facilities.

All expected practices should be applied to any data centre constructed from 2011 onwards, specifically all practices marked as “Entire data centre”, “New software”, “New IT equipment” and “New build or retrofit” which are within the applicants control.

Practices are marked in the expected column as;

Category	Description
Entire Data Centre	Expected to be applied to all existing IT, Mechanical and Electrical equipment within the data centre
New Software	Expected during any new software install or upgrade
New IT Equipment	Expected for new or replacement IT equipment
New build or retrofit	Expected for any data centre built or undergoing a significant refit of the M&E equipment from 2010 onwards
Optional practices	Practices without a background colour are optional for participants

Note that existing IT equipment moved from another data centre is not expected to comply with the New IT Equipment or New Software practices. New or replacement IT equipment excludes the direct replacement of failed hardware with like for like as part of normal

operations. New software install or upgrade refers to major upgrades of software and not the application of service packs and patches in normal management and use.

Retrofit is intended to describe major disruptive works in the data centre which present the opportunity at little incremental cost to implement these additional practices. Examples of retrofit would be (a) when the power to the data floor is shut off and the IT equipment and racks removed it is expected that practice 5.1.1 Contained hot or cold aisle would be implemented (b) if the CRAC units are being upgraded or replaced it is expected that practice 5.5.1 Variable speed fans would be implemented as part of this change.

2.3 Application and Assessment

The best practices form part of the application and assessment for Participant status. This process is described in the main Code of Conduct document.

2.4 Value of Practices

Each practice has been assigned a qualitative value to indicate the level of benefit to be expected from an action and the relative priorities that should be applied to them. These values are from 1 to 5 with 5 indicating the maximum value. These values are not intended to be totalled to provide an overall 'operator score' and should not be mistaken for quantitative. This would require large scale data on the effects of each practice or technology which is not yet available as well as a complex system of scoring representing the combinational increase or reduction of individual practice values within that specific facility

2.5 Applicability of Expected Practices

It is understood that not all operators will be able to implement all of the expected practices in their facilities due to physical, logistical, planning or other constraints. In these instances an explanation of why the expected action is not applicable or practical should be provided in the "Reason why this practice cannot be implemented in this data centre" column in the reporting form, alternative best practices from the supplement may be identified as direct replacements if they result in similar energy savings.

2.6 Type of Applicant

Each applicant should identify the type of operator that best describes their activity within the data centre for which they are completing the form on the "Data Centre Information" tab as;

Type	Description
Operator	Operates the entire data centre from the physical building through to the consumption of the IT services delivered.
Colo provider	Operates the data centre for the primary purpose of selling space, power and cooling capacity to customers who will install and manage IT hardware.
Colo customer	Owns and manages IT equipment located in a data centre in which they purchase managed space, power and cooling capacity.
Managed service provider (MSP)	Owns and manages the data centre space, power, cooling, IT equipment and some level of software for the purpose of delivering IT services to customers. This would include traditional IT outsourcing.
Managed service provider in Colo	A managed service provider which purchases space, power or cooling in this data centre.

Table 2-1 Types of applicants

The type of operator serves two purposes, first it assists the secretariat in the assessment of an application and second it will be included in the listing for data centres which achieve participant status on the Code of Conduct website.

2.7 Applicants who do not control the entire data centre

It is understood that not all operators are responsible for all aspects of the IT environment defined within the best practices. This is not a barrier to Participant status but the operator should sign as a Participant and act as an Endorser for those practices outside of their control.

The following sections are included to provide guidance to operators with partial control of the data centre on which practices they are expected to Implement and which they are expected to Endorse.

It is suggested that you download the application form, select your type of operator and then your areas of responsibility whilst reading this document to understand how this categorisation guides practice implementation.

2.7.1 Guidance to operators with partial control of the data centre

The best practice tab of the reporting form provides guidance for each of the minimum expected practices on whether these are considered to apply to each of these example types of operator, in which cases responsibility is to be shared and how that may be implemented. This may be found in the columns labelled "Guidance to operators with partial control of the data centre".

2.7.2 Areas of Responsibility

Operators' areas of responsibility are defined as;

Area	Description
Physical building	The building including security, location and maintenance.
Mechanical and electrical plant	The selection, installation, configuration, maintenance and management of the mechanical and electrical plant.
Data floor	The installation, configuration, maintenance and management of the main data floor where IT equipment is installed. This includes the floor (raised in some cases), positioning of CRAC units and PDUs, basic layout of cabling systems (under floor or overhead).
Racks	The installation, configuration, maintenance and management of the racks into which rack mount IT equipment is installed.
IT equipment	The selection, installation, configuration, maintenance and management of the physical IT equipment.
Operating System / Virtualisation	The selection, installation, configuration, maintenance and management of the Operating System and virtualisation (both client and hypervisor) software installed on the IT equipment. This includes monitoring clients, hardware management agents etc.
Software	The selection, installation, configuration, maintenance and management of the application software installed on the IT equipment.
Business practices	The determination and communication of the business requirements for the data centre including the importance of systems, reliability availability and maintainability specifications and data management processes.

Table 2-2 Areas of responsibility

An example of Participant responsibility would be a collocation provider who does not control the IT equipment should actively endorse the practices relating to IT equipment to their customers. This might include the provision of services to assist customers in adopting those practices. Equally an IT operator using collocation should request their collocation provider to implement the practices relating to the facility.

An applicant should mark their responsibility for each of these areas on the "Data Centre Information" tab of the reporting form as "Y", "N", or "Partial".

Note that these boundaries of responsibility do not apply within organisations. An applicant is considered to control an area if a parent, subsidiary or group company owns or controls the area. For example, if another division of the same group of companies operates a colo facility within which the applicant operates equipment as a service provider this is considered to be a managed service provider with responsibility for the physical building, mechanical and electrical plant, data floor and racks, not a managed service provider in colo.

2.7.3 Implement or Endorse

Each operator should determine which of the practices apply to them based on their areas of responsibility. The table below provides an overview for common types of Participant;

	Operator	Colo provider	Colo customer	MSP in Colo	MSP
Physical building	Implement	Implement	Endorse	Endorse	Implement
Mechanical & electrical plant	Implement	Implement	Endorse	Endorse	Implement
Data floor and air flow	Implement	Implement & Endorse	Implement & Endorse	Implement	Implement
Racks and rack air flow	Implement	Implement & Endorse	Implement & Endorse	Implement	Implement
IT equipment	Implement	Endorse	Implement	Implement	Implement
Operating System & Virtualisation	Implement	Endorse	Implement	Implement	Implement
Software	Implement	Endorse	Implement	Implement & Endorse	Implement & Endorse
Business practices	Implement	Endorse	Implement	Endorse	Endorse

Table 2-3 Areas of responsibility for common applicant types

The reporting form contains logic to assist Applicants in determining which of the Expected Practices they should Endorse and which they should Implement based upon the areas of the data centre that are within their control. An Applicant should select “Y”, “N” or “Partial” for each of the identified areas of control on the “Data Centre Information” tab of the reporting form. The form will then mark each Expected Practice with “I” Implement, “E” Endorse or “I & E” Implement and Endorse action in the columns labelled “Expected Status based on responsibility areas” to provide guidance to the applicant.

