Ref. Ares(2016)5650345 - 29/09/2016

European Commission

Final Report

Ecodesign Technical Assistance Study on Standards for Lot 9 Enterprise Servers and Enterprise Data Storage

Intertek

Written by Intertek PLC
June 2016

Authors:

Ms Catriona McAlister, Intertek PLC

Mr Anson Wu, Intertek PLC

Mr Jonathan Wood, Intertek PLC

Mr Bob Harrison, Intertek PLC

Mr Stephen Fernandes, Intertek PLC

Ms Fiona Brocklehurst, Intertek PLC

EUROPEAN COMMISSION

Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs
Directorate C - Industrial Transformation and Advanced Value Chains
Unit C.1 – Clean Technologies and Products

Contact: Davide Polverini

E-mail: davide.polverini*@ec.europa.eu*

GROW-ECODESIGN@ec.europa.eu

European Commission
B-1049 Brussels

EUROPEAN COMMISSION

Final Report

Ecodesign Technical Assistance Study on Standards for ENTR Lot 9 Enterprise Servers and Enterprise Data Storage

Directorate-General for Internal Market, Industry, Entrepreneurship and SMEs
2016 ET-02-16-987-EN-N

Europe Direct is a service to help you find answers to your questions about the European Union.

Freephone number (*)**:**

00 800 6 7 8 9 10 11

(*) The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

LEGAL NOTICE

This document has been prepared for the European Commission however it reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

More information on the European Union is available on the Internet (http://www.europa.eu).

Luxembourg: Publications Office of the European Union, 2016

ISBN 978-92-79-61834-5
doi: 10.2873/782959

© European Union, 2016
Reproduction is authorised provided the source is acknowledged.

Printed in Belgium

TABLE OF CONTENTS

1. INTRODUCTION 9
2. POLICY CONTEXT 10
2.1. European initiatives 10
2.2. International and industry initiatives 11
3. STANDARDISATION CONTEXT 13
3.1. Basic principles 13
3.2. Entities involved 13
3.3. The standardisation process 14
4. PRODUCT SCOPE 16
4.1. Enterprise servers 16
4.2. Enterprise storage 16
4.3. Definitions 16
5. PERFORMANCE PARAMETERS CONSIDERED 18
6. RELEVANT EC STANDARDISATION REQUESTS 20
6.1. Standardisation request related to the computer regulation 20
6.2. Mandate M/462 on telecommunications infrastructure 20
6.2.1. Activities requested 20
6.2.2. Organisation of standardisation work 20
6.2.3. Review of standardisation activities 21
6.2.4. Standardisation workplan 21
6.3. Mandate M/543 on generic standards, which cover ecodesign requirements related to material efficiency 21
7. STANDARDISED TEST METHOD GAP ANALYSIS ON ENTERPRISE SERVERS AND ENTERPRISE DATA STORAGE 23
7.1. Relevant European standards (EN) 27
7.1.1. CENELEC (CLC TC 215): EN 50600 Series 27
7.2. Relevant international standards 27
7.2.1. ISO/IEC 30134-4 SO/IEC 30134-5 (ITEE and ITEU for servers) 27
7.2.2. IEC TR 62635:2012 28
7.2.3. ISO 7779:2010 28
7.3. Relevant national standards / initiatives 28
7.3.1. United Kingdom: British Standards Institute ZZ/1 Publicly Available Specification (PAS) 141:2011 28
7.3.2. United States: NSF/ANSI 426 28
7.3.3. United States: NIST Special Publication 800-88 29
7.3.4. United Kingdom CESG standards on data sanitation 29
7.4. Relevant industry standards and international initiatives 29
7.4.1. SPEC: SPECpower_ssj2008 29
7.4.2. SPEC: SERT V1.1.1 29
7.4.3. SNIA Emerald™ Power Efficiency Measurement Specification 30
7.4.4. EPRI & Ecova: Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7 31
7.4.5. ASHRAE TC 9.9 2011: Thermal Guidelines for Data Processing Environments, 4th Edition (2015) Equipment environmental specifications for air cooling 31
7.4.6. ECMA: ECMA-74 13th edition (June 2015) (based on ISO 3741, ISO 3744, ISO 3745, ISO 11201) 32
7.4.7. IEEE 1680.4 Servers 32
7.4.8. ANSI ATIS: 060015.2013 (TEER) 32

8. STANDARDS ENGAGEMENT 33
8.1. Standards activities engaged with 33
8.2. *Beta testing / Evaluation Programme SERT v1.1.1* 33
8.2.1. Description 33
8.2.2. Objectives *34*
8.2.3. Results *34*
9. CONCLUSIONS 35

APPENDICES

APPENDIX 1 Coverage and status of key standards / initiatives for enterprise servers and data storage

APPENDIX 2 Server and Storage Standards Listing

APPENDIX 3 Practical Insights on SERT™ testing for Enterprise Servers

APPENDIX 4 Metrics Paper

List of Figures

Figure 1 Ecodesign preparatory and adoption procedure..... 10

Figure 2 - Standards currently available..... 24

Figure 3- Traffic light summary of standards gaps for enterprise servers (ES) and data storage 26

Figure 4 - Summary of ASHRAE Thermal Guideline Classes..... 31

List of Tables

Table 1 Entities involved in standardisation 14

Table 2 - Standardisation organisations 14

Table 3 – Parameters to assess for standards availability 19

Table 4 – Coverage and status of key standards / initiatives for enterprise servers (ES) and data storage 23

Table 5 – Standards engagement activities during project 33

Table 6 - Ongoing priority standardisation activities..... 35

1. INTRODUCTION

The Ecodesign Technical Assistance Study on Standards for Enterprise Servers and Data Storage (DG ENTR Lot 9) is a European Commission study, led by Intertek.

The study aims to provide technical assistance to support standardisation-related tasks for equipment under the scope of the DG ENTR Lot 9 (namely enterprise servers and data storage). The focus of this project is upon the provision of technical assistance to facilitate the establishment of the foundation standards (or, when necessary, transitional methods) that will be necessary for implementing measures addressing Lot 9 products, should the Commission decide to proceed with such measures. Standardised approaches to measurement are necessary in order that manufacturers can assess their compliance with any requirements that may be defined in regulation, and so that national bodies can assess market compliance of products on their markets.

The study methodology centres upon an assessment of the need for standards – identifying parameters and existing standards and identifying gaps. The priority is to facilitate work towards:

A robust, durable standardised method for measuring the energy efficiency of servers (especially rack servers but also blade servers)

A robust, durable standardised method for measuring the energy efficiency of data storage devices

The study includes interaction with the relevant standardisation processes and a consideration of how metrics might be built upon the identified standards. This second component includes testing of rating tools and measurement approaches in order to provide recommendations to the standardisation processes underway to ensure repeatability, consistency and robustness.

The final deliverables include the following:

- "Standardised Test Method Gap Analysis" (Appendix 1)
- "Server and Storage Standards Listing" (Appendix 2)
- "Practical Insights on SERT testing for Enterprise Servers" report. (Appendix 3)
- "White Paper: Investigation of potential approaches to energy efficiency metrics for enterprise servers, based upon the SERT rating tool" (Appendix 4)

This document presents the results of a gap analysis into standardised test methods for measurement and calculation, which could be used to support the implementation of a potential future EU Ecodesign Regulation on servers and storage equipment. It considers the parameters that could be described in Annex II of a Regulation, explores how these would need to be supported by standardised test methods.

2. POLICY CONTEXT

2.1. European initiatives

The Ecodesign Directive is a key European Union (EU) sustainability policy, addressing both competitiveness and sustainable development in line with Europe's 2020 Strategy. The directive aims to improve upon environmental performance of energy related products across the EU, by establishing a framework to set ecodesign requirements or to encourage manufacturer voluntary agreements.

DG Growth and DG Energy are responsible for the Ecodesign directive. The first step toward an ecodesign regulation is the identification of a product on the ecodesign working plan - an indicative list of product groups that are considered as priorities for the adoption of implementing measures. This is followed by a preparatory study which explores the options to improve the environmental performance of the product and provides the necessary information to prepare for the next phases in the policy process such as the impact assessment, the consultation forum, and the possible draft implementing measures or voluntary agreement.

The Working Plan for 2012-2014 identified Enterprise servers and data storage as a key product area to be addressed, with initial estimated potential savings of 135 PJ/year as of 2030. As a result, the preparatory study "DG ENTR Lot 9" covering enterprise servers, data storage and ancillary equipment was initiated. The ecodesign preparatory and adoption procedures are illustrated in Figure 1.

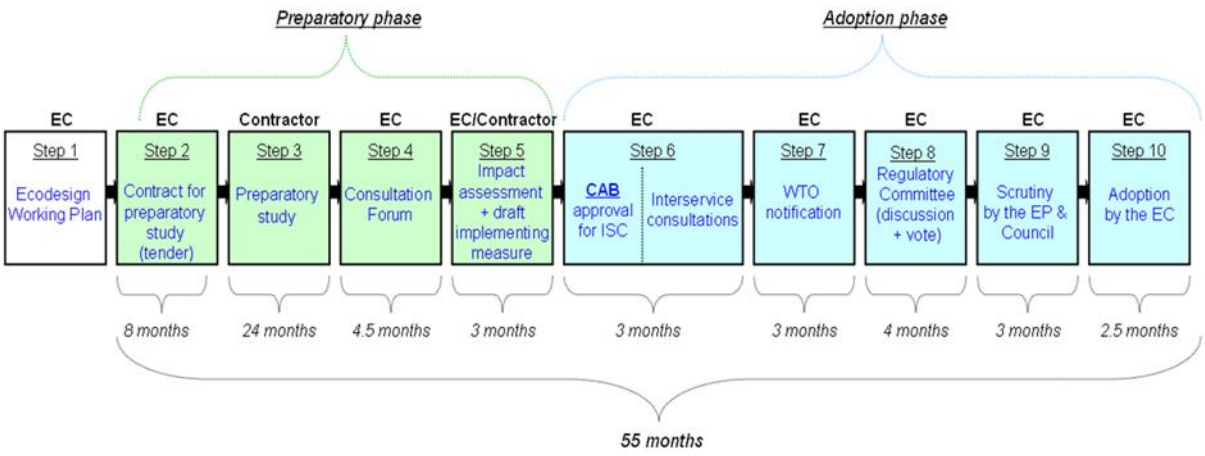


Figure 1 Ecodesign preparatory and adoption procedure

The Lot 9 preparatory study was completed in September 2015, and the final report published in November 2015. Work has now progressed to step 4/5 of the above diagram. A dedicated impact assessment study was commenced in October 2015 in order to analyse various potential policy options, with regard to servers and data storage devices. In parallel with this activity, this technical assistance contract on standardisation gaps is intended to develop measurement methods for the energy efficiency/product performance of servers and data storage devices.

The Commission has already regulated some aspects of servers through the Commission Regulation (EU) No 617/2013 of 26 June 2013, implementing Directive 2009/125/EC of the European Parliament and of the Council with regard to ecodesign requirements for computers and computer servers. This regulation only addresses efficiency requirements for power supply units in a sub-set of servers. The Commission is due to review the Ecodesign Regulation on computers towards the end of 2015.

Other overarching EU policy initiatives of relevance not addressed in further in this study include:

- The Waste Electrical and Electronic Equipment (WEEE) Directive (2012/19/EU)
- The REACH Regulation (No 1907/2006)

- The Restriction of Hazardous Substances (RoHS) Directive (2011/65/EU)
- The Electromagnetic Compatibility Directive (2004/108/EC)
- Low Voltage Directive (2006/95/EC)
- Regulation (EU) No 1275/2008 on standby and off mode electric power consumption of electrical and electronic household and office equipment

2.2. International and industry initiatives

International voluntary policy initiatives approach energy efficiency varying objectives and approaches. Those which have begun to establish measurement methods and standards include:

The voluntary ENERGY STAR® label (United States (US) Environmental Protection Agency (EPA) and European Commission): Addresses data centre products such as enterprise servers, data storage and large network equipment. ENERGY STAR programme requirements for servers v2.0 was implemented in the US on the 16th December 2013. The US EPA v3.0 server specification was launched in March 2016. New criteria for data storage are in development.

ENERGY STAR previously developed a testing methodology and performance standard for server idle power. This is limited to 1-2 socket servers which covers the largest sector of the server market.

The EU Code of Conduct for Data Centres (European Commission Joint Research Centre): Takes a holistic approach to the operation and selection of equipment for use in data centres, providing a means of outlining energy efficient best practice and putting in place voluntary targets for signatories to meet. The Code of Conduct does not specify test methods for IT equipment, but encourages selection tailored to the specific data centre application, and references ENERGY STAR for Servers as a possible solution for procuring efficient IT equipment.

Blue Angel (The Federal Ministry for the Environment, Nature Conservation, Building and Nuclear Safety): This eco-label claims to provide the demand side (public sector or industry) with a reliable means of including ecological criteria in procurement contracts for external data centre services. The most recent Blue Angel eco-label for data centre services was implemented in February 2015 and includes both technical and information disclosure requirements, as well as recommendations on further energy saving opportunities.

Triple E programme (The Sustainable Energy Authority of Ireland): The Triple E is a searchable listing of energy efficient equipment that meet minimum criteria listed under the programme. The programme covers a range of server and storage products used in data centres with the most recent specifications developed in 2010.

Certified Energy Efficient Data Centre Award (CEEDA) (Datacenter Dynamics Ltd): CEEDA provides an audited and certified assessment of the implementation of energy efficiency best practices within a data centre. It delivers an operational and deployment roadmap for further improving performance and enables demonstration of conformance to a benchmark. Most of the current best practices are derived from the European Code of Conduct for Data Centres, with metrics included developed by The Green Grid, which depending on the assessment type may include: Power usage effectiveness (PUE); carbon usage effectiveness (CUE); water usage effectiveness (WUE) and energy reuse effectiveness (ERE).

80 PLUS certification (Ecova): 80 PLUS is an electric utility-funded incentive programme to integrate more energy-efficient power supplies into desktop computers and servers. The performance specification requires power supplies of 80% or greater energy-efficiency.

Top Runner Program in Japan (Energy Conservation Center Japan): The Top Runner Programme was introduced in 1999 to reduce energy consumption in Japan. The programme includes energy efficiency requirements for a range of different products types including servers. The range of servers covered is very wide, including mainframes, blade and 1-4 socket rack servers. The same metric is used across all computing products and is based on theoretical maximum central processing Unit (CPU) performance (CTP), idle power and standby power. It does not consider Random Access Memory (RAM) or hard drives.

These relatively recent policies continue to mature as knowledge builds and industry responds to demand for greater efficiency. Almost all of the efforts to date have been focussed on energy

efficiency as it is considered the highest lifecycle impact of these products, which are under continuous use in relatively high power consumption modes over their entire lifetime. This means that other environmental parameters such as hazardous chemical, and recyclability have mostly gone unaddressed. However, it should be noted that the NSF International (US Green Electronics Council) and Institute of Electrical and Electronics Engineers (IEEE) standards are seeking to address some of the wider environmental cycle impacts associated with servers in order to feed into an EPEAT specification for green procurement.

In addition, industry associations and partnerships such as The Green Grid work to encourage greater efficiency within data centres and provide a means of recognition for those who achieve the specified levels.

3. Standardisation context

3.1. Basic principles

The concept of a standard is well established. Recognised definitions of standards are shown below:

ISO website: A standard is a document, established by a consensus of subject matter experts and approved by a recognised body that provides guidance on the design, use or performance of materials, products, processes, services, systems or persons.

Formal definition of a Standard (ISO/IEC Guide 2): Document, established by consensus and approved by a recognised body, that provides, for common and repeated use, rules, guidelines or characteristics for activities or their results, aimed at the achievement of the optimum degree of order in a given context.

Definition from regulation 1025/2012 on standardisation: 'standard' means a technical specification, adopted by a recognised standardisation body, for repeated or continuous application, with which compliance is not compulsory, and which is one of the following:

(a) 'international standard' means a standard adopted by an international standardisation body;

(b) 'European standard' means a standard adopted by a European standardisation organisation;

(c) 'harmonised standard' means a European standard adopted on the basis of a request made by the Commission for the application of Union harmonisation legislation;

(d) 'national standard' means a standard adopted by a national standardisation body;

Standards are not the same as regulations. They are voluntary but are often necessary to support the implementation of regulation in that they describe how attributes of products should be measured in a clear and reproducible manner.

Implementing measures require clear, robust and appropriately harmonised measurement standards (or draft transitional methods – explained later) fairly applied to all products in scope. Without the foundation of standards, enforcement of regulation becomes impossible and laws have no force behind them.

Priorities in the creation of standards are:

- Robustness
- Clarity
- Applicability
- Avoidance of loopholes
- Coherence with other standards
- Complementary to legal requirements.

3.2. Entities involved

There are various different bodies involved in standardisation as listed in Table 1.

Table 1 Entities involved in standardisation

Type	Examples
Government organisations	Codex, IMO, UN/ECE.
National standardisation bodies	BSI, DIN, AFNOR
European standardisation organisations (ESOs)	CEN, CENELEC and ETSI
International standardisation bodies	ISO, IEC and ITU
Industry consortia	The Green Grid, Ecova Plug Load Solutions (80 Plus)

In the European Union, only standards developed by the ESOs (see Table 2) are recognised as 'European Standards'. The ESOs closely cooperate in the interest of European harmonisation, creating both standards requested by the market and harmonised standards in support of European legislation.

Table 2 - Standardisation organisations

ESO	CEN	CENELEC	ETSI
Description	European Committee for Standardisation. The main body for developing standards in Europe in all areas except telecommunications (ETSI) and electrotechnical (CENELEC).	European Committee for Electrotechnical Standardisation. CENELEC coordinates closely with CEN via the CEN-CENELEC Management Centre (CCMC) on strategic matters of common interests.	European Telecommunications Standards Institute. Produces globally-applicable standards for Information and Communications Technologies (ICT), including fixed, mobile, radio, converged, broadcast and internet technologies.
International equivalent	ISO (the International Organization for Standardisation)	IEC (the International Electrotechnical Commission)	ITU-T (the International Telecommunication Union and telecommunication standardisation sector)

There are agreements to recognise international standards against the particular needs within the EU for standards where a need has not been recognised or prioritised at the international level. Many CEN and CENELEC standards are identical to ISO and IEC standards – around 31% of CEN standards are identical to ISO due to the Vienna Agreement[1], and around 60% of the CENELEC standards are substantially identical to IEC due to the Dresden agreement[2].

3.3. The standardisation process

The standardisation process within ecodesign usually involves the European Commission making a formal standardisation request (SR) to ESOs to develop product-specific standards relevant to aspects of performance in accordance with Regulation (EU) No 1025/2012 on European standardisation[3]. This directive, one of the foundations of the single market act[4], has the goal of modernising the European standards process to enable more standards to be produced, faster and with greater inclusivity. It provides the general framework for European standardisation policy and places obligations on the recognised European Standardisation Organisations to meet the standardisation principles of transparency, openness, impartiality and consensus, effectiveness and relevance, coherence, and development dimension. Standardisation requests are created for each new implementing measure under the Ecodesign Directive.

[1] International Organization for Standardization (ISO) and European Committee for Standardization (CEN), agreement on technical co-operation between ISO and CEN (Vienna agreement), http://boss.cen.eu/ref/Vienna_Agreement.pdf
[2] IEC - CENELEC Agreement on common planning of new work and parallel voting, http://www.iec.ch/about/globalreach/partners/regional/iec_cenelec_agreement.htm
[3] paragraphs 1 and 2 of Article 10
[4] (SMA) COM 2011

Implementing measures will usually reference a product-specific "harmonised" standard, meaning a specification adopted by a recognised standards body under a mandate from the Commission[5]. A harmonised standard is deemed to exist when ESO members have formally presented the standards produced or identified conformity with the mandate.

Where a harmonised standard does not exist, transitional measuring methods and verification procedures can be detailed in a separate communication in the Official Journal of the European Union (OJEC), which can then be referenced in Commission guidance to accompany ecodesign directives for products. Such a communication would typically list out test methods in tabular form for each directive requirement. Transitional methods would ultimately be replaced by harmonised standards, which would also be published in the OJEC in accordance with Articles 9 and 10 of Directive 2009/125/EC.

European standards can be split into two main types – prescriptive (state requirements) and non-prescriptive (provide advice or information). In the ESO process, the full (EN) standard is the most prescriptive, and is usually what is referred to as a harmonised standard. It guarantees the commitment of national standards bodies (NSBs) who must adopt the standard at a national level and remove/modify any conflicting standards (even if the country voted against the draft). EN standards may take 2 to 4 years to develop and must be reviewed at the latest 5 years from publication.

[5] in accordance with the procedure laid down in Directive 98/34/EC of the European Parliament and of the Council of 22 June 1998 laying down a procedure for the provision of information in the field of technical standards and regulations (1), for the purpose of establishing a European requirement, compliance with which is not compulsory.

4. Product Scope

The product scope for this gap analysis is limited to the scope published in the Task 7: Scenarios report[6] published as part of the Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC (DG ENTR Lot 9) - Enterprise servers and data equipment.

The Task 7 report defined the scope of the Preparatory study as limited to "enterprise servers" and "enterprise storage" products. Networking equipment was excluded.

4.1. Enterprise servers

The report specified "enterprise servers" as including products that are:

Defined as computer servers according to the definition of the ENERGY STAR® specification for computer servers (version 2.0),

Modular and having different form factors,

Marketed and sold through enterprise channels.

"Enterprise servers" were not considered to include products that are:

Intended for private end-users (domestic) or embedded (machinery) applications

The report's authors did not explicitly remove from the scope enterprise server types such as mainframes, high performance computer systems, resilient servers or server appliances. However, they strongly recommended that the technical, economical and operational feasibility of ecodesign measures for these products should be reviewed in detail. In particular, these products could be difficult to as they could be highly customised and used for mission-critical computing processes in which functional or operational requirements take priority over environmental performance.

4.2. Enterprise storage

The Task 7 report specified "enterprise storage" as including products that are:

Defined as storage product according to the definition of the ENERGY STAR® specifications for data centre storage equipment (version 1.0)

Marketed and sold through enterprise channels

"Enterprise storage" was not considered to include products that are:

Private (domestic) and portable data storage products, computer servers, computers with storage capacities, and network equipment.

Whilst more specialist enterprise storage equipment such as Online 5 or 6 was not explicitly removed from the scope of the preparatory study, it was noted that due to the specialist nature of these product types, the environmental performance of these products may be of significantly less concern than operational performance.

4.3. Definitions

The Task 7 report proposed adopting the product definitions already used in other EU Regulations, such as the Ecodesign Regulation (EU) No 617/2013 on computers and computer servers in order to guarantee a harmonised approach. Where products are not defined in an existing Regulation, it was suggested that definitions be aligned with the relevant ENERGY STAR specification (Enterprise Servers Specification Version 2.06 and Data Centre Storage Eligibility Criteria Version 1.07)[7].

6 Preparatory study for implementing measures of the Ecodesign Directive 2009/125/EC DG ENTR Lot 9 - Enterprise servers and data equipment (June 2015) Task 7 Draft report. available at www.ecodesign-servers.eu/

7 It should be noted that depending on the timing of any Ecodesign Regulation measures, reference to product definitions in newer ENERGY STAR specification may be more appropriate.

Whilst definitions could be based off the foundation of ENERGY STAR, it is likely that they would need to be refined in order to be sufficiently robust for the purposes of a standard supporting any regulation. Due to the voluntary nature of the ENERGY STAR programme, definitions are able to have a degree of flexibility. The language may be insufficiently detailed to ensure that i) all products covered under the scope meet the definition and ii) all product types intended to be outside scope are explicitly excluded.

Further definitions (not defined in the Task 7 report) would need to be developed for standardisation purposes to address factors such as:

- Excluded products
- Power modes
- Individual components where additional allowances may apply

5. Performance parameters considered

Among the potential requirements envisaged for enterprise servers and data storage devices, the preparatory study identified:

- Product information requirements (on product performance, operating conditions, etc..)
- Requirements on product hardware components (e.g. on the efficiency of the internal power supply units)
- Requirements on product software components and configuration (e.g. software which supports virtualization)
- Requirements on some product operating conditions, in particular the temperature
- Requirements on product material efficiency (reuse, recycling).

As a result, the saving potential at the level of servers and data storage devices was been estimated to be in the order of 17 TWh by 2030.

Building upon this list with a knowledge of wider standardisation initiatives where appropriate, the key parameters for which standards needed to be identified have been outlined in Table 3.

Table 3 – Parameters to assess for standards availability

Impact area	Parameter	Product	Source / Explanation
Energy performance in operation	Active State (power demand / rating)	ES, DS	Preparatory study task 7
	Idle State (power demand/ rating)	ES, DS	Preparatory study task 7
	Energy proportional operation (dynamic range)	ES, DS	Preparatory study task 7
	Overall energy performance (TEC type approach)	ES, DS	Preparatory study task 7
Product hardware / software / configuration	Power Supply Efficiency	ES, DS	Preparatory study task 7
	Power Supply Power Factor	ES, DS	Preparatory study task 7
	Capacity Optimizing Methods (COMs)	DS	Preparatory study task 7
	Reusability of components (Firmware availability)	ES, DS	Preparatory study task 7
Product operating conditions	Operating temperature and humidity	ES, DS	Preparatory study task 7
	Acoustic noise	ES, DS	Preparatory study task 7
Material efficiency	Removability of external enclosures, PCBs, processors, data storage devices and batteries with common tools	ES, DS	Preparatory study task 7
	Ease of dismantling, reuse and recycling at the end-of-life.	ES, DS	Preparatory study task 7
	Data sanitisation [8]	ES, DS	Preparatory study task 7
	Critical raw material (CRM) content	ES, DS	Preparatory study task 7
	Postconsumer recycled content of CRM	ES, DS	JRC Science and Policy Report, Environmental Footprint and Material Efficiency Support for product policy, Analysis of material efficiency requirements of enterprise servers
	Replacement component availability	ES, DS	JRC as previous
	Reduction of surplus parts by default	ES, DS	JRC as previous
	Hardware functionality testing software tools	ES, DS	JRC as previous

[8] Data sanitisation is the complete removal or all data from a storage component or equipment to make it unrecoverable by forensic methods. Sanitisation can be destructive or non-destructive to the hardware equipment. While it is not a direct environmental aspect,

6. Relevant EC Standardisation Requests

6.1. Standardisation request related to the computer regulation

As previously mentioned, "Commission Regulation (EU) No 617/2013 of 26 June 2013 describing ecodesign requirements for computers and computer servers" addresses efficiency requirements for power supply units in a sub-set of servers.

In relation to this regulation, the Commission has already issued:

A Commission Communication on transitional methods for measurement.

A standardisation request (mandate 545[9]) for the ESOs to develop harmonised standards which will incorporate relevant measurement and calculation methods.

The standardisation request includes standards to enable the measurement of power supply efficiency. It has been accepted by the CENELEC Technical Board[10] and the standardisation work within CENELEC has been initiated.

6.2. Mandate M/462 on telecommunications infrastructure

6.2.1. Activities requested

The European Commission recognised that further action was needed in order to improve the energy efficiency and offset the growth of the telecommunications infrastructure. Therefore, in 2010, the Commission published the EU standardisation request M/462 addressed to CEN, CENELEC and ETSI. The standardisation request called for two distinct phases of work:

In Phase 1 of the standardisation request, the European standardisation organisations CEN, CENELEC and ETSI, in cooperation with other relevant standards organisations, were invited:

To analyse the economic environment and the political context for efficient energy use,

To identify the potential role of ICT standardisation in achieving efficient energy use,

To identify existing and/or ongoing standardisation and consensus-building activities on the issue within national, regional and international standardisation organisations, formal or otherwise, and to assess their relevance for achieving European policy objectives in this domain,

To identify consistencies, relations, dependencies, hierarchy (taxonomy), overlaps and gaps in ICT standardisation work related to efficient energy use,

To establish a standardisation work programme with a view to filling the gaps, taking into account relevant regulatory initiatives, R&D projects or standardisation activities carried out by relevant fora and consortia.

In Phase 2, the work programme is agreed and the standardisation activities are launched following consultation of the Member States on the results of Phase 1.

In 2011 the ESO's published the "Framework Document for ESO Response to EU Mandate M/462", which aimed to respond to the Phase 1 standardisation request requirements and provide a gap-analysis of the relevant existing and ongoing standardisation activities.

Whilst a number of standards will be delivered under the M462 workplan, this does not necessarily mean that these standards would automatically be referenced by any regulatory measures, should these be developed. Each standard would need to be considered for suitability on a case by case basis.

6.2.2. Organisation of standardisation work

The Joint Coordination Group established between CEN, CENELEC and ETSI in response to Mandate M/462 is coordinating the standardisation work for this request. In addition, they will also coordinate with the CEN, CENELEC and ETSI "Green Data Centres" group (CG GDC), as a first step

9 http://ec.europa.eu/growth/tools-databases/mandates/index.cfm?fuseaction=search.detail&id=566

10 https://www.cenelec.eu/aboutcenelec/whatwestandfor/supportlegislation/europeanmandates.html

in focussing on energy efficiency. ETSI's Technical Committee "Environmental Engineering" (TC EE) is responsible for defining the environmental and infrastructural aspects for telecommunication equipment in various types of installations. They are therefore involved in engineering aspects of standards such environmental conditions (climatic, thermal, acoustic, etc.), equipment (physical requirements of racks, sub-racks and cabinets including thermal matters), power supply requirements, and eco-environmental matters (energy efficiency, environmental impact analysis, alternative energy sources). The key activities of TC EE addressing eco-environmental matters are:

reduction of power consumption of telecommunication equipment and related infrastructure;

determination of the environmental impact of telecommunication equipment.

Cooperation of TC EE with other Technical Bodies and with external organizations is managed within ETSI through the Operational Co-ordination Group (OCG). Several external liaisons have been established with standardisation bodies including IEC, CENELEC and the ITU-T, and other organisations / research projects. CENELEC supports the ETSI standardisation activity by producing standards in the energy efficiency field for components, infrastructure designs and infrastructure installation which are applicable to the delivery of ICT within customer premises (which may also be applicable to the needs of operator's sites). An example of this is the development of the EN 50600 series which, in part, addresses the installation of appropriate infrastructure to enable the energy efficiency of data centres to be measured and monitored. CEN is not actively involved as the most relevant CEN activities (development of standards addressing life cycle assessment) lie outside the direct focus of operational energy efficiency defined by the Mandate M.462.

6.2.3. Review of standardisation activities

The ESO review of standardisation activities relevant to Mandate M/462 included those at a European and an international level, covering relevant documentation being produced by ESOs and other standards organisations, fora and consortia. Of particular relevance is the CEN-CLC-ETSI established Coordination Group to develop standards for Energy Efficiency within Data Centres and associated infrastructure. This review of standardisation activities used as a basis the ESO response to M/462, building upon this with more recent insights, as detailed in section 7.

6.2.4. Standardisation workplan

The standardisation workplan put forward by the ESOs is not broken down to the level of specific areas or standards, but states a general ambition to have published European Standards (ENs) covering each of the subject areas ("Operation", "Test" and "KPI") for areas including:

Network operator sites / Facilities / Data centres

IT Equipment / Servers and storage

In particular, in relation to KPIs, the following activities have been since highlighted by ESOs as necessary for ecodesign purposes:

Measurement Process for Energy Efficiency KPI for Servers

Measurement method and Process for Energy Efficiency KPI for Storage equipment[11]

The original objective of ESO's was to produce the required standards within 3 years from the approval of the standardisation program (around 2014), but it is likely that this deadline has now been delayed as some of the standardisation work is still underway.

ETSI has to date led much of this standardisation work. Due to the telecommunications focus of ETSI however, these activities have more concentrated on network efficiency considerations than on server and storage energy efficiency. For standards to be applicable for Lot 9 products, a product-specific focus on energy efficiency will be necessary.

6.3. Mandate M/543 on generic standards, which cover ecodesign requirements related to material efficiency

The M/543 standardisation request was issued on 17.12.2015. It aims to contribute to the implementation of the Commission's action plan on the Circular Economy through development of generic standards related to material efficiency aspects (such as recyclability, recoverability and

11 http://docbox.etsi.org/Workshop/2015/201506_EEWORKSHOP/SESSION01_Setting_the_Scene/Mandate_462_RodolpheWouters_EC.pdf

reusability, durability, reversible disassembly and end of life extraction time) that could by applied to any product group listed in the Article 16 of Ecodesign Directive.