There are many instances where the responsibility for a practice will need to be shared between supplier and customer, for example the installation of IT equipment in the correct orientation in a hot / cold aisle layout data centre. In this case both parties should Implement the practice themselves and Endorse it to the other party(ies).

2.7.4 Marking practice status

An applicant should mark each practice in the application form as one of;

Mark	Description
No mark	Not implemented, a reason why the practice is not applicable to the applicant should be provided if this is an expected practice.
Committed Date	Not yet implemented but a program is in place to implement the practice by the specified date. This should be within 36 months of the application date.
I	Implemented practice.
E	Endorsed practice, this practice cannot be implemented by the applicant as it is outside their area of responsibility but is endorsed to their suppliers or customers.
I & E	This practice is partially within the control of the applicant. The applicant has implemented the practice as far as practical and endorsed the practice to their customers or suppliers. Include a description of the actions taken to endorse the practice.

Table 2-4 Marking practices in the application form

3 Data Centre Utilisation, Management and Planning

It is important to develop a holistic strategy and management approach to the data centre. This will enable the Participant to effectively deliver reliability, economic, utilisation and environmental benefits.

3.1 Involvement of Organisational Groups

Ineffective communication between the disciplines working in the data centre is a major driver of inefficiency as well as capacity and reliability issues.

No	Name	Description	Expected	Value
3.1.1	Group involvement	Establish an approval board containing representatives from all disciplines (software, IT, M&E). Require the approval of this group for any significant decision to ensure that the impacts of the decision have been properly understood and an effective solution reached. For example, this could include the definition of standard IT hardware lists through considering the M&E implications of different types of hardware. This group could be seen as the functional equivalent of a change board.	Entire Data Centre	5

3.2 General Policies

These policies apply to all aspects of the data centre and its operation.

No	Name	Description	Expected	Value
3.2.1	Consider the embedded energy in devices	Carry out an audit of existing equipment to maximise any unused existing capability by ensuring that all areas of optimisation, consolidation and aggregation are identified prior to new material investment.	Entire Data Centre	3

3.3 Resilience Level and Provisioning

One of the most significant sources of inefficiency in data centres is the over provisioning of space, power or cooling and the facilities being run at part capacity. Monolithic, as opposed to modular design of facilities also represents a significant and frequently unnecessary capital expenditure. Further, as the level of resilience of the data centre increases the inefficiencies due to fixed overheads increase and this is compounded by poor utilisation.

No	Name	Description	Expected	Value
3.3.1	Build resilience to business requirements	Only the level of resilience and therefore availability actually justified by business requirements and impact analysis should be built, or purchased in the case of a collocation customer. 2N infrastructures are frequently unnecessary and inappropriate. If only a single level of resilience is available in the data centre an increased resilience or availability for critical services can be obtained by splitting the IT platform across multiple sites.	New build or retrofit	3
3.3.2	Consider multiple levels of resilience	It is possible to build a single data centre to provide multiple levels of power and cooling resilience to different floor areas. Many co-location providers already deliver this, for example, optional 'grey' power feeds without UPS or generator back up.	New build or retrofit	3
3.3.3	Lean provisioning of power and cooling for a maximum of 18 months of data floor capacity	The provisioning of excess power and cooling capacity in the data centre drives substantial fixed losses and is unnecessary. Planning a data centre for modular (scalable) expansion and then building out this capacity in a rolling program of deployments is more efficient. This also allows the technology 'generation' of the IT equipment and supporting M&E infrastructure to be matched, improving both efficiency and the ability to respond to business requirements.	New build or retrofit	3
3.3.4	Design to maximise the part load efficiency once provisioned	The design of all areas of the data centre should be maximise the achieved efficiency of the facility under partial fill and variable IT electrical load. This is in addition to one off modular provisioning and considers the response of the infrastructure to dynamic loads. e.g. Variable Frequency (or speed) Drive for pumps and fan units.	New build or retrofit	3
3.3.5	Design effective resilience	Utilise appropriate levels of resilience at the data centre, IT equipment, software and network levels to achieve the required service resilience and availability. High resilience at the physical level is rarely an effective overall solution	No	4

4 IT Equipment and Services

The IT equipment creates the demand for power and cooling in the data centre, any reductions in power and cooling used by or provisioned for the IT equipment will have magnified effects at the utility energy supply.

Note that the specifications of IT equipment operating temperature and humidity ranges in this section do not indicate that the data floor should be immediately operated at the upper bound of these ranges, this is addressed in section 5.3. The purpose of the equipment environmental specifications in this section is to ensure that new equipment is capable of operating under the wider ranges of temperature and humidity thus allowing greater flexibility in operating temperature and humidity to the operator.

4.1 Selection and Deployment of New IT Equipment

Once IT equipment is purchased and installed in the data centre it typically spends several years in the data centre consuming power and creating heat. The appropriate selection of hardware and deployment methods can provide significant long term savings.

No	Name	Description	Expected	Value
4.1.1	Multiple tender for IT hardware – Power	Include the Energy efficiency performance of the IT device as a high priority decision factor in the tender process. This may be through the use of Energy Star or SPECPower type standard metrics or through application or deployment specific user metrics more closely aligned to the target environment which may include service level or reliability components. The power consumption of the device at the expected utilisation or applied workload should be considered in addition to peak performance per Watt figures.	New IT Equipment	5
4.1.2	Multiple tender for IT hardware – Basic operating temperature and humidity range at equipment intake	Include the operating temperature and humidity ranges at equipment intake of new equipment as high priority decision factors in the tender process. The minimum range, at the air intake to servers, is 18°C-27°C and 5.5°C dew point up to 15°C dew point & 60% RH. The current relevant standard is the ASHRAE Recommended range for Class 1 Data Centers as described by ASHRAE in “2008 ASHRAE Environmental Guidelines for Datacom Equipment”.	New IT Equipment	4

4.1.3	Multiple tender for IT hardware – Extended operating temperature and humidity range	<p>Starting 2012 new IT equipment should be able to withstand the air inlet temperature and relative humidity ranges of 5°C to 40°C and 5% to 80% RH, non-condensing respectively, and under exceptional conditions up to +45°C as described in ETSI EN 300 019 Class 3.1.</p> <p>All vendors should indicate the maximum allowable temperature and humidity for all equipment to maximise the efficiency opportunities in refrigeration and free cooling.</p> <p>It should be noted that where equipment with differing environmental requirements is not segregated, the equipment with the more restrictive temperature range will influence the cooling conditions and corresponding energy consumption for all of the IT Equipment.</p> <p>From 40°C to 45°C intake temperature it is acceptable for equipment to implement performance reduction mechanisms to continue delivering the intended service at lower speed whilst preventing damage. These mechanisms should not reduce performance below 80% of the nominal for that device. Where such performance reduction mechanisms are used a clear description of the operating parameters and performance impact should be provided.</p>	New IT Equipment from 2012	5
4.1.4	Select equipment suitable for the data centre – Power density	<p>Select and deploy equipment at the design power density (per rack or sq m) of the data centre to avoid running the cooling system outside design parameters.</p> <p>Note that increasing power density may create cooling and air flow management problems reducing both capacity and efficiency. Power and cooling need to be considered as capacity constraints in addition to physical space.</p>	No	3
4.1.5	Select equipment suitable for the data centre - Air flow direction	<p>When selecting equipment for installation into racks ensure that the air flow direction matches the air flow design for that area. This is commonly front to rear or front to top.</p> <p>If the equipment uses a different air flow direction to that defined for the area into which it is installed (such as right to left when the rack is intended to be front to back) it should only be used with a correction mechanism such as ducts, or special racks that divert the air flow to the defined direction.</p> <p>Uncorrected equipment with non standard air flow will compromise the air flow management of the data centre and therefore restrict temperature set points. It is possible to mitigate this issue by segregating such equipment as per practice 5.1.10.</p>	New IT Equipment	4
4.1.6	Enable power management features	Formally change the deployment process to include the enabling of power management features on IT hardware as it is deployed. This includes BIOS, operating system and driver settings.	New IT Equipment	5