Standards delivered under the work programme of M/543 can provide a foundation from which product specific standards can be developed, but are not a pre-requisite for the development of product-specific material efficiency standards (i.e. it is not expected that all product-specific standardisation work on material efficiency be stalled until the M/543 standards are delivered).

Therefore, this analysis, whilst not focusing on material efficiency aspects, has included details of standards relating to parameters that may be relevant to enterprise servers and data storage.

7. Standardised Test Method Gap Analysis on Enterprise Servers and Enterprise Data Storage

This section summarises the results of the 2015 gap analysis. Key standards were examined to determine which relevant parameters they addressed, what the current status of each standard was, the degree of harmonisation, and how relevant it was to supporting ecodesign for Lot 9 products.

The diagram in Figure 2**.** illustrates the interactions of the standards currently available. A further overview of the status of the standards and initiatives assessed is shown in Table 4.

Table 4 – Coverage and status of key standards / initiatives for enterprise servers (ES) and data storage

	Published standard in use by industry addressing most aspects of a parameter	Standard in draft that may be suitable to address a parameter	Shortlisted standards not suitable.	Total
European Standards	0	0	4	4
International Standards	2	1	0	3
National standards/initiatives	5	1	1	7
Industry standards/ International initiatives	5	2	1	8
Total	14	4	6	22

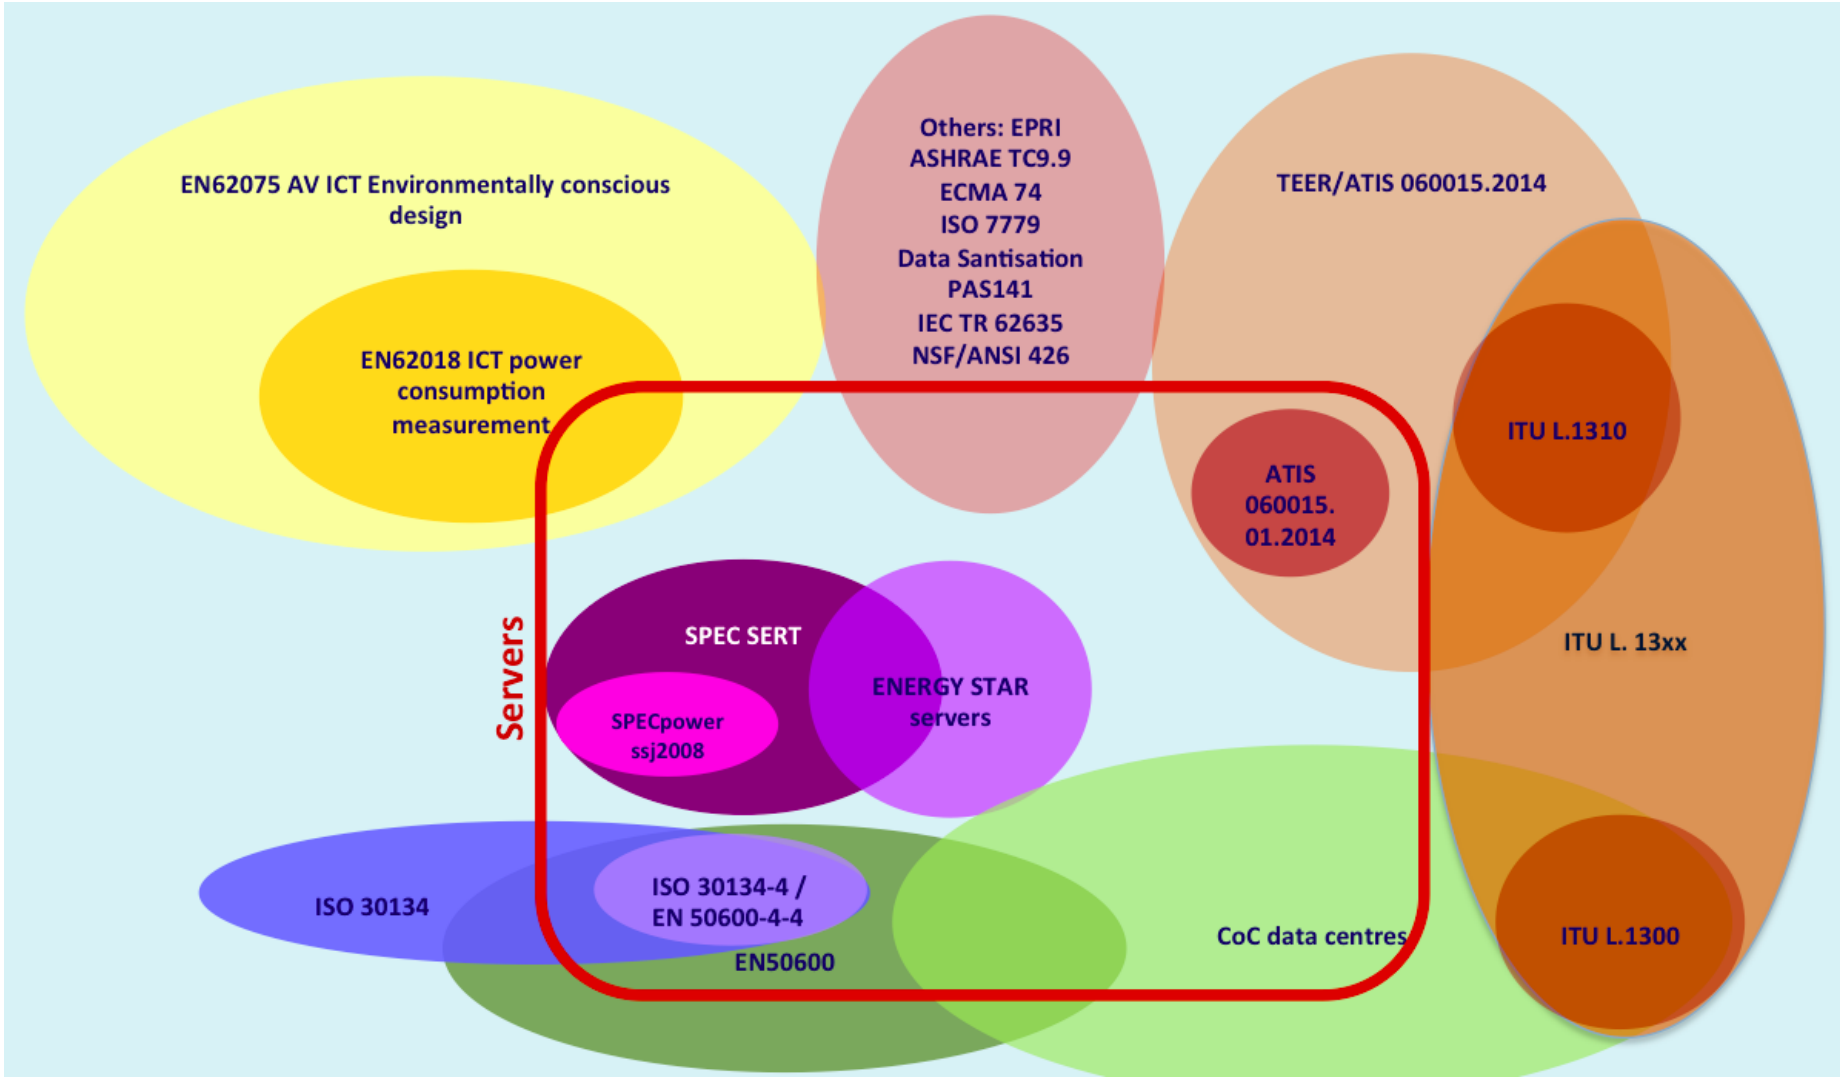


Figure 2 - Standards currently available

A detailed breakdown of the standards available and what parameters they address is contained in Appendix 1. Whilst 10 standards have already been identified that could address most aspects of a parameter, some parameters are better covered than others, and there are substantial gaps in some areas. These gaps are highlighted in the analysis of Figure 3 (Green signifies standards established and adequate, red that for a high priority area there is insufficient current coverage of standards).

The gap analysis highlighted the following considerations:

For both enterprise servers and data centre storage, the critical area of focus is the ability to assess and rate energy performance. Whilst some standards exist that could support these areas to some degree, further necessary work is underway to improve upon these approaches. In particular, it is necessary to ensure that the existing standards meet the priorities for creating a standard (robustness clarity etc described in section 3.1) and that sufficient data is available to validate the applicability of the test method over the range of server configurations and form the basis of a meaningful efficiency metric.

Test approaches for power supply efficiency and power factor are relatively well established although not harmonised.

Test approaches for material efficiency aspects, are for the most part not well-defined, but not a key priority at this stage. However, data sanitisation is relatively well supported.

Further information on the status of the most relevant standards is contained in the following section.

Parameter	Published standard in use by industry addressing most aspects of a parameter	Standard in draft that may be suitable to address a parameter	Shortlisted standards not suitable.	Ability to verify via test
Server active state and idle state power	ESTAR servers v2.0 SERT V1.1.0	ESTAR servers v3.0 **(2016)** ISO/IEC 30134 **(est 2017)**	SPECpower_ssj2008	Yes
Server overall energy performance (TEC)	ESTAR servers v2.0 SERT V1.1.0	ESTAR servers v3.0 **(2016)**	SPECpower_ssj2008	Yes but not fully validated for regulatory purposes
Data storage active state and idle state power, and overall energy performance (TEC)	SNIA Emerald V2.x	ESTAR storage v2.0 **(no timeline)**	ENERGY STAR storage v1.0 EN 50600	Yes but complex and test standard revisions not all compatible
Energy proportional operation	SERT V1.1.0		EN 62018	Yes – servers only
Data storage COMs	ESTAR storage v1.0 SNIA Emerald V2.x			Yes
Power Supply Efficiency	EPRI protocol v 6.7		EN 300 132-3	Yes
Power Supply Power Factor	EPRI protocol v 6.7			Yes
Firmware availability		NSF/ANSI 426 **(est 2016)**		Limited. Verification proc.
Operating temperature & humidity	ASHRAE guidelines 4th edition table 2.3		EN 300 019-1	Yes
Acoustic noise	ISO 7779:2010, ECMA 74		EN 62075	Yes
Removability of external enclosures, PCBs, processors, data storage devices and batteries with common tools	IEC TR 62635:2012 PAS 141:2011	NSF/ANSI 426 **(est 2016)**	EN 62075	Limited. Verification proc.
Ease of dismantling, reuse and recycling at the end-of-life.	IEC TR 62635:2012 PAS 141:2011	NSF/ANSI 426 **(est 2016)**	EN 62075	Limited. Verification proc.
Data sanitisation	NIST 800-88 rev1 CESG various PAS 141:2011		EN 62075	Yes
CRM content		NSF/ANSI 426 **(est 2016)**	EN 62075	Limited. Verification proc.
Postconsumer recycled content of CRM				
Replacement components availability				
Reduction of surplus parts by default				
Hardware functionality testing s/ware tools				

Figure 3 Traffic light summary of standards gaps for enterprise servers (ES) and data storage

7.1. Relevant European standards (EN)

7.1.1. CENELEC (CLC TC 215): EN 50600 Series

CENELEC is establishing a set of standards on data centres (DC), under the EN 50600 series. The standards are designed to be fair, consistent and comparable. The target audience is the average Small and medium enterprise (SME) DCs that might not have the high level of expertise compared to a large data centre and therefore benefit from design and operation guidelines.

The intention is to address the complexities of DC in a holistic form compared to previous work which has already covered discrete components and systems such as cabling, UPS, fire systems and access control.

The Activities are built on three pillars:

1. Design – building, power, environmental control, IT cabling and security.
2. Operation and Management
3. KPIs – to assess resource and energy efficiency including subsystems and possibly components which will be linked to the EU CoC for DCs.

The first two pillars are almost finished and activities are now concentrated on the KPIs for pillar 3 in close cooperation with ISO/IEC JTC 1/SC39 WG who are developing ISO/IEC 30134-4 (defining what a KPI should comprise). Whilst there is no formal collaborative agreement, standards development is coordinated between the CENELEC and ISO activities by the same experts sitting on both sets of committees.

The development of standards relating to energy efficiency of servers is being actively pursued via an accelerated process. Standards should be available in 2017 and will include PUE which has been handed over from The Green Grid. Building upon an existing white paper[12], a standard providing the rationale for the use of SERT, considering how to approach revision control, and suggesting how it could be used as a standard for evaluating servers is likely to be issued in a separate publication ISO/IEC 30134-7.

7.2. Relevant international standards

7.2.1. ISO/IEC 30134-4 SO/IEC 30134-5 (ITEE and ITEU for servers)

ISO/IEC 30134-4 is a project under the ISO/IEC Joint Technical Committee (JTC) 1/SC 39 addressing ITEE (IT energy efficiency) and ISO/IEC 30134-5 addresses ITEU (IT Energy Utilisation) KPIs for servers. Utilisation is not a relevant parameter for the product since it will depend on the application by the end-user. The work is being led by the Japanese and Korean representatives in the working group.

It is still in early stages of addressing some of the complex issues in the area. Whilst the first internal committee draft report was completed in spring 2015, further work is necessary to take this forward and transform the findings into KPIs.

The current draft defines KPIs and describes application but leaves it to the user to pick the test. There is no focus on any particular testing tool as there are so many different options covering different use cases and architecture (Linpack, SERT etc), however, the current focus is on the efficiency at maximum performance. This may present a problem because peak efficiency does not occur at maximum performance and servers are almost never used at this load level.

One possible long term solution is to arrive at a well-defined testing approach combining a number of different workloads to provide a representative picture of energy performance, however, there has been no final decision made regarding this. Such a deliverable would likely be developed iteratively and would not be expected for a few years or in the first edition.

[12] http://www.thegreengrid.org/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE

7.2.2. IEC TR 62635:2012

The IEC Technical Report (TR) provides a methodology for information exchange involving electronic and electrical equipment manufacturers and recyclers. The report also identifies how recyclability and recoverability rates should be calculated in order to provide accurate information to recyclers. It is envisaged that this information enables appropriate and optimized end of life (EoL) treatment operations, provides sufficient information to characterize activities at EoL treatment facilities.

7.2.3. ISO 7779:2010

The ISO 7779 standard specifies procedures for measuring and reporting the noise emission of information technology and telecommunications equipment.

7.3. Relevant national standards / initiatives

Note: The EU (voluntary) Code of Conduct on Data Centres and the EPA ENERGY STAR® Program Requirements for Computer Servers v2.0/v3.0 are not detailed further in this section as they have already been addressed in section 2.2.

7.3.1. United Kingdom: British Standards Institute ZZ/1 Publicly Available Specification (PAS) 141:2011

PAS 141 is a process management specification for the re-use of used and waste electrical and electronic equipment (UEEE and WEEE). The specification was developed by industry experts working with the UK Department for Business, Innovation and Skills (BIS).

The main aims of PAS 141 are to:

Improve the standards for the re-use and refurbishment of electrical and electronic equipment that has reached the end of its first useful life in the UK; and

Address the demand from consumers for assurance that the used electrical products they buy are electrically safe to use and functionally fit for purpose.

PAS 141 provides the following.

A framework for the testing, treatment and provision of re-use electrical and electronic equipment in the UK;

Reassurance that used equipment is electrically safe to use and functionally fit for purpose;

A method of differentiating legitimate exports from illegal exports of WEEE under the guise of being sent abroad for re-use.

A PAS 141 Certification Scheme was launched on the 27th February 2013.

7.3.2. United States: NSF/ANSI 426

(Linked to and IEEE 1680.4) The NSF 426 standard development process is a USA based initiative to develop a set of environmental criteria for servers which address multiple environmental impact categories. The final standard will be American National Standards Institute (ANSI) accredited.

The purpose of the NSF standard for servers is to establish product environmental performance criteria and corporate performance metrics that exemplify environmental leadership in the market. The scope of the standard is limited to "servers" that are covered under the ENERGY STAR Program Requirements for Computer Servers Version 2.0.

The standard provides a framework and consistent set of performance objectives for manufacturers in the design and manufacture of servers and server components. The standard establishes measurable criteria across multiple environmental impact categories including energy efficiency, management of substances, preferable materials use, product packaging, design for repair, reuse, and recycling, product longevity, responsible end-of-service/end-of-life management, life cycle assessments, and corporate responsibility.

Latest developments suggest that the NSF and IEEE 1680.4 standards will be combined into a single standard. Negotiations on this combination process are on-going at the time of writing.