4.1.7	Provision to the as configured power	Provision power and cooling only to the as-configured power draw capability of the equipment, not the PSU or nameplate rating. Note that this may require changes to the provisioning if the IT equipment is upgraded internally.	New IT Equipment	3
4.1.8	Energy Star hardware	The Energy Star Labelling programs for IT equipment should be used as a guide to server selection where and when available for that class of equipment. Operators who are able to determine the in use energy efficiency of hardware through more advanced or effective analysis should select the most efficient equipment for their scenario.	No	3
4.1.9	Energy & temperature reporting hardware	<p>Select equipment with power and inlet temperature reporting capabilities, preferably reporting energy used as a counter in addition to power as a gauge. Where applicable, industry standard reporting approaches should be used such as IPMI, DCMI and SMASH.</p> <p>To assist in the implementation of temperature and energy monitoring across a broad range of data centres all devices with an IP interface should support one of;</p> <ul style="list-style-type: none"> • SNMP polling of inlet temperature and power draw. Note that event based SNMP traps and SNMP configuration are not required • IPMI polling of inlet temperature and power draw (subject to inlet temperature being included as per IPMI 2.0 rev 4) • An interface protocol which the operators' existing monitoring platform is able to retrieve inlet temperature and power draw data from without the purchase of additional licenses from the equipment vendor <p>The intent of this practice is to provide energy and environmental monitoring of the data centre through normal equipment churn.</p>	No	3
4.1.10	Control of equipment energy use	Select equipment which provides mechanisms to allow the external control of its energy use. An example of this would be the ability to externally restrict a server's maximum energy use or trigger the shutdown of components, entire systems or sub-systems	No	5
4.1.11	Select free standing equipment suitable for the data centre – Air flow direction	<p>When selecting equipment which is free standing or supplied in custom racks the air flow direction of the enclosures should match the air flow design in that area of the data centre. This is commonly front to rear or front to top.</p> <p>Specifically the equipment should match the hot / cold aisle layout or containment scheme implemented in the facility.</p> <p>Uncorrected equipment with non standard air flow will compromise the air flow management of the data centre and therefore restrict temperature set points. It is possible to mitigate this compromise by segregating such equipment as per practice 5.1.10</p>	New IT Equipment	4

4.2 Deployment of New IT Services

The service architecture, software and deployment of IT services have an impact at least as great as that of the IT hardware.

No	Name	Description	Expected	Value
4.2.1	Deploy using Grid and Virtualisation technologies	Processes should be put in place to require senior business approval for any new service that requires dedicated hardware and will not run on a resource sharing platform. This applies to servers, storage and networking aspects of the service.	New IT Equipment	5
4.2.2	Reduce IT hardware resilience level	Determine the business impact of service incidents for each deployed service and deploy only the level of hardware resilience actually justified.	New IT Equipment	4
4.2.3	Reduce hot / cold standby equipment	Determine the business impact of service incidents for each IT service and deploy only the level of Business Continuity / Disaster Recovery standby IT equipment and resilience that is actually justified by the business impact.	New IT Equipment	4
4.2.4	Select efficient software	Make the energy use performance of the software a primary selection factor. Whilst forecasting and measurement tools and methods are still being developed, approximations can be used such as the (under load) power draw of the hardware required to meet performance and availability targets. This is an extension of existing capacity planning and benchmarking processes. See "Further development of software efficiency definitions" in section 11.	New Software	4
4.2.5	Develop efficient software	Make the energy use performance of the software a major success factor of the project. Whilst forecasting and measurement tools and methods are still being developed approximations, can be used such as the (under load) power draw of the hardware required to meet performance and availability targets. This is an extension of existing capacity planning and benchmarking processes. Performance optimisation should not be seen as a low impact area to reduce the project budget. See "Further development of software efficiency definitions" in section 11.	New Software	4
4.2.6	Incentives to develop efficient software	If outsourcing software development then include the energy use of the software in the bonus / penalty clauses of the contract. Whilst forecasting and measurement tools and methods are still being developed approximations, can be used such as the (under load) power draw of the hardware required to meet performance and availability targets. This is an extension of existing capacity planning and benchmarking processes. Performance optimisation should not be seen as a low impact area to reduce the project budget. See "Further development of software efficiency definitions" in section 11.	No	4

4.2.7	Eliminate traditional 2N hardware clusters	Determine the business impact of short service incidents for each deployed service and replace traditional active / passive server hardware clusters with fast recovery approaches such as restarting virtual machines elsewhere. (this does not refer to grid or High Performance Compute clusters)	No	4
-------	--	--	----	---

4.3 Management of Existing IT Equipment and Services

It is common to focus on new services and equipment being installed into the data centre but there are also substantial opportunities to achieve energy and cost reductions from within the existing service and physical estate.

No	Name	Description	Expected	Value
4.3.1	Audit existing physical and service estate	Audit the existing physical and logical estate to establish what equipment is in place and what service(s) it delivers. Consider the implementation of an ITIL type Configuration Management Data base and Service Catalogue.	No	4
4.3.2	Decommission unused services	Completely decommission and remove, the supporting hardware for unused services	Entire Data Centre	5
4.3.3	Virtualise and archive legacy services	Servers which cannot be decommissioned for compliance or other reasons but which are not used on a regular basis should be virtualised and then the disk images archived to a low power media. These services can then be brought online when actually required	No	5
4.3.4	Consolidation of existing services	Existing services that do not achieve high utilisation of their hardware should be consolidated through the use of resource sharing technologies to improve the use of physical resources. This applies to servers, storage and networking devices.	No	5
4.3.5	Decommission low business value services	Identify services whose business value is low and does not justify the financial or environmental cost, decommission or archive these services	No	4
4.3.6	Shut down idle equipment	Servers, networking and storage equipment that is idle for significant time and cannot be virtualised and archived should be shut down or put into a low power 'sleep' state. It may be necessary to validate the ability of legacy applications and hardware to survive these state changes without loss of function or reliability.	No	3
4.3.7	Control of system energy use	Consider resource management systems capable of analysing and optimising where, when and how IT workloads are executed and their consequent energy use. This may include technologies that allow remote deployment or delayed execution of jobs or the movement of jobs within the infrastructure to enable shutdown of components, entire systems or sub-systems. The desired outcome is to provide the ability to limit localised heat output or constrain system power draw to a fixed limit, at a data centre, row, rack or sub-DC level	No	4

4.4 Data Management

Storage is a major growth area in both cost and energy consumption within the data centre. It is generally recognised that a significant proportion of the data stored is either unnecessary or duplicated nor requires high performance access and that this represents an organisational challenge. Some sectors have a particular issue due to very broad and non specific data retention directions from governments or regulating bodies. Where there is little structure to the data storage, implementation of these regulations can cause large volumes of data not required by the regulations to be unnecessarily heavily protected and archived.