7.3.3. United States: NIST Special Publication 800-88

The NIST document aims to assist in the development of effective media sanitization programmes with proper and applicable techniques and controls for sanitization and disposal decisions based on different levels of data security required.

The publication assists with decision making when media require disposal or reuse. It also provides guidance for information disposition, sanitization, and control decisions. The publication provides reference to applicable techniques and controls for data sanitisation based on different levels of data security needs.

7.3.4. United Kingdom CESG standards on data sanitation

CESG is the Information Security arm of the UK Government Communications Headquarters (GCHQ), and the National Technical Authority for Information Assurance within the UK. Their main role is to provide technical assistance concerning Information Security in Government.

There are three main standards the CESG is responsible for in the area of data sanitisation. These are:

CPA Security Characteristics for Data Sanitisation - Flash Based Storage

CAS Sanitisation Requirements Version 2.0 Nov 2014

HMG Information Assurance (IA) Standard No. 5 - Secure Sanitisation Version 5.0

The Her Majesty's Government (HMG) IA Standard No. 5 identifies how to destroy data depending on its sensitivity, where it is located and the media on which it is stored. The CAS Sanitisation Requirements Version 2.0 is a certification scheme to which commercial sanitisation services may subscribe, therefore demonstrating compliance with HMG IA Standard No. 5 when serving Government customers. The CPA Security Characteristics for Data Sanitisation - Flash Based Storage document includes requirements for sanitisation of all Flash-based storage media (e.g. solid state hard drives).

Note there are a number of other national initiatives addressing data sanitisation. These are not discussed in detail, as the NIST Publication 800-88 document provides a good overview of data sanitisation methods and procedures.

7.4. Relevant industry standards and international initiatives

There are already a number of measurement methods available for servers, storage and other equipment to measure and report energy use in an accurate and reproducible manner. However, there are large gaps in coverage and many are still not finalised. Regardless, their development to date has established a technical expertise in the industry that can be called upon in the formal development of transitional and harmonised standards.

7.4.1. SPEC: SPECpower_ssj2008

SPECpower is the initial rating tool developed by the SPEC group. It was earmarked for use in the first ENERGY STAR server specification, but unresolved complexities meant that SPEC recommended delaying on its inclusion. Efforts subsequently shifted to development of the SERT tool. SPECpower only measures efficiency under a very limited conditions of the SPEC ssj_2008 test which tests the CPU and RAM. In fact, ssj_2008 is now effectively a worklet for the hybrid workload that comprises part of the SERT tool). The power is measured at different server utilisation levels, from 0 to 100% and gives a power consumption level and performance rating at each level

The SPECpower tool is referenced in the Irish Triple E program. This aggregates the performance and power for low (10-30%), mid (40-60%) and high (70-100%) utilisation levels and produces three performance/power ratio. The specifications set minimum ratios for each utilisation level.

7.4.2. SPEC: SERT V1.1.1

The Server Efficiency Rating Tool (SERT tool) was created by the Standard Performance Evaluation Corporation (SPEC). SPEC is a non-profit organisation open to all parties but requires membership fees. SPEC has over 50 members which includes almost all the main ICT hardware manufacturers and a number of software and internet companies. There are also SPEC Associates and a Research Group which include approximately 100 other organisations, in particular universities in USA, Japan and Germany.

SERT is a software tool for measuring server energy efficiency. Central design considerations underpinning the SERT tool and making it a promising candidate for use in policy measures include; reproducibility of results, fairness, verifiability and usability.

The SERT tool is intended to be economical and easy to use with the minimum equipment and skill requirements. It has a graphical user interface for easy configuration, and after setup the process is automated to minimise the time necessary for testing. SERT is hardware and OS agnostic, meaning that it supports various hardware platforms and operating systems. and has the ability to run on a wide range of server specifications and configurations even as these continue to expand. Servers should be tested in their "as shipped" or "out of the box" state (although it may be necessary to pre-configure some RAID and other storage settings). Target run time is around five hours, although this will vary with server generations.

The SERT tool simulates a variety of common types of work via worklets. These are essentially software simulations of real working environments tailored to test discrete system components (e.g. processors, memory and storage) and subsystems (e.g. RAM and CPU). A range of worklets is necessary to ensure platform neutrality, as performance of different server architectures will vary with different workloads. For ease of comparability, the worklet results are normalised against results for an arbitrarily selected baseline server model.

Via the current worklet approach, SERT provides results representative of real working environments. One of these worklets is the ssj_2008 test from SPECpower. Each worklet provides a numerical output which can then be combined and interpreted into an overall pass/fail conclusion in relation to the requirements of a particular policy. Manufacturers are discouraged from quoting numerical values for specific worklets in isolation for marketing purposes as taken in isolation these values can be misrepresentative.

Results are provided in both machine (XML) and human readable (HTML/TXT) forms, accompanied by summary and detail reports. Purchase price for the SPEC software ranges from $900 (not for profit reduced rate) to $3,000. Charges are to cover the costs of providing support on the tool.

As previously referenced, there is an ISO standard under development which recognises but does not endorse the SERT tool as a possible option to measure efficiency. In addition, SERT has been used as a foundation standardised testing tool to support policy measures on server energy efficiency. Current policy interest in SERT includes the United States (US EPA / ENERGY STAR), Korea, China, and Australia / New Zealand. SPEC are working closely with ENERGY STAR toward the revision of the server specification Usually, the policy maker gathers the data, does the analysis and defines their approach and metric, and then SPEC can customise the tool accordingly.

For further details on SERT, please see the white paper in Appendix 4.

7.4.3. SNIA Emerald™ Power Efficiency Measurement Specification

Storage Networking Industry Association (SNIA) Emerald is a test method specification for storage equipment. It measures the power consumption under a variety of use cases to give an overall efficiency metric. Due to the large differences in storage equipment design compared to servers with no clear market sector to focus on, the metric has a very detailed classification system based on the designed size and use of the storage system.

Emerald is designed to allow comparison of products within the same classification. Where applicable, five main tests are performed using vdbench, the first 'four corners' are designed to test the extremes of performance under small random data accesses and streaming data. The results report power, data throughout rate and latency.

The fifth test is called the hot band which simulates more realistic usage where a variety of data accesses are performed but concentrated around bands of frequently accessed data. This requires a lot of set up and as a result testing can be very complex and take days.

In addition, idle power is measured as well as tests to check for the presence of COMs, but not their effectiveness. While the test measurement covers all use cases, some of these may not be relevant for the product be used.

Emerald does not produce as much test result data compared to SERT, and therefore there are fewer variables to consider. However, the test report template requires in depth description of the system configuration since this can affect performance. There is currently not enough data to analyse fully and determine how detailed a metric can be developed based on performance,

latency, throughout etc. which also fairly treats various configurations such as capacity. Results across different classes of products cannot be compared.

SNIA Emerald is developed and maintained by SNIA, a non-profit, international organisation of manufacturers, systems integrators, developers, systems vendors, industry professionals, and end users.

SNIA is being used by ENERGY STAR to develop specifications, however, the lack of data and difficulty in testing has meant progress so far has been slower than servers.

7.4.4. EPRI & Ecova: Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7

The Electric Power Research Institute (EPRI) is a widely used testing method for internal power supplies which is used in 80plus scheme, ENERGY STAR and other efficiency policies. This protocol was established in 2004 but has integrated the Server Test Protocol since 2008, including test methods for direct current (DC-DC) power supplies. It includes instructions to measure the power supply efficiency and power factor at various load levels.

7.4.5. ASHRAE TC 9.9 2011: Thermal Guidelines for Data Processing Environments, 4th Edition (2015) Equipment environmental specifications for air cooling

The American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) defines a range of humidity and temperature operating conditions which many data centre and IT equipment manufacturers adhere to. Since data centres will often contain a very mixed range of IT equipment types and brands, this helps ensure the entire system interoperates reliability. The Thermal Guidelines were first published in 2004, and then updated in 2008, 2011 and 2015 (and are likely to continue being updated with similar frequency). The current ranges, set by committee including data centre operators and equipment manufacturers, are shown in Figure 4 below:

Equipment Environment Specifications for Air Cooling							
	Product Operation[b,c]					Product Power Off[c,d]	
Class [a]	Dry-Bulb Temperature[e,g], °C	Humidity Range, Noncondensing[h, i, k, l]	Maximum Dew Point[k], °C	Maximum Elevation[e, j, m], m	Maximum Rate of Change[f], °C/h	Dry-Bulb Temperature, °C	Relative Humidity[k], %
Recommended (Suitable for all four classes; explore data center metrics in this book for conditions outside this range.)							
A1 to A4	18 to 27	–9°C DP to 15°C DP and 60% rh					
Allowable							
A1	15 to 32	–12°C DP and 8% rh to 17°C DP and 80% rh	17	3050	5/20	5 to 45	8 to 80
A2	10 to 35	–12°C DP and 8% rh to 21°C DP and 80% rh	21	3050	5/20	5 to 45	8 to 80
A3	5 to 40	–12°C DP and 8% rh to 24°C DP and 85% rh	24	3050	5/20	5 to 45	8 to 80
A4	5 to 45	–12°C DP and 8% rh to 24°C DP and 90% rh	24	3050	5/20	5 to 45	8 to 80
B	5 to 35	8% to 28°C DP and 80% rh	28	3050	N/A	5 to 45	8 to 80
C	5 to 40	8% to 28°C DP and 80% rh	28	3050	N/A	5 to 45	8 to 80

* For potentially greater energy savings, refer to the section "Detailed Flowchart for the Use and Application of the ASHRAE Data Center Classes" in Appendix C for the process needed to account for multiple server metrics that impact overall TCO.

Figure 4 - Summary of ASHRAE Thermal Guideline Classes

The temperature range has an impact on the design and energy consumption of the server, particularly the internal cooling system such as the heatsink and fans. A larger range can reduce the data centre cooling costs but since the internal computing components are often unchanged and the same amount of heat produced must still removed, it can shift the energy consumption into the server, and require larger, more energy consuming fans as well as larger, heavier heatsinks which may impact the overall lifecycle. This aspect, however, was already analysed in the

Lot 9 preparatory study, with particular regards to servers; it was concluded that the overall energy saving is significantly positive (see in particular Task 6).

7.4.6. ECMA: ECMA-74 13th edition (June 2015) (based on ISO 3741, ISO 3744, ISO 3745, ISO 11201)

This Ecma Standard specifies procedures for measuring and reporting the noise emission of information technology and telecommunications equipment. Densely installed server equipment in a data centre can emit high levels of noise. Best practice requires that a data centre (which is often considered to be an industrial space) should only be occupied for service and maintenance purposes, however, this may not always be true for SMEs.

7.4.7. IEEE 1680.4 Servers

The IEEE 1680 series are a USA based series of standards which focus on IT products and which include environmental performance criteria across multiple environmental impact categories. The standard defines environmental performance criteria for computer servers as defined in the ENERGY STAR Server specifications, including managed servers and blade servers, relating to reduction or elimination of environmentally sensitive materials, materials selection, design for end of life, lifecycle extension, energy conservation, end of life management, corporate performance, and packaging.

The IEEE series of standards (apart from IEEE 1680.4) are used to inform the development of EPEAT specifications. The developers of the IEEE 1680.4 and NSF 426 standards are working together to produce a single multi-attribute environmental standard for servers.

7.4.8. ANSI ATIS: 060015.2013 (TEER)

This is a network efficiency test method used to determine the efficiency of the networking component of servers or networking equipment. It provides a measure of the data throughput per unit of power. Network interfaces have traditionally been poorly energy managed and historically consumed the same amount of energy regardless of how much data was passing through. While this is a small proportion of server power for slower network interfaces, 10Gb ethernet interfaces can consume around 5W each.

The Telecommunications Energy Efficiency Ratio (TEER) is defined in the ANSI ATIS 060015.01.2014 standard. which addresses "Energy efficiency for telecommunication equipment: Methodology for measurement and reporting - Server requirements". This standard references SPEC and specifies how to measure network efficiency only for a server, if TEER is requested. It does not cover server efficiency in general.

8. Standards engagement

8.1. *Standards activities engaged with*

During this project, an effort was made to engage with the various standardisation processes. Activities are summarised in Table 5.

Table 5 – Standards engagement activities during project

Country / Region	Standardisation Group	Relevant initiatives	Contact
Europe	CEN-CENELEC-ETSI coordination group on Green Data Centres (CG GDC)	CLC/TC 215 JTC 1/SC 39 Mandate M/462 The Green Grid	Teleconference with CENELEC contact 7th August. Project team and European Commission attendance and presentation at 13th April coordination group meeting in London.
United States	US Environmental Protection Agency (EPA)	ENERGY STAR	Teleconferences 1st September 2015 and 29th January 2016, and remote attendance and presentation at meeting on 19th November to discuss server developments and data analysis in relation to SERT and active mode.
International	SPEC	SERT	Teleconferences 6th July and 20th January. Attendance and presentation at European SPEC Symposium on 18th March 2016 Also included the set up of a "Beta Testing" programme, see below.
China	CNIS	Chinese server metric development	Meeting 14th September 2015 to discuss activities and potential for coordination
Korea	Kemco	Korean server metric development	Emails September 2015 and March 2016 but did not identify correct person or no responses

8.2. *Beta testing / Evaluation Programme SERT v1.1.1*

8.2.1. Description

In order to engage stakeholders on the use of SERT, a beta testing programme was launched in collaboration with SPEC. This enabled stakeholders to have temporary/trial access to the version 1.1.1 SERT tool for free (it is normally necessary to purchase a license for $2,800).

8.2.2. Objectives

The intention of the trial was to provide stakeholders with the opportunity to become familiar with the SERT testing tool, at the same time as providing the Lot 9 study with insights on:

- Ease of use,
- Procedural refinements/clarifications
- Potential directions for development of an EU metric based on SERT, should such an approach be deemed appropriate.
- Focus areas included:
- Ease of set-up and use
- Execution experiences
- Run rule clarifications
- Load characteristics
- Measurement characteristics
- Bug discovery

8.2.3. Results

Unfortunately, despite publicising the trial on multiple occasions, no stakeholders came forward to take part in the trial – this is possibly because most of them are already involved in SERT development. However, the trial license proved useful to the project team in the testing activity carried out and detailed in the testing report contained in Appendix 3.

9. Conclusions

There is much standardisation activity currently underway in the area of server standards. Enterprise storage standards for energy efficiency measurement are not so advanced. This project has endeavoured to follow, map and facilitate the standards process related to lot 9 products, and the information in it should be up to date as at April 2016.

In addition to the gap analysis activity, the project has also contributed an in-depth analysis into server metrics, documented in the white paper in Appendix 4. In order to gain deeper practical insights on testing experiences with the SERT tool, Intertek also carried out testing on a number of servers. Insights and recommendations stemming from this activity are detailed in the testing report document contained in Appendix 3.

A list of potential standards references that could be used to support any server policy requirements, should these be defined, is contained in Appendix 2. On going priority standardisation activities that will continue once this project has completed include:

Table 6 - Ongoing priority standardisation activities

Standardisation area	Observations	Key standards activities and expected delivery
Server idle / active metric approach	SPEC and ENERGY STAR will continue working on this in 2016. The European Commission will provide input to the ENERGY STAR specification development process via the EU-US ENERGY STAR agreement.	Server active power test methodology, SERT SPEC and ENERGY STAR Servers v3.0 expected to formally commence in early 2016 and complete at some point in late 2016 or early 2017. Server KPIs and ITEE in ISO 30314-4 and CENELEC EN 50600 expected to deliver in early 2017.
Storage idle / active metric approach	Product complexity poses significant challenges to develop and finalise metrics.	SNIA Emerald and ENERGY STAR Storage v2.0 specification are expected post 2016
Material efficiency considerations	NSF/IEEE are the most active in the standard development process for material efficiency considerations. The EU joint research centre is also active in the area.	NSF 426/IEEE 1680.4 Standard for Servers is expected to be delivered in 2016.