No	Name	Description	Expected	Value
4.4.1	Data management policy	Develop a data management policy to define which data should be kept, for how long and at what level of protection. Communicate the policy to users and enforce. Particular care should be taken to understand the impact of any data retention requirements,	Entire Data Centre	3
4.4.2	Separate user logical data storage areas by retention and protection policy	Provide users with multiple data storage areas which are clearly identified by their retention policy and level of data protection. Communicate this policy to users to enable them to store data in an area which matches the required levels of protection and retention. This is particularly valuable where strong retention requirements exist as it allows data subject to those requirements to be separated at source presenting substantial opportunities for cost and energy savings. Where possible automate the application of these policies.	No	3
4.4.3	Separate physical data storage areas by protection and performance requirements	Create a tiered storage environment utilising multiple media types delivering the required combinations of performance, capacity and resilience. Implement clear guidelines on usage of storage tiers with defined SLAs for performance and availability. Consider a tiered charging model based on usage at each tier.	No	4
4.4.4	Select lower power storage devices	When selecting storage hardware evaluate the energy efficiency in terms of the service delivered per Watt between options. This may be deployment specific and should include the achieved performance and storage volume per Watt as well as additional factors where appropriate, such as the achieved levels of data protection, performance availability and recovery capability required to meet the business service level requirements defined in the data management policy. Evaluate both the in use power draw and the peak power of the storage device(s) as configured, both impact per device cost and energy consumption through provisioning.	No	3
4.4.5	Reduce total data volume	Implement an effective data identification and management policy and process to reduce the total volume of data stored. Consider implementing 'clean up days' where users delete unnecessary data from storage.	No	4

4.4.6	Reduce total storage volume	Implement the data management policy to reduce the number of copies of data, both logical and physical (mirrors). Implement storage subsystem space saving features, such as space efficient snapshots / copies or compression. Implement storage subsystem thin provisioning features where possible.	No	4
-------	-----------------------------	--	----	---

5 Cooling

Cooling of the Data Centre is frequently the largest energy loss in the facility and as such represents a significant opportunity to improve efficiency.

5.1 Air Flow Management and Design

The objective of air flow management is to minimise bypass air, which returns to the CRAC units without performing cooling and the resultant recirculation and mixing of cool and hot air increasing equipment intake temperatures. To compensate, CRAC unit air supply temperatures are frequently reduced or air flow volumes increased, which has an energy penalty. Addressing these issues will deliver more uniform equipment inlet temperatures and allow set points to be increased (with the associated energy savings) without the risk of equipment overheating. Implementation of air management actions alone does not result in an energy saving – they are enablers which need to be tackled before set points can be raised.

No	Name	Description	Expected	Value
5.1.1	Design – Contained hot or cold air	<p>There are a number of design concepts whose basic intent is to contain and separate the cold air from the heated return air on the data floor;</p> <ul style="list-style-type: none"> • Hot aisle containment • Cold aisle containment • Contained rack supply, room return • Room supply, Contained rack return, (inc. rack chimneys) • Contained rack supply, Contained rack return <p>This action is expected for air cooled facilities over 1kW per square meter power density. Note that the in rack cooling options are only considered to be containment where the entire data floor area is cooled in rack, not in mixed environments where they return cooled air for remix with other air flow.</p>	New build or retrofit	5
5.1.2	Rack air flow management – Blanking Plates	Installation of blanking plates where there is no equipment to reduce cold air passing through gaps in the rack. This also reduces air heated by one device being ingested by another device, increasing intake temperature and reducing efficiency.	Entire Data Centre	3
5.1.3	Rack air flow management – Other openings	<p>Installation of aperture brushes (draught excluders) or cover plates to cover all air leakage opportunities in each rack. This includes;</p> <ul style="list-style-type: none"> • floor openings at the base of the rack • Gaps at the sides, top and bottom of the rack between equipment or mounting rails and the perimeter of the rack 	New build or retrofit	3

5.1.4	Raised floor air flow management	Close all unwanted apertures in the raised floor. Review placement and opening factors of vented tiles. Maintain unbroken rows of cabinets to prevent bypass air – where necessary fill with empty fully blanked racks. Managing unbroken rows is especially important in hot and cold aisle environments. Any opening between the aisles will degrade the separation of hot and cold air.	Entire Data Centre	3
5.1.5	Design – Return plenums	Consider the use of return plenums to return heated air from the IT equipment to the air conditioning units	No	3
5.1.6	Design – Contained hot or cold air – Retrofit	Where hot / cold aisle separation is already in use but there is no containment of hot or cold air it is possible to retrofit to provide basic separation.	No	3
5.1.7	Raised floor air flow management – Obstructions	Review the placement and level of obstruction created by cabling, cable trays and other structures in the air flow paths, these obstruct airflow and create turbulence, increasing the resistance and increasing the energy requirements of air movement. The use of overhead cabling trays for signalling can substantially reduce these losses.	No	2
5.1.8	Design – Hot / cold aisle	As the power densities and air flow volumes of IT equipment have increased it has become necessary to ensure that equipment shares an air flow direction, within the rack, in adjacent racks and across aisles. The hot / cold aisle concept aligns equipment air flow to create aisles between racks that are fed cold air from which all of the equipment draws intake air in conjunction with hot aisles with no cold air feed to which all equipment exhausts air.	New IT Equipment New build or retrofit	3
5.1.9	Design – Raised floor or suspended ceiling height	It is common to use the voids in the raised floor, suspended ceiling or both in a data centre to feed cold air to equipment or extract hot air from the equipment. Where they are used, increasing the size of these spaces can significantly reduce fan losses moving the air.	No	3
5.1.10	Equipment segregation	Deploy groups of equipment with substantially different environmental requirements and / or equipment airflow direction in a separate area. Where the equipment has different environmental requirements it is preferable to provide separate environmental controls. This objective of this practice is to address the issue of the data centre cooling plant settings being constrained by the IT equipment with the most restrictive environmental range or poor air flow control as this compromises the efficiency of the entire data centre.	No	3
5.1.11	Provide adequate free area on rack doors	Solid doors can be replaced (where doors are necessary) with partially perforated doors to ensure adequate cooling airflow which often impede the cooling airflow and may promote recirculation within the enclosed cabinet further increasing the equipment intake temperature.	New IT Equipment New build or retrofit	3

5.2 Cooling Management

The data centre is not a static system and the cooling systems should be tuned in response to changes in the facility thermal load.

No	Name	Description	Expected	Value
5.2.1	Scalable or modular installation and use of cooling equipment	Cooling plant should be installed in a modular fashion allowing operators to shut down unnecessary equipment. This should then be part of the review at each cooling load change. Design to maximise the part load efficiency as described in 3.3	No	3
5.2.2	Shut down unnecessary cooling equipment	If the facility is not yet fully populated or space has been cleared through consolidation non variable plant such as fixed speed fan CRAC units can be turned off in the empty areas. Note that this should not be applied in cases where operating more plant at lower load is more efficient, e.g. variable speed drive CRAC units.	No	3
5.2.3	Review of cooling before IT equipment changes	The availability of cooling including the placement and flow of vented tiles should be reviewed before each IT equipment change to optimise the use of cooling resources.	Entire Data Centre	2
5.2.4	Review of cooling strategy	Periodically review the IT equipment and cooling deployment against strategy.	Entire Data Centre	2
5.2.5	Review CRAC Settings	Ensure that CRAC units in occupied areas have appropriate and consistent temperature and relative humidity settings to avoid units working against each other. For example many CRAC units now have the option to connect their controls and run together when installed in the same area.	No	3
5.2.7	Dynamic control of building cooling	It is possible to implement control systems that take many factors including cooling load, data floor air temperature and external air temperature into account to optimise the cooling system, (e.g. chilled water loop temperature) in real time.	No	3
5.2.8	Effective regular maintenance of cooling plant	Effective regular maintenance of the cooling system is essential to maintain the design operating efficiency of the data centre. e.g. belt tension, condenser coil fouling (water or air side), evaporator fouling, filter changes etc.	No	2

5.3 Temperature and Humidity Settings

Facilities are often overcooled with air temperatures (and hence chilled water temperatures, where used) colder than necessary resulting in an energy penalty. Increasing the set range for humidity can substantially reduce humidifier loads. Reviewing and addressing air management issues as described in sections 5.1 and 5.2, is required before set points can be changed in order to avoid risk to operational continuity, expert advice should be sought before changing the environmental range for the facility. An increase in chilled water temperature set points provides enhanced efficiency for free cooling economisers and a reduction in compressor energy consumption. Unnecessary humidifier loads generated by water loop or evaporator temperatures below the working dew point causing dehumidification-humidification loops should be eliminated through adjustment of the lower humidity set point.

The specification of wider operating humidity and temperature ranges for the data floor should be performed in conjunction with changes in IT equipment procurement policy, over time narrow tolerance equipment will be naturally cycled out and replaced.

No	Name	Description	Expected	Value
5.3.1	Review and if possible raise target IT equipment intake air temperature	<p>Data Centres should be designed and operated <i>at their highest efficiency</i> within the current environmental range of 18°C to 27°C.</p> <p>The current, relevant standard is the ASHRAE <i>Recommended</i> range for Class 1 Data Centers, as described by ASHRAE in “2008 ASHRAE Environmental Guidelines for Datacom Equipment”. Operations in this range enable energy savings by reducing or eliminating overcooling. This range applies to legacy data centres with existing equipment. Note that other best practices for airflow management (containment, hot aisle/cold aisle, blanking plates, and sealing leaks) may need to be implemented at the same time to ensure successful operations.</p> <p>Note that the ASHRAE TC9.9 specification also provides for an allowable temperature range of 15°C to 32°C, see “2008 ASHRAE Environmental Guidelines for Datacom Equipment” for further information.</p> <p>Note that some IT equipment may exhibit significant increases in fan power consumption as intake temperature is increased. Validate that your IT equipment will not consume more energy than is saved in the cooling system.</p>	Entire Data Centre	4

5.3.2	Review and increase the working humidity range	<p>Reduce the lower humidity set point(s) of the data centre within the ASHRAE range (5.5°C dew point) to eliminate loop de-humidification and re-humidification.</p> <p>Review and if practical increase the upper humidity set point(s) of the data floor within current environmental range of 15°C dew point & 60% RH to decrease the humidity control loads within the facility.</p> <p>The current, relevant standard is the ASHRAE <i>Recommended</i> range for Class 1 Data Centers, as described by ASHRAE in “2008 ASHRAE Environmental Guidelines for Datacom Equipment”.</p> <p>Note that the ASHRAE TC9.9 specification also provides for an allowable operating humidity range of 20% to 80% RH, see “2008 ASHRAE Environmental Guidelines for Datacom Equipment” for further information.</p>	Entire Data Centre	4
5.3.3	Expanded IT equipment inlet environmental conditions (temperature and humidity)	Where appropriate and effective, Data Centres can be designed and operated within the air inlet temperature and relative humidity ranges of 5°C to 40°C and 5% to 80% RH, non-condensing respectively, and under exceptional conditions up to +45°C as described in ETSI EN 300 019, Class 3.1.	Starting 2012	5
5.3.4	Review set points of air and water temperatures	Once air management issues have been addressed and IT equipment target temperatures agreed these temperatures can be increased (using less energy) without increasing server inlet temperatures beyond acceptable levels. Note that some IT equipment may use more power under increased inlet temperatures.	Entire Data Centre	3
5.3.5	Review and if possible raise chilled water loop temperature	<p>Review and if possible increase the chilled water temperature set points to maximise the use of free cooling economisers and reduce compressor energy consumption.</p> <p>Where a DX system is used the evaporator temperatures should be reviewed.</p>	Entire Data Centre	4

5.4 Cooling Plant

The cooling plant typically represents the major part of the energy used in the cooling system. This is also the area with the greatest variation in technologies.

5.4.1 Free and Economised Cooling

Free or economised cooling designs use cool ambient conditions to meet part or all of the facilities cooling requirements hence compressor work for cooling is reduced or removed, which can result in significant energy reduction. Economised cooling can be retrofitted to some facilities. The opportunities for the utilisation of free cooling are increased in cooler climates and where increased temperature set points are used

No	Name	Description	Expected	Value
5.4.1.1	Direct air free cooling	External air is used to cool the facility. Chiller systems are present to deal with humidity and high external temperatures if necessary. Exhaust air is re-circulated and mixed with intake air to avoid unnecessary humidification / dehumidification loads.	No	5
5.4.1.2	Indirect air free cooling	Re circulated air within the facility is primarily passed through a heat exchanger against external air to remove heat to the atmosphere.	No	5
5.4.1.3	Direct water free cooling	Chilled water cooled by the external ambient air via a free cooling coil. This may be achieved by dry (/adiabatic) coolers or by evaporative assistance through spray onto the dry (/adiabatic) coolers.	No	4
5.4.1.4	Indirect water free cooling	Chilled water is cooled by the external ambient conditions via a heat exchanger which is used between the condenser and chilled water circuits. This may be achieved by dry (/adiabatic) coolers, evaporative assistance through spray onto the dry (/adiabatic) coolers or cooling towers.	No	4
5.4.1.5	Sorption cooling (absorption / adsorption)	Waste heat produced as a by-product of power generation or other processes close to the data centre is used to power the cooling system in place of electricity. This is frequently part of a Tri Gen combined cooling heat and power system. These systems should be assessed for viability over their full life time against an optimised economised cooling plant over the realistic operating range of load and climate and considering likely efficiency changes in the utility supply.	No	2

5.4.2 High Efficiency Cooling Plant

The next preference cooling technology is the use of high efficiency cooling plant. Designs should operate efficiently at system level and employ efficient components. This demands an effective control strategy which optimises efficient operation, without compromising reliability.