Appendix 1: Coverage and status of key standards / initiatives for enterprise servers and data storage

Body	Standard	Topic	Relevant parameters addressed	Status	Ecodesign suitability	Relevant products
European standardisation deliverables (EN)						
ETSI EE2	EN 300 132-3	Power supply interface	None identified	Published. Undergoing revision.	Not suitable	N/A
ETSI EE	EN 300 019-1 series	Environmental conditions for telecoms	Operating temperature and humidity	Published.	Not data centre appropriate. ASHRAE more suited to data centres	ES, DS
CENELEC TC108	EN 62075 (IEC 62075)	Environmentally conscious design	Acoustic noise Removability of external enclosures, PCBs, processors, data storage devices and batteries. Ease of dismantling, reuse and recycling at the end-of-life. Data sanitisation	Published.	General concepts. Insufficient detail for product specific purposes, except possibly data deletion.	ES, DS
CENELEC TC108	EN 62018 (IEC 62018)	Power consumption of ICT	Energy proportional operation (dynamic range) Lower power modes (other than off mode)	Published	Defines power modes, but insufficient detail for product-specific purposes.	ES, DS
CENELEC TC 215	EN 50600 Series	Data centre design & operation	Overall energy performance (all power modes or TEC type approach)	In draft (estimated 2016)	May be suitable, but depends on level of detail delivered in EN.	ES, DS
International standardisation deliverables (ISO / IEC)						
ISO/IEC JTC 1/SC 39	ISO/IEC 30134-4	IT energy efficiency and energy utilisation	Active State (power demand / rating)	In draft (estimated early 2017)	May be suitable – depends on level of detail and appropriateness of performance levels addressed. Current	ES

					focus on efficiency at max load (not common use of servers)	
IEC	IEC TR 62635:2012	Guidelines for end-of-life information and recyclability rate calculation.	Removability of external enclosures, PCBs, processors, data storage devices and batteries with common tools Ease of dismantling, reuse and recycling at the end-of-life.	Published	Yes. Whilst a technical report rather than a full standard, this provides adequate approach to be applicable to the product groups.	ES, DS
ISO	ISO 7779:2010	Measurement of airborne noise emitted by information technology and telecommunications equipment	Acoustic noise	Published	Yes. Established standard in use by industry addressing many aspects	ES, DS
National standards / initiatives						
United Kingdom BSI ZZ/1	PAS 141:2011	Reuse of used and waste electrical and electronic equipment (process management).	Removability of external enclosures, PCBs, processors, data storage devices and batteries with common tools Ease of dismantling, reuse and recycling at the end-of-life. Data sanitisation	Published	Likely to be suitable as reference material.	ES, DS
EU European Commission (voluntary)	2010 Best Practices for the EU Code of Conduct on Data Centres	Energy efficient best practice and voluntary targets	None identified	Published	Not suitable. Data centre not product-specific focus.	N/A
United States, IEEE/ANSI	IEEE 1680.4	Environmental impacts of servers	Same as NSF 426	Unpublished (unclear if completion likely)	Applicability unclear as access to standard not possible unless within working group. Will be merged with the NSF 426 standard	ES
United	ENERGY STAR®	Energy efficiency of	Active State (power demand / rating)	Published.	Established standard in	ES

States EPA	Program Requirements for Computer Servers v2.0/v3.0	servers	Idle State (power demand/ rating) Overall energy performance (all power modes or TEC type approach)	Under revision (estimated 2016).	use by industry addressing most aspects	
United States EPA	ENERGY STAR® specification for data centre storage v1.0	Energy efficiency of data centre storage	Capacity Optimizing Methods (COMs)	Published. Under revision (estimated 2017).	Suitable to address COMs but not other energy efficiency aspects as yet.	DS
United States NSF/ANSI	New Standard - NSF/ANSI 426, Draft 1, Issue 1	Environmental impacts of servers	Removability of external enclosures, PCBs, processors, data storage devices and batteries with common tools Ease of dismantling, reuse and recycling at the end-of-life. Critical raw material (CRM) content Postconsumer recycled content of CRM Replacement components availability Reduction of surplus parts by default Hardware functionality testing software tools	In draft (estimated 2016).	Likely to be suitable. To be merged with IEEE 1680.4.	ES
United States, ANSI ATIS	060015.01. 2014	Measurement and reporting - Server network energy efficiency: telecommunications Energy Efficiency Ratio (TEER)	None identified	Published.	Network efficiency (i.e. data throughput per unit power) was not identified as a key parameter. Approach does not cover server efficiency in general	ES
United States NIST	NIST Special Publication 800-88 Revision 1	Guidelines for Media Sanitization	Data sanitisation	Published	Yes. Established standard in use by industry addressing most aspects.	ES, DS
United Kingdom CESG	CPA Security Characteristics for Data Sanitisation - Flash Based	Data sanitisation	Data sanitisation	Published	Yes. Established standard in use by industry addressing most aspects, although NIST may provide a more	ES, DS

	Storage CAS Sanitisation Requirements Version 2.0 Nov 2014 HMG IA Standard No. 5 - Secure Sanitisation Version 5.0				succinct source.	
Industry standards / International initiatives						
SPEC	SPECpower_ssj2008	Energy efficiency rating for servers	Active State (power demand / rating) Idle State (power demand/ rating) Energy proportional operation (dynamic range) Overall energy performance (TEC type approach)	Published	Superseded by SERT, this tool addressed some aspects, especially idle state and CPU-RAM active power demand.	ES
SPEC	SERT V1.1.0	Energy efficiency rating for servers	Active State (power demand / rating) Idle State (power demand/ rating) Energy proportional operation (dynamic range) Overall energy performance (TEC type approach)	Published. (Next revision estimated 2017 or later)	Yes. Established standard in use by industry addressing most aspects.	ES
SNIA	SNIA Emerald™ Power Efficiency Measurement Specification	Power demand of storage equipment	Active State (power demand / rating) Idle State (power demand/ rating) Energy proportional operation (dynamic range) Overall energy performance (TEC type approach) COMs	Published.	Yes. Established standard in use by industry addressing most aspects.	DS
EPRI & Ecova	Generalized Test Protocol Revision 6.7	Energy efficiency of internal power supplies	Power Supply Efficiency Power Supply Power factor	Published	Yes. Established standard in use by industry.	ES, DS
ASHRAE	Thermal Guidelines 4th Edition, air cooling	Thermal Guidelines for Data Processing Environments	Operating temperature and humidity	Published	Yes. Established standard in use by industry.	ES, DS

	environmental specifications					
ECMA	ECMA-74 13th edition (June 2015)	Noise emissions	Acoustic noise	Published	Yes. Established standard in use by industry. Based on ISO 3741, ISO 3744, ISO 3745, ISO 11201.	ES, DS
ITU-T SG5 Q17		Recommendations for energy efficiency metrics, best practice and measurement for telecommunication equipment	None identified	In draft (Expected 2016)	ITU is working to identify gaps and assist standardisation and harmonisation of existing efforts.	ES, DS

Appendix 2: Server and Storage Standards Listing

1. Server Standards

Parameter	Source	Reference Test Method / Title	Notes
Active State Energy Efficiency (incorporating Idle State and Energy proportional design / Dynamic Range considerations)	SPEC	Server Efficiency Rating Tool (SERT) V1.1.1	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz). The metric for server energy efficiency based upon SERT tool test results is listed below: **High utilisation** $server\ efficiency = \left\{\frac{0.6}{CPU\ eff} + \frac{0.4}{memory\ eff}\right\}^{-1}$ **Low utilisation** $server\ efficiency = \boxed{\left(\frac{idle\ power}{\max power} + 1.5\right)}_{Dynamic\ range} \times \left\{\frac{0.6}{CPU\ eff} + \frac{0.4}{memory\ eff}\right\}^{-1}$ where at a workload level: $workload\ efficiency = no.worklets \times \left\{\sum \frac{1}{worklet\ efficiency}\right\}^{-1}$ and at a worklet level: $worklet\ efficiency = \frac{\sum performance\ at\ each\ utilisation\ level}{\sum power\ at\ each\ utilisation\ level}$ For the purposes of supporting information, the average server performance can be calculated as shown below: $server\ perf = \left\{\frac{0.6}{CPU\ perf} + \frac{0.4}{memory\ perf}\right\}^{-1}$ $server\ power = \left\{\frac{0.6}{CPU\ perf} + \frac{0.4}{memory\ perf}\right\}^{-1} \times \left\{\frac{0.6 \times CPU\ power}{CPU\ perf} + \frac{0.4 \times memory\ power}{memory\ perf}\right\}$

Parameter	Source	Reference Test Method / Title	Notes
			Where at the workload level: $workload\ perf = no.worklets \times \left\{\sum \frac{no.utilisation\ levels}{worklet\ perf\ sum}\right\}^{-1}$ $workload\ power = \left\{\sum \frac{no.worklet\ utilisation\ lvl}{worklet\ perf\ sum}\right\}^{-1} \times \sum \frac{worklet\ power\ sum}{worklet\ perf\ sum}$
Power Supply Efficiency	EPRI and Ecova	Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz) Initiatives specifying requirements for this parameter include: • 80 Plus programme (EPRI and Ecova) • Ecodesign Regulation (EU) No.617/2013
Power Supply Power Factor	EPRI and Ecova	Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz) Initiatives specifying requirements for this parameter include: • 80 Plus programme (EPRI and Ecova) • Ecodesign Regulation (EU) No.617/2013
Operating temperature		Not available	There are no specific test methods for testing this aspect in relation to the operation of a product, but a method for measuring temperatures during testing is provided in: • CENELEC EN 50564:2011 Electrical and electronic household and office equipment. Measurement of low power consumption In addition, initiatives specifying potential ranges for this parameter in the context of a data centre include: • American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) Thermal Guidelines for Data Processing Environments, 3rd Edition, in Table 2.3 specify different conditions for data centre operation.

Parameter	Source	Reference Test Method / Title	Notes
Operating humidity		Not available	There are no specific test methods for testing this aspect in relation to the operation of a product, but initiatives specifying requirements for this parameter include: • American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) Thermal Guidelines for Data Processing Environments, 3rd Edition, in Table 2.3 specify different conditions for data centre operation.
Acoustic noise	ECMA	Standard ECMA-74 Measurement of Airborne Noise emitted by Information Technology and Telecommunications Equipment 13th edition (June 2015)	Initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers
Secure deletion of data	CESG	HMG IA Standard No. 5 - Secure Sanitisation Version 5.0	Initiatives specifying requirements for this parameter include: • CPA Security Characteristics for Data Sanitisation - Flash Based Storage • CPA Security Characteristic Overwriting Tools for Magnetic Media Version 2.1 • CAS Sanitisation Requirements Version 2.0 Nov 2014
Removability of external enclosures/casings to increase material recovery rate		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive)

Parameter	Source	Reference Test Method / Title	Notes
			• Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Removability of printed circuit boards (including main boards and memory cards), processors, data storage devices (such as HDD or SSD) and batteries		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive) • Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Ease of dismantling, reuse and recycling at the end-of-life of the product.		Not available	The closest solution to a test procedure for design for recyclability is contained in • European Commission: COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) There are no specific test methods for the other aspects, but initiatives specifying requirements for this parameter for ICT equipment include: • PAS 141:2011 - Reuse of used and waste electrical and electronic equipment (UEEE and WEEE) – Process

Parameter	Source	Reference Test Method / Title	Notes
			management – Specification • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive) Article 15: Information for treatment facilities • Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Critical raw material (CRM) content		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • Draft NSF 426 Servers / IEEE 1680.4
Postconsumer recycled content of CRM		Not available	No test methods or initiatives specifying requirements for this parameter were identified.
Firmware availability and compatibility		Not available	No test methods or initiatives specifying requirements for this parameter include: • Draft NSF 426 Servers / IEEE 1680.4
Restriction of Substances of Very High Concern (SVHCs)	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)

Parameter	Source	Reference Test Method / Title	Notes
Restrictions on the presence of specific hazardous substances	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)
Restrictions based on CLP hazard classifications	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)
Data storage drive reliability and protection	Telcordia	SR-332 - Reliability Prediction Procedure for Electronic Equipment	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)

Storage Standards

Parameter	Source	Reference/Title	Notes
Active State Energy Efficiency (incorporating Idle State and Energy proportional design / Dynamic Range considerations)	SNIA	SNIA Emerald™ Power Efficiency Measurement Specification Version 2.0.2: Section 7.3 General Requirements and Definitions and Section 7.4.3: Active Test	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz)
Power Supply Efficiency	EPRI and Ecova	Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz) Initiatives specifying requirements for this parameter include: • 80 Plus programme (EPRI and Ecova) • Ecodesign Regulation (EU) No.617/2013
Power Supply Power Factor	EPRI and Ecova	Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7	Testing should be conducted at an appropriate EU voltage and frequency (e.g. 230v, 50Hz) Initiatives specifying requirements for this parameter include: • 80 Plus programme (EPRI and Ecova) • Ecodesign Regulation (EU) No.617/2013
Capacity Optimizing Methods (COMs)	US EPA and SNIA	US EPA ENERGY STAR v1.0 specification for data centre storage and SNIA EmeraldTM Power Efficiency Measurement Specification Version 2.0.2: Section 7.4.5 Capacity Optimization Test	
Operating temperature		Not available	There are no specific test methods for testing this aspect in relation to the operation of a product, but a method for measuring temperatures during testing is provided in: • CENELEC EN 50564:2011 Electrical and electronic household and office equipment. Measurement of low power

Parameter	Source	Reference/Title	Notes
			consumption In addition, initiatives specifying potential ranges for this parameter in the context of a data centre include: • American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) Thermal Guidelines for Data Processing Environments, 3rd Edition, in Table 2.3 specify different conditions for data centre operation.
Operating humidity		Not available	There are no specific test methods for testing this aspect, but initiatives specifying requirements for this parameter include: • American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) Thermal Guidelines for Data Processing Environments, 3rd Edition, in Table 2.3 specify different conditions for data centre operation.
Acoustic noise	ECMA	Standard ECMA-74 Measurement of Airborne Noise emitted by Information Technology and Telecommunications Equipment 13th edition (June 2015)	Initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers
Secure deletion of data	CESG	HMG IA Standard No. 5 - Secure Sanitisation Version 5.0	Initiatives specifying requirements for this parameter include: • CPA Security Characteristics for Data Sanitisation - Flash Based Storage • CPA Security Characteristic Overwriting Tools for Magnetic Media Version 2.1 • CAS Sanitisation Requirements Version 2.0 Nov 2014
Removability of external enclosures/casings to increase material recovery rate		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers

Parameter	Source	Reference/Title	Notes
			• COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive) • Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Removability of printed circuit boards (including main boards and memory cards), processors, data storage devices (such as HDD or SSD) and batteries		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive) • Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Ease of dismantling, reuse and recycling at the end-of-life of the product.		Not available	The closest solution to a test procedure for design for recyclability is contained in • European Commission: COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers

Parameter	Source	Reference/Title	Notes
			(January 2016) There are no specific test methods for the other aspects, but initiatives specifying requirements for this parameter for ICT equipment include: • PAS 141:2011 - Reuse of used and waste electrical and electronic equipment (UEEE and WEEE) – Process management – Specification • COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016) • European Community Directive (2012/19/EU) - The Waste Electrical and Electronic Equipment Directive (WEEE Directive) Article 15: Information for treatment facilities • Article 15: Information for treatment facilities" requirements of the EU WEEE Directive (2012/19/EU). • Draft NSF 426 Servers / IEEE 1680.4 (and 1680.1 Computers)
Critical raw material (CRM) content		Not available	There are no specific test methods for testing this aspect but initiatives specifying requirements for this parameter for ICT equipment include: • Draft NSF 426 Servers / IEEE 1680.4
Postconsumer recycled content of CRM		Not available	No initiatives specifying requirements for this parameter were identified.
Firmware availability and compatibility		Not available	Initiatives specifying requirements for this parameter include: • Draft NSF 426 Servers / IEEE 1680.4

Parameter	Source	Reference/Title	Notes
Restriction of Substances of Very High Concern (SVHCs)	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)
Restrictions on the presence of specific hazardous substances	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)
Restrictions based on CLP hazard classifications	CENELEC	EN 62321 - Determination of certain substances in electrotechnical products	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)
Data storage drive reliability and protection	Telcordia	SR-332 - Reliability Prediction Procedure for Electronic Equipment	Initiatives specifying requirements for this parameter include: • COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers (January 2016)