No	Name	Description	Expected	Value
5.4.2.2	Chillers with high COP	Make the Coefficient Of Performance of chiller systems through their likely working range a high priority decision factor during procurement of new plant.	New build or retrofit	3
5.4.2.3	Cooling system operating temperatures	Evaluate the opportunity to decrease condensing temperature or increase evaporating temperature; reducing delta T between these temperatures means less work is required in cooling cycle hence improved efficiency. These temperatures are dependent on required internal air temperatures (see Temperature and Humidity Settings).	Entire Data Centre	3
5.4.2.4	Efficient part load operation	Optimise the facility for the partial load it will experience for most of operational time rather than max load. e.g. sequence chillers, operate cooling towers with shared load for increased heat exchange area	New build or retrofit	3
5.4.2.5	Variable speed drives for compressors, pumps and fans	Reduced energy consumption for these components in the part load condition where they operate for much of the time.	No	2
5.4.2.6	Select systems which facilitate the use of economisers	Select systems which facilitate the use of cooling economisers. In some buildings it may be possible to use air side economisers others may not have sufficient available space and may require a chilled liquid cooling system to allow the effective use of economised cooling.	No	4

5.6 Computer Room Air Conditioners

The second major component of most cooling systems is the air conditioner units within the computer room. The computer room side of the chiller plant is frequently poorly designed and poorly optimised in older facilities.

No	Name	Description	Expected	Value
5.6.1	Variable Speed Fans	<p>Many old CRAC units operate fixed speed fans which consume substantial power and obstruct attempts to manage the data floor temperature.</p> <p>Variable speed fans are particularly effective where there is a high level of redundancy in the cooling system, low utilisation of the facility or highly variable IT electrical load. These fans may be controlled by factors such as the return air temperature or the chilled air plenum pressure.</p> <p>Note that CRAC units with fixed speed compressors have minimum flow requirements which constrain the minimum operating load.</p>	New build or retrofit	4
5.6.2	Control on CRAC unit supply air temperature	Controlling on supply temperature ensures the server supply air (key temperature to control) is satisfactory without possible over cooling of air which may result when controlling on return temperature (where sensor location may impact)	No	2
5.6.3	Run variable speed CRAC units in parallel	It is possible to achieve efficiency gains by running CRAC units with variable speed fans in parallel to reduce the total electrical power necessary to achieve the required air movement as electrical power is not linear with air flow. Care should be taken to understand any new failure modes or single points of failure that may be introduced by any additional control system.	No	4
5.6.4	Direct liquid cooling of IT devices	<p>In place of air cooling it is possible to directly liquid cool part or all of some IT devices. This can provide a more efficient thermal circuit and allow the coolant liquid system temperature to be substantially higher, further driving efficiency, allowing for increased or exclusive use of free cooling or heat re use.</p> <p>Note that this practice applies to devices which deliver cooling fluid directly to the heat removal system of the components such as water cooled heat sinks or heat pipes and not the delivery of cooling liquid to an internal mechanical refrigeration plant or in chassis air cooling systems.</p>	No	4
5.6.5	Sequencing of CRAC units	<p>In the absence of variable speed fans it is possible to turn entire CRAC units on and off to manage the overall air flow volumes.</p> <p>This can be effective where there is a high level of redundancy in the cooling system, low utilisation of the facility or highly variable IT electrical load.</p>	No	2

5.7 Reuse of Data Centre Waste Heat

Data Centres produce significant quantities of waste heat, whilst this is typically at a relatively low temperature there are some applications for reuse of this energy. As IT equipment utilisation is increased through consolidation and virtualisation the exhaust temperature is likely to increase which will provide greater opportunity for waste heat to be re-used. Directly liquid cooled IT equipment is likely to provide a further improvement in the return temperature of coolant.

No	Name	Description	Expected	Value
5.7.1	Waste heat re-use	It may be possible to provide low grade heating to industrial space or to other targets such as swimming pools directly from heat rejected from the data centre. This can ameliorate an energy use elsewhere, reducing the total energy use of the data centre and the client of the waste heat.	No	3
5.7.2	Heat pump assisted waste heat re-use	Where it is not possible to directly re use the waste heat from the data centre due to the temperature being too low it can still be economic to use additional heat pumps to raise the temperature to a useful point. This can supply office, district and other heating.	No	2

6 Data Centre Power Equipment

The other major part of the facility infrastructure is the power conditioning and delivery system. This normally includes uninterruptible power supplies, power distribution units and cabling but may also include backup generators and other equipment.

6.1 Selection and Deployment of New Power Equipment

Power delivery equipment has a substantial impact upon the efficiency of the data centre and tends to stay in operation for many years once installed. Careful selection of the power equipment at design time can deliver substantial savings through the lifetime of the facility.

No	Name	Description	Expected	Value
6.1.1	Modular UPS deployment	It is now possible to purchase modular (scalable) UPS systems across a broad range of power delivery capacities. Physical installation, transformers and cabling are prepared to meet the design electrical load of the facility but the sources of inefficiency (such switching units and batteries) are installed, as required, in modular units. This substantially reduces both the capital cost and the fixed overhead losses of these systems. In low power environments these may be frames with plug in modules whilst in larger environments these are more likely to be entire UPS units.	New build or retrofit	3
6.1.2	High efficiency UPS	High efficiency UPS systems should be selected, of any technology including electronic or rotary to meet site requirements.	New build or retrofit	3
6.1.3	Use efficient UPS operating modes	UPS should be deployed in their most efficient operating modes such as line interactive. Technologies such as Rotary and High Voltage DC (direct current) can also show improved efficiency as there is no dual conversion requirement.	New build or retrofit	2
6.1.4	Code of Conduct compliant UPS	Select UPS systems compliant with the EU Code of Conduct for UPS where that UPS technology is included. Rotary UPS are not included in the UPS Code of Conduct. Note that this does not suggest that rotary UPS should not be used, simply that rotary technology is not covered by the UPS Code of Conduct.	No	2

6.2 Management of Existing Power Equipment

No	Name	Description	Mandatory	Value
6.2.1	Reduce engine-generator heater	When using engine heaters to keep generators ready for rapid starts, consider reducing the engine heater set-point. Block heaters for the Standby Generators should be controlled to only	No	2

	temperature set-point	operate when the temperature conditions warrant it		
--	-----------------------	--	--	--

7 Other Data Centre Equipment

Energy is also used in the non data floor areas of the facility in office and storage spaces. Energy efficiency in non-data centre areas should be optimised based on relevant building standards, such as relevant EU standards, LEED, BREEAM etc.

7.1 General practices

These general practices apply to the data floor and may be extended to the remainder of the building if no sustainable building standard is in use.

No	Name	Description	Expected	Value
7.1.1	Turn off Lights	Lights should be turned off, preferably automatically whenever areas of the building are unoccupied, for example switches which turn off lighting a specified time after manual activation. Motion detector activated lighting is generally sufficient to support security camera systems.	Entire Data Centre	1
7.1.2	Low energy lighting	Low energy lighting systems should be used in the data centre.	New build or retrofit	1

8 Data Centre Building

The location and physical layout of the data centre building is important to achieving flexibility and efficiency. Technologies such as fresh air cooling require significant physical plant space and air duct space that may not be available in an existing building.

8.1 Building Physical Layout

The physical layout of the building can present fundamental constraints on the applicable technologies and achievable efficiencies.

No	Name	Description	Expected	Value
8.1.1	Locate M&E plant outside the cooled area	Heat generating Mechanical and Electrical plant should be located outside the cooled areas of the data centre wherever possible to reduce the loading on the data centre cooling plant.	No	2
8.1.2	Select a building with sufficient ceiling height	Insufficient ceiling height will obstruct the use of efficient air cooling technologies such as raised floor, suspended ceiling or ducts in the data centre.	No	3
8.1.3	Optimise orientation of the data centre	Optimise the layout and orientation of the building to reduce the insolation heat loads and optimise the efficiency of heat transfer.	No	1
8.1.4	Facilitate the use of economisers	The physical layout of the building should not obstruct the use of economisers (either air or water)	No	3
8.1.5	Location and orientation of plant equipment	Cooling equipment, particularly dry (/adiabatic) coolers should be located in an area of free air movement to avoid trapping it in a local hot spot. Ideally this equipment should also be located in a position on the site where the waste heat does not affect other buildings and create further demand for air conditioning.	No	2
8.1.6	Minimise direct solar heating	Minimise solar heating of the cooled areas of the data centre by providing shade or increasing the albedo (reflectivity) of the building through the use of light coloured roof and wall surfaces. Shade may be constructed, provided by trees or "green roof" systems.	No	2

8.2 Building Geographic Location

Whilst some operators may have no choice of the geographic location for a data centre it nevertheless impacts achievable efficiency, primarily through the impact of external climate.

No	Name	Description	Expected	Value
8.2.1	Locate the Data Centre where waste heat can be reused	Locating the data centre where there are available uses for waste heat can save substantial energy. Heat recovery can be used to heat office or industrial space, hydroponic farming and even swimming pools.	No	2
8.2.2	Locate the Data Centre in an area of low ambient temperature	Free and economised cooling technologies are more effective in areas of low ambient external temperature and or humidity. Note that most temperature climates including much of Northern, Western and Central Europe present significant opportunity for economised cooling.	No	3
8.2.3	Avoid locating the data centre in high ambient humidity areas	Free cooling is particularly impacted by high external humidity as dehumidification becomes necessary, many economiser technologies are also less effective.	No	1
8.2.4	Locate near a source of free cooling	Locating the data centre near a source of free cooling such as a river subject to local environmental regulation.	No	3
8.2.5	Co-locate with power source	Locating the data centre close to the power generating plant can reduce transmission losses and provide the opportunity to operate sorption chillers from power source waste heat.	No	2

9 Monitoring

The development and implementation of an energy monitoring and reporting management strategy is core to operating an efficient data centre.

9.1 Energy Use and Environmental Measurement

Most data centres currently have little or no energy use or environmental measurement capability; many do not even have a separate utility meter or bill. The ability to measure energy use and factors impacting energy use is a prerequisite to identifying and justifying improvements. It should also be noted that measurement and reporting of a parameter may also include alarms and exceptions if that parameter passes outside of the acceptable or expected operating range.

No	Name	Description	Expected	Value
9.1.1	Incoming energy consumption meter	Install metering equipment capable of measuring the total energy use of the data centre, including all power conditioning, distribution and cooling systems. Again, this should be separate from any non data centre building loads. Note that this is required for CoC reporting	Entire Data Centre	3
9.1.2	IT Energy consumption meter	Install metering equipment capable of measuring the total energy delivered to IT systems, including power distribution units. This may also include other power feeds where non UPS protected power is delivered to the racks. Note that this is required for CoC reporting.	Entire Data Centre	3
9.1.3	Room level metering of supply air temperature and humidity	Install metering equipment at room level capable of indicating the supply air temperature and humidity for the IT equipment.	No	2
9.1.4	CRAC / AHU unit level metering of supply or return air temperature and humidity	Collect data from CRAC / AHU units on supply or return (dependent upon operating mode) air temperature and humidity.	No	3
9.1.5	PDU level metering of IT Energy consumption	Improve visibility of IT energy consumption by metering at the Power Distribution Unit inputs or outputs.	No	3
9.1.6	PDU level metering of Mechanical and Electrical energy consumption	Improve visibility of data centre infrastructure overheads	No	3

9.1.7	Row or Rack level metering of temperature and humidity	Improve visibility of air supply temperature and humidity	No	2
9.1.8	Device level metering of temperature	Improve granularity by using built in device level metering of intake and / or exhaust air temperature as well as key internal component temperatures	No	4

9.2 Energy Use and Environmental Collection and Logging

Once data on energy use and environmental (temperature and humidity) conditions is available through the installation of measurement devices in needs to be collected and logged.

No	Name	Description	Expected	Value
9.2.1	Periodic manual readings	Entry level energy, temperature and humidity reporting can be performed with periodic manual readings of consumption meters, thermometers and hygrometers. This should occur at regular times, ideally at peak load. Note that energy reporting is required by the CoC reporting requirements, also that automated readings are considered to be a replacement for Participant status.	Entire Data Centre	3
9.2.2	Automated daily readings	Automated daily readings enable more effective management of energy use. Supersedes Periodic manual readings.	No	4
9.2.3	Automated hourly readings	Automated hourly readings enable effective assessment of how IT energy use varies with IT workload Supersedes Periodic manual readings and Automated daily readings.	No	4

9.3 Energy Use and Environmental Reporting

Energy use and environmental (temperature and humidity) data needs to be reported to be of use in managing the energy efficiency of the facility.

No	Name	Description	Expected	Value
9.3.1	Written report	Entry level reporting consists of periodic written reports on energy consumption and environmental ranges. This should include determining the averaged DCiE over the reporting period. Note that this is required by the CoC reporting requirements, also that this report may be produced by an automated system.	Entire Data Centre	3

9.3.2	Energy and environmental reporting console	An automated energy and environmental reporting console to allow M&E staff to monitor the energy use and efficiency of the facility provides enhanced capability. Averaged and instantaneous DCiE are reported. Supersedes written report	No	3
9.3.3	Integrated IT energy and environmental reporting console	An integrated energy and environmental reporting capability in the main IT reporting console allows integrated management of energy use and comparison of IT workload with energy use. Averaged, instantaneous and working range DCiE are reported and related to IT workload. Supersedes Written report and Energy and environmental reporting console. This reporting may be enhanced by the integration of effective physical and logical asset and configuration data.	No	4

9.4 IT Reporting

Utilisation of the IT equipment is a key factor in optimising the energy efficiency of the data centre.