Organisation	Programme	URL(s)
American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE)	American Society of Heating, Refrigerating, and Air-conditioning Engineers (ASHRAE) Thermal Guidelines for Data Processing Environments, 3rd Edition, in Table 2.3	https://www.ashrae.org/resources--publications/bookstore/datacom-series
British Standards Institute (BSI)	PAS 141:2011 - Reuse of used and waste electrical and electronic equipment (UEEE and WEEE) – Process management – Specification	http://shop.bsigroup.com/en/ProductDetail/?pid=000000000030245346
CENELEC	EN 50564:2011 Electrical and electronic household and office equipment. Measurement of low power consumption	http://shop.bsigroup.com/ProductDetail/?pid=000000000030192768
CENELEC	Mandate 545 - Commission Implementing Decision on a standardisation request to the European standardisation organisations as regards computers and computer servers in support of the implementation of Commission Regulation (EU) No 617/2013 of 26 June 2013, implementing Directive 2009/125/EC of the European Parliament and of the Council with regard to ecodesign requirements for computers and computer servers	http://ec.europa.eu/growth/tools-databases/mandates/index.cfm?fuseaction=search.detail&id=566
CENELEC	EN 62321-1:2013 Determination of certain substances in electrotechnical products. Introduction and overview	http://shop.bsigroup.com/ProductDetail/?pid=000000000030229131
Communications-Electronics Security Group (CESG)	CPA Security Characteristics for Data Sanitisation - Flash Based Storage	https://www.cesg.gov.uk/.../data_sanitisation_flash_based_storage.pdf

Communications-Electronics Security Group (CESG)	CAS Sanitisation Requirements Version 2.0 Nov 2014	https://www.cesg.gov.uk/publications/Documents/cas_sanitisation_service_requirement.pdf
Communications-Electronics Security Group (CESG)	CPA Security Characteristic Overwriting Tools for Magnetic Media Version 2.1	https://www.cesg.gov.uk/content/files/protected_files/document_files/CPA%20SC%20Overwriting%20Tools%20for%20Magnetic%20Media%20v2-1.pdf
Communications-Electronics Security Group (CESG)	HMG IA Standard No. 5 - Secure Sanitisation Version 5.0	http://www.cesg.gov.uk/aboutus/contactus/Pages/index.aspx
ECMA	Standard ECMA-74 Measurement of Airborne Noise emitted by Information Technology and Telecommunications Equipment 13th edition (June 2015)	http://www.ecma-international.org/publications/standards/Ecma-074.htm
ECMA	ECMA 370: 5th Edition / June 2015 - TED THE ECO DECLARATION	http://www.ecma-international.org/publications/standards/Ecma-370.htm
Ecova	Generalized Test Protocol for Calculating the Energy Efficiency of Internal Ac-Dc and Dc-Dc Power Supplies Revision 6.7	http://www.plugloadsolutions.com/80PlusPowerSupplies.aspx
European Commission (EC)	Waste Electrical and Electronic Equipment Directive (WEEE)	http://ec.europa.eu/environment/waste/weee/index_en.htm
European Commission (EC)	European Rare Earths Competency Network (ERECON)	http://ec.europa.eu/growth/sectors/raw-materials/specific-interest/erecon/index_en.htm
European Commission (EC)	Raw Materials Initiative	http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=C:520
European Commission (EC)	COMMISSION DECISION of 9 June 2011 on establishing the ecological criteria for the award of the EU Ecolabel for personal computers	http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32011D0337&from=EN

European Commission (EC)	COMMISSION DECISION of XXX establishing the ecological criteria for the award of the EU Ecolabel for personal, notebook and tablet computers Date: 22 Jan 2016	http://ec.europa.eu/transparency/regcomitology/index.cfm?do=search.documentdetail&GhBtyQRBEECxQIbXLck+VK7dtxVGV2+ZcjrWVPNcE3IjYPYKb9Q5l4ombI50qVxG
International Standards Organisation (ISO)	ISO/IEC JTC 1/SC 39 - Sustainability for and by Information Technology	http://www.iso.org/iso/home/store/catalogue_tc/catalogue_tc_browse.htm?commid=654019&development=on
Standard Performance Evaluation Corporation (SPEC)	Server Efficiency Rating Tool (SERT)	https://www.spec.org/sert/
Telcordia	SR-332 Reliability Prediction Procedure for Electronic Equipment	http://telecom-info.telcordia.com/site-cgi/ido/docs.cgi?ID=SEARCH&DOCUMENT=SR-332&#ORD
US Environmental Protection Agency (EPA)	ENERGY STAR Enterprise Servers Specification Version 2.0	https://www.energystar.gov/products/spec/enterprise_servers_specification_version_2_0_pd

Appendix 3: Practical Insights on SERT™ testing for Enterprise Servers

Insights from testing carried out under the Ecodesign Technical Assistance Study on Standards for Lot 9 Enterprise Servers and Enterprise Data Storage

The authors would like to thank Broadberry for their kind collaboration on this testing activity, without which this report would not have been possible.

TABLE OF CONTENTS

APPENDIX 3: PRACTICAL INSIGHTS ON SERT™ TESTING FOR ENTERPRISE SERVERS

1. TESTING SCOPE 2
   - *1.1. Testing logistics* 3
   - *1.2. Testing equipment* 3
   - *1.3. Products tested* 4
   - *1.4. SERT test execution* 4
   - *1.5. Test conditions* 4
   - *1.6. Testing Insights* 5
   - *1.7. Testing costs* 6
2. GUIDANCE ON TESTING 7
   - *2.1. Testing set-up* 7
   - *2.2. Testing implementation* 7
3. INTERPRETATION OF RESULTS 9
4. SUMMARY RESULTS 10
   - *4.1. Run 1- Full Test* 10
   - *4.2. Run 2 - Consistency check part-run* 11
   - *4.3. Run 3 – Full Test* 12
   - *4.4. Run 4 – Full Test* 13

List of Figures

Appendix 3 Figure 1 Schematic for SERT controller and system under test (SPEC) ... 5

List of Tables

Appendix 3 Table 1 - Products tested 4

Appendix 3 Table 2 - Summary of testing insights 5

Appendix 3 Table 3- Costs of testing 6

1. Testing scope

The focus of this testing activity was upon enterprise servers. Storage testing was discounted due to the complexities, time overhead and immaturity of the test methodologies. Whilst the testing only addresses a subset of server models of one manufacturer, the focus of the activities was not on gathering a representative data set, but rather on gaining insights into the following (where logistically possible):

- **Ease of use:** Analysis of the test process itself, to consider ease of use and provide insights on potential clarifications or refinements to testing implementation. It is the intention of SPEC that carrying our SERT testing should not require highly technical expertise.
- **Real-world relevance:** Consideration of the extent to which the tool results can be considered representative of real-world usage.
- **Repeatability:** Investigation into repeatability via multiple tests on a single sample, and comparative tests on multiple samples of the same model.
- **Configuration definition:** Relative variations in SERT results with different product configurations and consideration of options for definition of configurations for testing.
- **Storage:** Variability in results due to different storage media – for example PCIe very high performance SSDs.
- **Game-ability:** Assessment of how easily results can be influenced by testing set-up in order to achieve improved results.
- **Test conditions:** Consideration of the influence of environmental parameters / test conditions i.e. temperature, humidity.
- **Testing costs:** Assessment of resources necessary, in terms of time and financial burden in order to carry out testing using the tool.

1.1. Testing logistics

Two server products provided to Intertek by the company Broadberry were tested between March and May 2016. All testing was carried out by experienced technicians. Some tests were carried out at the Intertek laboratories in Milton Keynes, UK. Other tests were carried out at Broadberry premises located in London.

1.2. Testing equipment

- Power analyser: Yokogawa WT310
- AC stabilised and conditioned AC mains Power Supply: Kikusui PCR 1000L, used at:
  - Voltage: 230V.a.c. ±5%
  - Frequency: 50Hz ± 1%
  - THD of the voltage waveform: < 5%
- Thermometer: Temperature@lert USB TM-STD30
- Humidity meter: Vaisala HMI41
- Air speed meter: Airflow TA430

The equipment necessary for testing was readily available in the testing laboratory, with the exception of a specific approved temperature sensor (TM-SDT30), which had to be purchased from the USA at a cost of $200.

It is estimated that the total cost for a manufacturer to purchase the testing equipment listed above would be in the region of 9,000 Euros, although if manufacturers already test products in house, they may only need to purchase the thermometer and air speed meter at a cost of under 2,000 Euros. In addition to the equipment cost, it is also necessary to purchase the SERT software at a cost of approximately 2,450 Euros. Further details on the cost of testing are contained in Appendix 3 Table 3.

1.3. Products tested

Tests included the latest "basic building brick" server, as well as a larger system worth over £20,000.

Details of the products tested are contained in Appendix 3 Table 1 below:

Sample number	Model number	CPU		Memory		Storage	
		Number of CPU cores	Frequency (GHz)	Memory Modules / dimms (MB)	RAM (GB)	Number of drives	Type of drive
1	X9SRE/X9SRE-3F/X9SRi/X9SRi-3F	4	3.7	4	64	3	HDD
2	X10DRi	8	3.2	4	64	10	HDD

Appendix 3 Table 1 - Products tested

1.4. SERT test execution

SERT tests can be carried out directly by manufacturers or on manufacturer's behalves by an external testing laboratory.

In the case of the US ENERGY STAR label, prior to associating the label with any server product, it is necessary to obtain written certification of ENERGY STAR qualification from an EPA recognised Certification Body based on testing in an EPA recognised testing laboratory[13].

For EU ENERGY STAR registered products and for ecodesign conformity purposes, testing by certified bodies is not necessary and can be carried out directly by manufacturers, so the coverage of certified laboratories in Europe is lower.

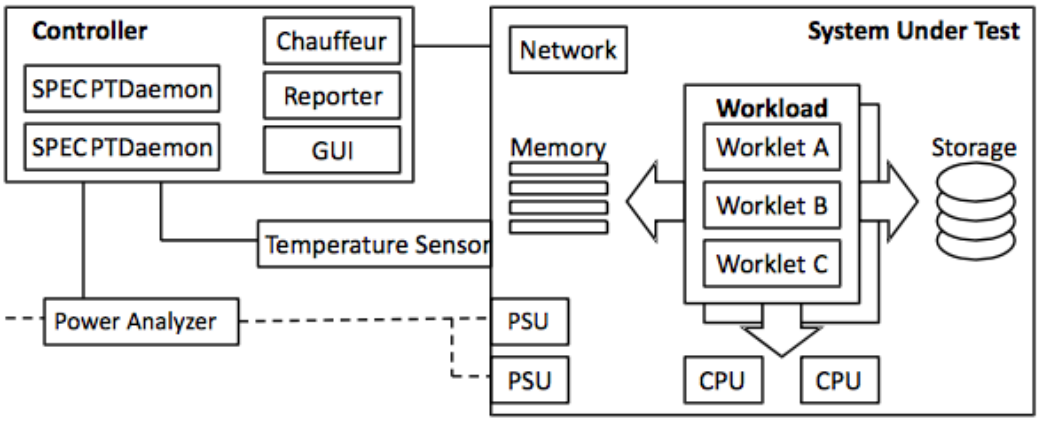
1.5. Test conditions

It is important that the physical test environment is representative of typical user environments. The temperature range can have an impact on the design and energy consumption of the server, particularly the internal cooling system such as the heatsink and fans. In particular, unusually low temperatures (below 20°C for a data centre) may result in artificially lower power demand and improved performance during the tests. The SERT tool must therefore be run within constrained environmental conditions, specified as follows:

- Ambient temperature lower limit: 20°C
- Ambient temperature upper limit: within documented operating specification of the SUT (but it is likely that servers will be tested as close to the lower limit as possible as this is where they perform most efficiently).
- Elevation and Humidity: within documented operating specification of the SUT
- No overt direction of air flow in the vicinity of the measured equipment in a way that would be inconsistent with normal data centre practices.

[13] A list of EPA-recognized laboratories and certification bodies can be found at https://www.energystar.gov/index.cfm?fuseaction=recognized_bodies_list.show_RCB_search_form

Compliance with these conditions is validated as shown in Appendix 3 Figure 1, by the use of a temperature sensor in the testing rig.



Appendix 3 Figure 1 Schematic for SERT controller and system under test (SPEC)

1.6. Testing Insights

Key testing observations are detailed in Appendix 3 Table 2:

Testing consideration	Observation
Ease of Use	After initial set-up, and with some modifications of the Graphical User Interface (GUI) it is clear that the tests can be carried out by an individual with novice skills level. The auto detection of configuration on the first System Under Test (SUT) was not very effective, but worked well on the second SUT. If detection does not work correctly, parameters can be corrected.
Repeatability	Additional runs were undertaken on the same sample to test repeatability. These results suggest a high level of repeatability in the worklet summary graphical results. A full test suite (valid run) repeated after several days of other testing on Sample 1 showed a maximum deviation of less than 1.5%in the key summary efficiency scores of CPU, Storage, Hybrid, Memory and Idle.
Real-world relevance	It is considered that results provide a reasonable approximation of the server usage under each workload, except for storage. There is a barrier to storage worklets providing representative results due to the pre-requisite of RAID disabling. There is also potential for very large variations in storage efficiency results as there are such differences between energy performance of SSD, SSD PCIe and HDD.
Game-ability:	GUI discovery data automatically qualifying the SUT technical and physical characteristics can be modified and has to be checked and corrected where errors occur. From the testing performed for this report, there is no indication that the modification of this data or errors in discovery could modify the efficiency score generated for the SUT.
Test conditions	Testing carried out at significantly different input airflow temperatures (averaging 20 °C and 30 °C) indicated that there was no significant impact on the values determined by the SERT for the worklet results summary.

Appendix 3 Table 2 - Summary of testing insights

In summary, the benchmark appears to hold a great deal of potential for the assessment of the energy efficiency of servers. A close collaboration with SPEC on tool development could ensure the suitability of the tool to support any energy efficiency initiatives on servers within the EU.

1.7. Testing costs

Server manufacturers can choose either to outsource their testing to an external lab or test internally in their own laboratory set up. The option chosen would depend on the number of models and configurations to be tested, balanced against the cost of equipment, software licenses and technician time.

There is a lack of established information on the commercial costs of testing in an external laboratory, but one processor manufacturer has indicated their willingness to test servers containing their processors on behalf of the manufacturers for potentially a relatively small (200 to 300 EUR) fee.

For in-house testing, a breakdown of the resources required in terms of equipment and laboratory technician time are outlined in Appendix 3 Table 3.

Appendix 3 Table 3- Costs of testing

Type of cost	Frequency	Detail	Number of	Approx. total cost (EUR)
Equipment	One off	Power analyser	1	€3,700
Equipment	One off	Power Supply (AC mains conditioner providing standard voltage and harmonic content)	1	€3,500
Equipment	One off	Thermometer & Humidity meter	1	€1,000
Equipment	One off	Air speed meter	1	€900
Software License	One off	SERT software purchase		€2,450
Calibration	Annual	Per measurement instrument (varies by instrument and source)	1	€1,300
			TOTAL initial outlay	**€12,850**
Labour	Per server	Set-up time	0.75 days	€375
Labour	Per server	Testing time	1.20 days	€600
Labour	Per server	Documentation	0.25 days	€125
			TOTAL per server cost	**€1,100**

Note: Cost of technician time is based upon a 7 hour day at a typical rate of 500 Euros. Technician time includes a full storage drive configuration check (e.g. examination of RAID settings and reconfiguration as required), and installation of SERT, Java, and measurement instrumentation software. Labour associated with testing time assumes a confirmatory short worklet run (e.g. "storage random") is performed and delivered with viable results before a complete SERT run. It is assumed that a complete run monitored occasionally allows the technician to perform other activities.

2. Guidance on Testing

The following guidance is based on the Server Efficiency Rating Tool (SERT) User Guide 1.1.1 dated January 2016, and on the following equipment used during testing :

- An Intel® Celeron™ N2050 1.6 GHz based PC running Microsoft ® Windows 10™ 64-bit as the controller, which sends the worklets to the SUT and interfaces with the power meter and temperature sensor.
- A Yokogawa WT310e power meter with a National Instruments GPIB to USB interface to measure the power consumption of the SUT
- A temp@lert sensor TM-STD30 for monitoring the minimum temperature at the air inlet of the SUT

2.1. Testing set-up

Controller setup

- The controller PC does not need to be highly specified, but needs to be stable, be set up (hardware interface and drivers) to communicate with the power meter and thermometer (e.g. via USB) and be able to communicate with the SUT via a 1Gb/s network interface.

SUT setup

- RAID must be disabled
- Write caching must be disabled on all storage drives
- The firewall must be disabled
- For a 64-bit environment, the SUT must have « Lock Pages in Memory » enabled

Controller & SUT common setup

- The controller PC and the SUT must have the SERT software installed, within the folder structure as specified in the User Guide.
- The controller PC and the SUT are required to have an installation of Java Runtime Environment. Note that the User Guide provides an example of setting up the Java path on the SUT, which implies the development version is required (the path contains the string « JDKL ») and this is not the case. The Java installation path will contain the string « JRE ».