No	Name	Description	Expected	Value
9.4.1	Server Utilisation	Reporting of the processor utilisation of the overall or grouped by service / location IT server estate. Whilst effective metrics and reporting mechanisms are still under development a basic level of reporting can be highly informative.	No	3
9.4.2	Network Utilisation	Reporting of the proportion of the overall or grouped by service / location network capacity utilised. Whilst effective metrics and reporting mechanisms are still under development a basic level of reporting can be highly informative.	No	3
9.4.3	Storage Utilisation	Reporting of the proportion of the overall or grouped by service / location storage capacity and performance utilised. Whilst effective metrics and reporting mechanisms are still under development a basic level of reporting can be highly informative. The meaning of utilisation can vary depending on what is considered available capacity (e.g., ports, raw v. usable data storage) and what is considered used (e.g., allocation versus active usage). Ensure the definition used in these reports is clear and consistent. Note that mixed incentives are possible here through the use of technologies such as de-duplication.	No	3

10 Practices to become minimum expected

The following practices are planned to become minimum expected practices in future updates of the Code. The update year of the code in which the practices will become expected is shown in the table.

No	Name	Description	Expected	Year
4.1.3	Multiple tender for IT hardware – Extended operating temperature and humidity range	<p>Starting 2012 new IT equipment should be able to withstand the air inlet temperature and relative humidity ranges of 5°C to 40°C and 5% to 80% RH, non-condensing respectively, and under exceptional conditions up to +45°C as described in ETSI EN 300 019 Class 3.1.</p> <p>All vendors should indicate the maximum allowable temperature and humidity for all equipment to maximise the efficiency opportunities in refrigeration and free cooling.</p> <p>It should be noted that where equipment with differing environmental requirements is not segregated, the equipment with the more restrictive temperature range will influence the cooling conditions and corresponding energy consumption for all of the IT Equipment.</p> <p>From 40°C to 45°C intake temperature it is acceptable for equipment to implement performance reduction mechanisms to continue delivering the intended service at lower speed whilst preventing damage. These mechanisms should not reduce performance below 80% of the nominal for that device. Where such performance reduction mechanisms are used a clear description of the operating parameters and performance impact should be provided.</p> <p>To address equipment types which cannot be procured to meet this specification exclusions are provided in practices 5.1.10 for all data centres and 5.1.12 for new build. Directly liquid cooled IT devices are addressed in practice 4.1.12.</p>	New IT Equipment	2012

4.1.9	Energy & temperature reporting hardware	<p>Select equipment with power and inlet temperature reporting capabilities, preferably reporting energy used as a counter in addition to power as a gauge. Where applicable, industry standard reporting approaches should be used such as IPMI, DCMI and SMASH.</p> <p>To assist in the implementation of temperature and energy monitoring across a broad range of data centres all devices with an IP interface should support one of;</p> <ul style="list-style-type: none"> • SNMP polling of inlet temperature and power draw. Note that event based SNMP traps and SNMP configuration are not required • IPMI polling of inlet temperature and power draw (subject to inlet temperature being included as per IPMI 2.0 rev 4) • An interface protocol which the operators' existing monitoring platform is able to retrieve inlet temperature and power draw data from without the purchase of additional licenses from the equipment vendor <p>The intent of this practice is to provide energy and environmental monitoring of the data centre through normal equipment churn.</p>	New IT equipment	2012
5.1.10	Equipment segregation	<p>Deploy groups of equipment with substantially different environmental requirements and / or equipment airflow direction in a separate area. Where the equipment has different environmental requirements it is preferable to provide separate environmental controls.</p> <p>This objective of this practice is to avoid the need to set the data centre cooling plant for the IT equipment with the most restrictive environmental range or poor air flow control as this compromises the efficiency of the entire data centre.</p>	New IT Equipment	2012
5.1.12	Separate environmental zones	<p>Where a data centre houses both IT equipment compliant with the extended range of practice 4.1.3 and other IT equipment which requires more restrictive temperature or humidity control separate areas should be provided. These areas should have separate environmental controls and may use separate cooling systems to facilitate optimisation of the cooling efficiency of each zone.</p> <p>Examples are equipment which;</p> <ul style="list-style-type: none"> • Requires tighter environmental controls to meet archival criteria such as tape • Requires tighter environmental controls to meet long warranty durations (10+ year) <p>The objective of this practice is to avoid the need to set the data centre cooling plant for the equipment with the most restrictive environmental range and therefore compromising the efficiency of the entire data centre.</p>	New build or retrofit	2012

11 Items under Consideration

This section contains suggested items that are under consideration for inclusion in the Best Practices.

No	Name	Description	Expected	Value
11.1	DC (direct current) power distribution	Consider the use of high voltage DC (direct current) power distribution within the data centre. This can reduce the overall power conversion and distribution losses within the facility.	No	2
11.2	Optimal Power Density	Guideline recommendations on the most efficient range for power density	No	2
11.3	Utilisation targets	Minimum or average targets for the utilisation of IT equipment (servers, networking, storage). This presents substantial risk when considered without the load to power profiles of the equipment and would require substantial work.	No	3
11.4	Further development of software efficiency definitions	There is much research and development needed in the area of defining, measuring, comparing and communicating software energy efficiency. Suggested examples of this are; Software could be made resilient to delays associated with bringing off-line resources on-line such as the delay of drive spin, which would not violate the service level requirements. Software should not gratuitously poll or carry out other unnecessary background "housekeeping" that prevents equipment from entering lower-power states, this includes monitoring software and agents.	No	3
11.5	Further development of storage performance and efficiency definitions	Storage performance has multiple dimensions, including throughput and latency, not all of which can be measured at the storage layer. Capacity also has multiple dimensions, allocation and usage, not all of which can be measured at the storage layer. Technologies such as de-duplication, compression, snapshots, and thin provisioning also need to be accounted for in a consistent and informative manner.	No	3

4.1.12	Operating temperature range – Liquid cooled IT equipment	<p>Devices whose primary cooling method is not air (directly liquid cooled) are not subject to the air environmental requirements specified in 4.1.3.</p> <p>These devices should be able to operate with supply coolant liquid temperatures equal to the air temperatures specified in 4.1.3.</p> <p>As described in 5.6.4 this practice applies to devices which deliver cooling fluid directly to the heat removal system of the components such as water cooled heat sinks or heat pipes and not the delivery of cooling liquid to an internal mechanical refrigeration plant or in chassis air cooling systems.</p>	No	4
4.1.13	Operating temperature range – Maximum full performance temperature	<p>In data centres where the IT equipment supply temperature is allowed to vary with external temperature to increase the available free cooling a key consideration is whether the full performance of IT equipment is available for the full range of IT intake temperatures.</p> <p>Operators using higher intake temperatures (35°C or above) should include the maximum temperature at which the IT equipment is able to operate without reducing performance, throughput or capacity for thermal reasons in their equipment selection process</p>	No	3