2.2. Testing implementation

The user must run the host environment on the SUT (a DOS window will open and remain opened during the testing).

The GUI is the most user-friendly way of using the SERT software. The Java Runtime Environment is required for this. The GUI will guide the user logically through the process.

The software will auto detect the SUT configuration. The auto-detected information will be preceded with an underscore (« _ »). It is not clear in the documentation, but the user must remove these underscores as they check the auto-detected details, otherwise the test will be invalidated.

The software will allow configuration settings in order to communicate with the test equipment and verify that the connection is working.

During the power measurements, the ranges used must be known in order to calculate the uncertainties of measurement per measurement made. Most power meters feature an auto-ranging function, but during the range-changing process, data is not captured. To overcome this, the SERT

software is able to perform a quick test run, in which the automatic ranges selected are recorded. During the main run, these ranges are manually selected prior to taking each measurement. In this respect, the Yokogawa WT310e is well supported and is the reason why it was selected for the Intertek tests.

If there are issues during testing, errors are logged. The SERT team are able to provide support, but to make this most effective the user should collect error logs by running a bat file. As the Intertek tests involved a Windows based controller, in this case a « collectlogfiles.bat /r » was run. This produced a zipped file within the SERTlog folder which can be sent to SERT using the form on

http://www.spec.org/sert/feedback/issuereport.html, or by email.

3. Interpretation of results

Refer to http://www.spec.org/sert/docs/SERT-Result_File_Fields.html and to the White Paper on "Investigation of potential approaches to energy efficiency metrics for enterprise servers, based upon the SERT™ rating tool" for further detail on how SERT results can be interpreted.

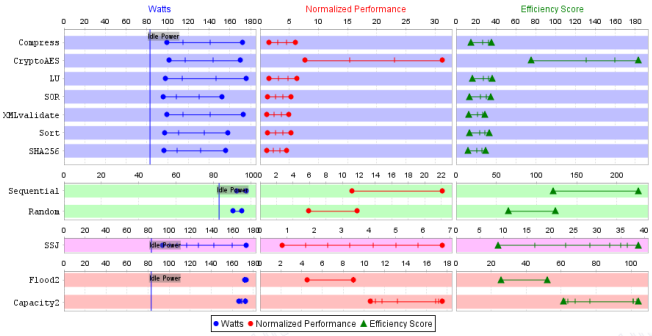
4. Summary results

4.1. Run 1- Full Test

Summary

Workload Efficiency Score				Idle Watts
CPU	Storage	Hybrid	Memory	
33.9	129.4	29.4	56.8	83.0

Worklet Summary



Workload	Worklet	Normalized Peak Performance	Watts at Lowest Load Level	Watts at Highest Load Level	∑ Normalized Performance	∑ Power (Watts)	Efficiency Score
CPU	Compress	6.137	99.9	172.2	15.336	528.1	29.039
	CryptoAES	31.306	101.3	170.4	77.535	533.5	145.341
	LU	6.421	98.5	176.1	16.040	536.0	29.927
	SOR	5.378	96.2	152.9	13.402	488.4	27.439
	XMLvalidate	5.022	99.4	173.3	12.554	529.3	23.719
	Sort	5.350	97.2	158.2	13.329	502.5	26.524
	SHA256	4.629	96.6	156.0	11.575	494.7	23.397
Storage	Sequential	22.177	92.4	97.7	33.342	190.1	175.393
	Random	11.826	90.8	95.4	17.763	186.2	95.406
Hybrid	SSJ	6.719	93.8	173.3	30.064	1,022.0	29.416
Memory	Flood2	9.025	172.3	173.6	13.544	345.9	39.150
	Capacity2	16.896	166.7	173.1	126.556	1,537.5	82.311
Idle	Idle	n/a	83.0	83.0	n/a	83.0	n/a

Aggregate SUT Data

# of Nodes	1	# of Processors	1
Total Physical Memory	64.0 GB	# of Cores	4
# of Storage Devices	4	# of Threads	8

System Under Test

Hardware per Node (1 Node)			
Hardware Vendor	Supermicro	Power Supply Quantity (active / populated / bays)	1 / 1 / 1
Model	X9SRE/X9SRE-3F/X9SRi/X9SRi-3F	Power Supply Details	1 x 680W, PWS-601-1H
Form Factor	1U	Power Supply Operating Mode	Standard
CPU Name	Intel(R) Xeon(R) CPU E5-1620 v2 @ 3.70GHz	Available Power Supply Modes	Standard
CPU Frequency	3701 MHz (up to 3500 MHz), SuperFast mode enabled	Disk Drive Bays (populated / available)	4 / 4
Number of CPU Sockets (populated / available)	1 / 3	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 100.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
CPU(s) Enabled	4 cores, 1 processors, 4 cores/processor	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 179.0 GB SATA Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup, Disk cache disabled in HarderSetup
Number of NUMA Nodes	3	Disk Drive	2 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 279.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
Hardware Threads	8 (2/core), Superthreading enabled	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Primary Cache	64 KB I + 64 KB D on chip per core	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 0 connected, 1 enabled in OS, 2 enabled in firmware 1000 Mbit/s
Secondary Cache	1 MB I+D on chip per chip	Management Controller or Service Processor	Yes
Tertiary Cache	10 MB I+D off chip per chip	Expansion Slots (populated / available)	1 / 3 PCI
Additional Cache	None	Optical Drives	No
Additional CPU Characteristics	None	Keyboard	Enhanced (101- or 102-key)
Total Memory Available to OS	64.0 GB	Mouse	USB Input Device
Total Memory Amount (populated / maximum)	64.0 GB / 5.0 GB	Monitor	Yes
Total Memory Slots (populated / available)	4 / 8	Additional Hardware	1 x ReallyFast Java Accelerator Card
Memory DIMMs	4 x 1GB 2Rx4 PC2-5300F ECC CL5; slots 1, 3, 6, and 8 populated		
Memory Operating Mode	Mirrored		

Software per Node (1 Node)			
Power Management	Enabled (see SUT Notes)	Boot Firmware Version	3.2a
Operating System (OS)	Microsoft Corporation Microsoft Windows Server 2012 R2 Datacenter Evaluation	Management Firmware Version	2.3.4.5
OS Version	6.3.9600	JVM Vendor	Oracle Corporation
Filesystem	NTFS	JVM Version	1.8.0_77-b03
Additional Software	None	Client Configuration ID	Intel_Win_HS17_1n

SUT Notes

- LOCK Pages in Memory enabled

Aggregate Electrical and Environmental Data

Line Standard	230V / 50 Hz / 1 phase / 2 wires
Elevation (m)	82
Minimum Temperature (°C)	24.7

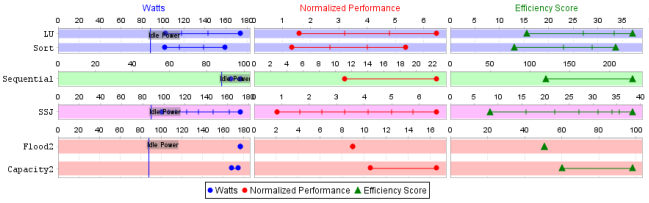
Average of recorded temperatures during test = 25.3°C

4.2. Run 2 - Consistency check part-run

Summary

Workload Efficiency Score				Idle Watts
CPU	Storage	Hybrid	Memory	
27.8	176.3	28.4	63.4	88.0

Worklet Summary



Workload	Worklet	Normalized Peak Performance	Watts at Lowest Load Level	Watts at Highest Load Level	∑ Normalized Performance	∑ Power (Watts)	Efficiency Score
CPU	LU	6.459	102.7	174.1	16.052	538.2	29.828
	Sort	5.376	102.0	159.5	13.407	516.3	25.965
Storage	Sequential	22.535	93.1	98.2	33.737	191.3	176.313
Hybrid	SSJ	6.629	98.2	172.6	29.808	1,050.4	28.378
Memory	Flood2	8.999	177.0	177.0	8.999	177.0	50.848
	Capacity2	16.580	168.5	175.1	27.179	343.6	79.104
Idle	Idle	n/a	88.0	88.0	n/a	88.0	n/a

Aggregate SUT Data

# of Nodes	1	# of Processors	1
Total Physical Memory	64.0 GB	# of Cores	4
# of Storage Devices	4	# of Threads	8

System Under Test

Hardware per Node (1 Node)

Hardware Vendor	Supermicro	Power Supply Quantity (active / populated / bays)	1 / 1 / 1
Model	X9SRE/X9SRE-3F/X9SRi/X9SRi-3F	Power Supply Details	1 x 680W, PWS-601-1H
Form Factor	1U	Power Supply Operating Mode	Standard
CPU Name	Intel(R) Xeon(R) CPU E5-1620 v2 @ 3.70GHz	Available Power Supply Modes	Standard
CPU Frequency	3701 MHz (up to 3500 MHz), SuperFast mode enabled	Disk Drive Bays (populated / available)	4 / 4
Number of CPU Sockets (populated / available)	1 / 3	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 100.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
CPU(s) Enabled	4 cores, 1 processors, 4 cores/processor	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 179.0 GB SATA Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup, Disk cache disabled in HarderSetup
Number of NUMA Nodes	3	Disk Drive	2 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 279.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
Hardware Threads	8 (2/core), Superthreading enabled	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Primary Cache	64 KB I + 64 KB D on chip per core	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 0 connected, 1 enabled in OS, 2 enabled in firmware 1000 Mbit/s
Secondary Cache	1 MB I+D on chip per chip	Management Controller or Service Processor	Yes
Tertiary Cache	10 MB I+D off chip per chip	Expansion Slots (populated / available)	1 / 3 PCI
Additional Cache	None	Optical Drives	No
Additional CPU Characteristics	None	Keyboard	Enhanced (101- or 102-key)
Total Memory Available to OS	64.0 GB	Mouse	USB Input Device
Total Memory Amount (populated / maximum)	64.0 GB / 5.0 GB	Monitor	Yes
Total Memory Slots (populated / available)	4 / 8	Additional Hardware	1 x ReallyFast Java Accelerator Card
Memory DIMMs	4 x 1GB 2Rx4 PC2-5300F ECC CL5; slots 1, 3, 6, and 8 populated		
Memory Operating Mode	Mirrored		

Software per Node (1 Node)

Power Management	Enabled (see SUT Notes)	Boot Firmware Version	3.2a
Operating System (OS)	Microsoft Corporation Microsoft Windows Server 2012 R2 Datacenter Evaluation	Management Firmware Version	2.3.4.5
OS Version	6.3.9600	JVM Vendor	Oracle Corporation
Filesystem	NTFS	JVM Version	1.8.0_77-b03
Additional Software	None	Client Configuration ID	Intel_Win_HS17_1n

SUT Notes

- LOCK Pages in Memory enabled

Aggregate Electrical and Environmental Data

Line Standard	230V / 50 Hz / 1 phase / 2 wires
Elevation (m)	82
Minimum Temperature (°C)	21.2

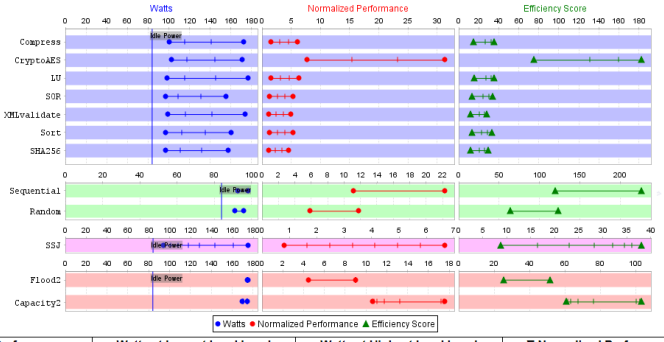
Average of recorded temperatures during test = 22.4°C

4.3. Run 3 – Full Test

Summary

Workload Efficiency Score				Idle Watts
CPU	Storage	Hybrid	Memory	
33.6	128.5	29.0	56.3	83.9

Worklet Summary



Workload	Worklet	Normalized Peak Performance	Watts at Lowest Load Level	Watts at Highest Load Level	∑ Normalized Performance	∑ Power (Watts)	Efficiency Score
CPU	Compress	6.176	100.8	172.7	15.413	530.9	29.033
	CryptoAES	31.353	102.8	171.4	77.775	536.8	144.893
	LU	6.331	99.0	177.0	15.907	537.4	29.602
	SOR	5.381	97.0	155.8	13.403	493.0	27.188
	XMLvalidate	4.963	99.5	173.9	12.403	532.5	23.291
	Sort	5.327	97.5	160.6	13.287	506.6	26.225
	SHA256	4.655	97.0	157.7	11.595	498.0	23.284
Storage	Sequential	22.308	92.9	98.4	33.438	191.3	174.837
	Random	11.845	91.4	96.1	17.724	187.5	94.515
Hybrid	SSJ	6.679	94.2	175.2	29.880	1,029.0	29.037
Memory	Flood2	9.019	174.5	175.6	13.540	350.0	38.680
	Capacity2	17.446	170.1	174.8	127.462	1,555.6	81.937
Idle	Idle	n/a	83.9	83.9	n/a	83.9	n/a

Aggregate SUT Data

# of Nodes	1	# of Processors	1
Total Physical Memory	64.0 GB	# of Cores	4
# of Storage Devices	4	# of Threads	8

System Under Test

Hardware per Node (1 Node)			
Hardware Vendor	Supermicro	Power Supply Quantity (active / populated / bays)	1 / 1 / 1
Model	X9SRE/X9SRE-3F/X9SRi/X9SRi-3F	Power Supply Details	1 x 680W, PWS-601-1H
Form Factor	1U	Power Supply Operating Mode	Standard
CPU Name	Intel(R) Xeon(R) CPU E5-1620 v2 @ 3.70GHz	Available Power Supply Modes	Standard
CPU Frequency	3701 MHz (up to 3500 MHz), SuperFast mode enabled	Disk Drive Bays (populated / available)	4 / 4
Number of CPU Sockets (populated / available)	1 / 3	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 100.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
CPU(s) Enabled	4 cores, 1 processors, 4 cores/processor	Disk Drive	1 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 179.0 GB SATA Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup, Disk cache disabled in HarderSetup
Number of NUMA Nodes	3	Disk Drive	2 x (Standard disk drives) LSI NMR8100-4i SCSI Disk Device 279.0 GB SAS Fixed hard disk media LSI NMR8100-4i SCSI Disk Device Controller Controller cache disabled in EasySetup
Hardware Threads	8 (2/core), Superthreading enabled	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Primary Cache	64 KB I + 64 KB D on chip per core	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 0 connected, 1 enabled in OS, 2 enabled in firmware 1000 Mbit/s
Secondary Cache	1 MB I+D on chip per chip	Management Controller or Service Processor	Yes
Tertiary Cache	10 MB I+D off chip per chip	Expansion Slots (populated / available)	1 / 3 PCI
Additional Cache	None	Optical Drives	No
Additional CPU Characteristics	None	Keyboard	Enhanced (101- or 102-key)
Total Memory Available to OS	64.0 GB	Mouse	USB Input Device
Total Memory Amount (populated / maximum)	64.0 GB / 5.0 GB	Monitor	Yes
Total Memory Slots (populated / available)	4 / 8	Additional Hardware	1 x ReallyFast Java Accelerator Card
Memory DIMMs	4 x 1GB 2Rx4 PC2-5300F ECC CL5; slots 1, 3, 6, and 8 populated		
Memory Operating Mode	Mirrored		

Software per Node (1 Node)			
Power Management	Enabled (see SUT Notes)	Boot Firmware Version	3.2a
Operating System (OS)	Microsoft Corporation Microsoft Windows Server 2012 R2 Datacenter Evaluation	Management Firmware Version	2.3.4.5
OS Version	6.3.9600	JVM Vendor	Oracle Corporation
Filesystem	NTFS	JVM Version	1.8.0_77-b03
Additional Software	None	Client Configuration ID	Intel_Win_HS17_1n

SUT Notes

- LOCK Pages in Memory enabled

Aggregate Electrical and Environmental Data

Line Standard	230V / 50 Hz / 1 phase / 2 wires
Elevation (m)	82
Minimum Temperature (°C)	24.7

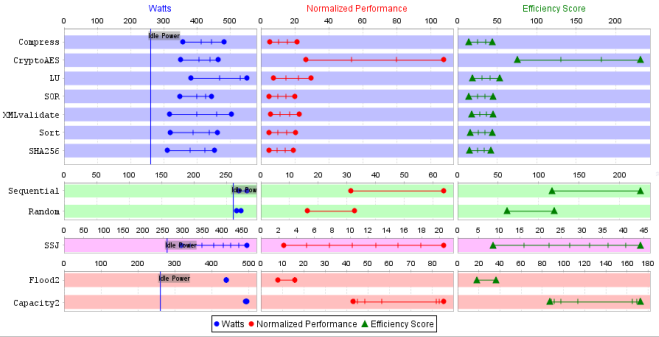
Average of recorded temperatures during test = 29.4°C

4.4. Run 4 – Full Test

Summary

Workload Efficiency Score				Idle Watts
CPU	Storage	Hybrid	Memory	
41.1	125.3	30.1	59.4	260.8

Worklet Summary



Workload	Worklet	Normalized Peak Performance	Watts at Lowest Load Level	Watts at Highest Load Level	∑ Normalized Performance	∑ Power (Watts)	Efficiency Score
CPU	Compress	21.269	359.5	482.0	53.092	1,701.8	31.198
	CryptoAES	107.723	352.0	464.4	268.024	1,663.8	161.092
	LU	29.812	383.7	551.8	74.255	1,936.0	38.355
	SOR	20.244	349.0	445.3	50.526	1,621.3	31.163
	XMLvalidate	22.687	317.8	505.3	56.643	1,683.5	33.646
	Sort	20.568	320.4	463.6	51.354	1,612.0	31.858
	SHA256	19.234	312.4	453.5	48.036	1,569.8	30.600
Storage	Sequential	63.834	269.3	282.3	95.201	551.6	172.581
	Random	32.746	266.5	273.6	49.114	540.1	90.930
Hybrid	SSJ	20.511	297.2	463.7	92.606	3,080.6	30.061
Memory	Flood2	16.058	440.8	440.8	24.118	881.5	27.359
	Capacity2	83.279	491.3	496.1	573.286	4,448.3	128.879
Idle	Idle	n/a	260.8	260.8	n/a	260.8	n/a

Aggregate SUT Data

# of Nodes	1	# of Processors	2
Total Physical Memory	63.9 GB	# of Cores	4
# of Storage Devices	11	# of Threads	8

System Under Test

Hardware per Node (1 Node)

Hardware Vendor	Supermicro	Power Supply Quantity (active / populated / bays)	1 / 2 / 1
Model	X10DRi	Power Supply Details	1 x 1400W, PWS-10400
Form Factor	1U	Power Supply Operating Mode	Standard
CPU Name	Intel(R) Xeon(R) CPU E5-2667 v3 @ 3.20GHz	Available Power Supply Modes	Standard
CPU Frequency	3201 MHz (up to 3500 MHz), SuperFast mode enabled	Disk Drive Bays (populated / available)	11 / 4
Number of CPU Sockets (populated / available)	2 / 3	Disk Drive	1 x (Standard disk drives) Kingston SNA-DC/U USB Device 112.0 GB SAS External hard disk media Kingston SNA-DC/U USB Device Controller Controller cache disabled in EasySetup
CPU(s) Enabled	4 cores, 2 processors, 4 cores/processor	Disk Drive	10 x (Standard disk drives) LSI MR9271-4i SCSI Disk Device 3,725.0 GB SATA Fixed hard disk media LSI MR9271-4i SCSI Disk Device Controller Controller cache disabled in EasySetup, Disk cache disabled in HarderSetup
Number of NUMA Nodes	3	Network Interface Cards	2 x Intel(R) 82599 10 Gigabit Dual Port Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Hardware Threads	8 (2/core), Superthreading enabled	Network Interface Cards	1 x Intel(R) Ethernet Connection I217-LM 0 connected, 1 enabled in OS, 2 enabled in firmware 1000 Mbit/s
Primary Cache	64 KB I + 64 KB D on chip per core	Network Interface Cards	2 x Intel(R) Ethernet Server Adapter X520-2 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Secondary Cache	1 MB I+D on chip per chip	Network Interface Cards	1 x Intel(R) I210 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Tertiary Cache	10 MB I+D off chip per chip	Network Interface Cards	6 x Intel(R) I350 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Additional Cache	None	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Additional CPU Characteristics	None	Network Interface Cards	1 x Intel(R) I350 Gigabit Network Connection 1 connected, 1 enabled in OS, 2 enabled in firmware 1000 Mbit/s
Total Memory Available to OS	63.9 GB	Network Interface Cards	2 x Realtek PCIe GBE Family Controller 1 connected, 0 enabled in OS, 2 enabled in firmware 0 Mbit/s
Total Memory Amount (populated / maximum)	63.9 GB / 5.0 GB	Management Controller or Service Processor	Yes
Total Memory Slots (populated / available)	4 / 8	Expansion Slots (populated / available)	1 / 3 PCI
Memory DIMMs	4 x 1GB 2Rx4 PC2-5300F ECC CL5; slots 1, 3, 6, and 8 populated	Optical Drives	No
Memory Operating Mode	Mirrored	Keyboard	Enhanced (101- or 102-key)
		Mouse	HID-compliant mouse
		Monitor	Yes
		Additional Hardware	1 x ReallyFast Java Accelerator Card

Software per Node (1 Node)

Power Management	Enabled (see SUT Notes)	Boot Firmware Version	2.0
Operating System (OS)	Microsoft Corporation Microsoft Windows 10 Pro	Management Firmware Version	2.3.4.5
OS Version	10.0.10586	JVM Vendor	Oracle Corporation
Filesystem	NTFS	JVM Version	1.8.0_73-b02
Additional Software	None	Client Configuration ID	Intel_Win_HS18_1

SUT Notes

- LOCK Pages in Memory enabled

Aggregate Electrical and Environmental Data

Line Standard	230V / 50 Hz / 1 phase / 2 wires
Elevation (m)	82
Minimum Temperature (°C)	27.4

Average of recorded temperatures during test = 28.9°C

Appendix 4 – Metrics Paper

Investigation of potential approaches to energy efficiency metrics for enterprise servers, based upon the SERT™ rating tool

TABLE OF CONTENTS

APPENDIX 4 – METRICS PAPER

1. WHY SERT? 6
2. INTRODUCTION TO THE SERT TOOL 7
   - *2.1. Ownership* 7
   - *2.2. Design and Operation* 7
   - *2.3. Test conditions* 8
   - *2.4. Results* 9
     - *2.4.1. Results by worklet:* 10
   - *2.5. Revisions and Licensing* 11
     - *2.5.1. License fee* 11
     - *2.5.2. Updates and Versions* 11
     - *2.5.3. Triggering of updates / new versions* 12
   - *2.6. SERT test execution* 12
   - *2.7. SERT scalability* 12
   - *2.8. Referencing SERT in standards* 12
   - *2.9. Applying SERT in policy* 13
     - *2.9.1. Data gathering* 13
     - *2.9.2. Power modifiers* 13
     - *2.9.3. Reporting* 13
3. STUDY METHODOLOGY FOR METRIC DEVELOPMENT 14
4. ASSESSMENT OF EXISTING METRIC APPROACHES 15
   - *4.1. Existing activities on server metrics by policy makers* 15
   - *4.2. Existing activities on server metrics by industry* 15
5. ANALYSIS OF DATA TO PROVIDE INSIGHTS ON SERT TOOL OPERATION 18
   - *5.1. Evidence base* 18
   - *5.2. Preliminary data analysis observations* 18
   - *5.3. Preliminary insights on combining workloads* 21
6. PRINCIPLES FOR METRIC DEVELOPMENT 22
7. FINAL METRIC 27
   - *7.1. Average Server Efficiency Metric* 27
   - *7.1.1. Stage 1 : Calculating the efficiency for each worklet* 27
   - *7.1.2. Stage 2: Calculating the efficiency for each workload* 27
   - *7.1.3. Stage 3: Combining workload components into an average server efficiency metric* 28
   - *7.1.4. Supplemental indicators for average server power consumption and performance* 29
     - *7.1.4.1. Stage 1 : Data collection at the worklet level* 29
     - *7.1.4.2. Stage 2: Calculating individual workload power consumption and performance* 29
     - *7.1.4.3. Combining workload components into the server performance and power consumption indicators* 30
8. SERVER EFFICIENCY METRIC EVALUATION AGAINST DEVELOPMENT PRINCIPLES 31
   - *8.1. Applying server efficiency metric to real server data* 31
   - *8.2. Evaluating metric in relation to requirements* 33
9. GUIDANCE ON PRODUCTS TO TEST 34

APPENDIX 4 ANNEX 1: TESTING METRIC FOR ACHIEVEMENT OF APPROPRIATE ENERGY EFFICIENCY RANKING OF PRODUCTS 37

Investigating dynamic range (idle/max ratio) between different performance servers ... 39

APPENDIX 4 ANNEX 2 : DEPLOYED POWER APPROACH AND TESTING METRIC FOR SYSTEM SCALABILITY UNDER DIFFERENT UTILISATIONS AND WORKLOADS USING REAL DATA. 42

Performance of metric at different utilisations levels 42

Performance of metric under different workloads 44

List of Figures

Appendix 4 Figure 1- Elements of the SERT tool 7

Appendix 4 Figure 2- Summary of Workloads and Worklets used in SERT (Source: SPEC) 8

Appendix 4 Figure 3- Schematic for SERT controller and system under test (SPEC) 9

Appendix 4 Figure 4 – Testing intervals typical to worklets 10

Appendix 4 Figure 5 - Sample SERT results (SPEC) 11

Appendix 4 Figure 6 - Metric development methodology 14

Appendix 4 Figure 7 - Specifications of servers in preliminary analysis 18

Appendix 4 Figure 8 - Preliminary data analysis observations 20

Appendix 4 Figure 9 – Stages in determining average server efficiency from SERT results 27

Appendix 4 Figure 10 - Power, performance and efficiency relationship 29

Appendix 4 Figure 11 - Efficiency / Performance comparison for analysis data set 32

Appendix 4 Figure 12 - Hypothetical efficiency curves modelled 38

Appendix 4 Figure 13 - Power-performance curve of servers with equal efficiency under metric with dynamic range weighting of +1 39

Appendix 4 Figure 14- Power-performance curve of servers with equal efficiency under dynamic range weighting of +1.5 40

Appendix 4 Figure 15 - Power-performance curve of servers with equal efficiency with no dynamic range weighting 41

Appendix 4 Figure 16 - Average deployed power plotted against efficiency metric 42

Appendix 4 Figure 17 - Deployed power at 50% fixed utilisation against efficiency metric with dynamic range weighting 1.5, based on The Green Grid data set 43

Appendix 4 Figure 18 - Deployed power at 50% fixed utilisation against efficiency metric without dynamic range weighting, based on The Green Grid data set 44

List of Tables

Appendix 4 Table 1 – Evolution of conceptual industry metrics 16

Appendix 4 Table 2 Metric requirements 22

Appendix 4 Table 3 - Evaluation of proposed metric 33

Appendix 4 Table 4- Hypothetical efficiency curves considered 37

Appendix 4 Table 5 - Relative performance of hypothetical curves 38

Appendix 4 Table 6 - Metric correlation with deployed power using balanced workload under various utilisation levels, with and without dynamic range .. 43

Appendix 4 Table 7 - Metric correlation with deployed power using CPU intensive workload under various utilisation levels, with and without dynamic range .. 45

Appendix 4 Table 8 - Metric correlation with deployed power using memory intensive workload under various utilisation levels, with and without dynamic range 45

1. Why SERT?

Whilst a number of measurement methods are available to measure and report server energy use in an accurate and reproducible manner, many are still not finalised. An analysis of standards for enterprise server and data storage products identified the availability of metrics for active energy efficiency for servers as a key gap. Therefore the goal of this study was to analyse and present the options for metrics for active state energy efficiency of enterprise servers.

The approach that has received widest attention to date by policy and standardisation initiatives is the use of the Standard Performance Evaluation Corporation (SPEC) SERT™[14] tool as a basis for active energy efficiency metrics. Therefore the development of metrics appropriate for policy use based on the SERT tool are the focus of this report.

An ISO standard under development[15] recognises but does not specifically endorse the SERT tool as a possible option to measure efficiency. Current policy interest in SERT includes the United States (US EPA / ENERGY STAR), Korea, China, and Australia / New Zealand. The usual policy approach to metric development based upon the SERT tool involves the policy maker gathering the data, carrying out the analysis and defining their approach and metric, so that where necessary SPEC can customise the tool accordingly.

[14] SPEC, SERT, SPECpower_ssj2008 and the SPEC logo are registered trademarks of the Standard Performance Evaluation Corporation. Copyright © 1988-2016 Standard Performance Evaluation Corporation (SPEC). All rights reserved.

15 ISO/IEC 30134-4 led by ISO/IEC JTC 1/SC39 WG addressing Key Performance Indicators (KPIs) to assess resource and energy efficiency for servers.

2. Introduction to the SERT tool

2.1. Ownership

The SERT tool was created by the Standard Performance Evaluation Corporation (SPEC). SPEC is a non-profit organisation open to all parties but requires membership fees. SPEC has over 50 members which includes almost all the main ICT hardware manufacturers and a number of software and internet companies. There are also SPEC Associates and a Research Group which include approximately 100 other organisations, in particular universities in USA, Japan and Germany.

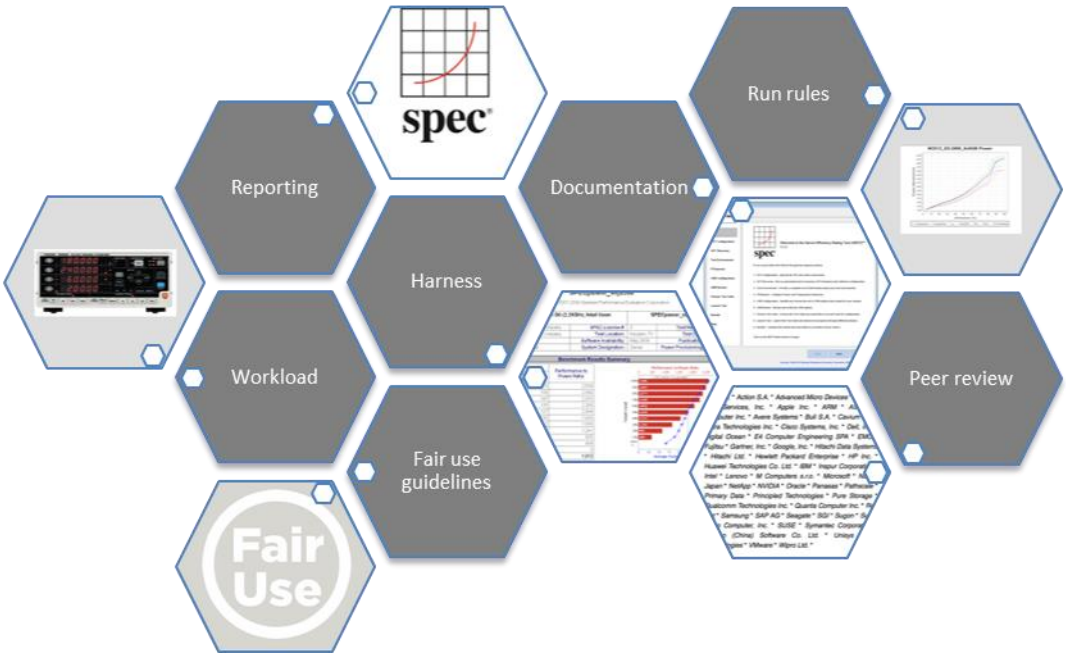
The SPECpower Committee leading development was established in 2006, and includes corporate (AMD, Dell, Fujitsu, HP, Intel, IBM, Microsoft and Oracle) and academic support (University of Wurzburg, Germany).

2.2. Design and Operation

The Server Efficiency Rating Tool (SERT tool) is a software tool for measuring server energy efficiency. Central design considerations underpinning the SERT tool and making it a promising candidate for use in policy measures include; reproducibility of results, fairness, verifiability and usability.

The SERT tool is intended to be economical and easy to use with the minimum equipment and skill requirements. It has a graphical user interface for easy configuration, and after setup the process is automated to minimise the time necessary for testing. SERT is hardware and OS agnostic, meaning that it supports various hardware platforms and operating systems. and has the ability to run on a wide range of server specifications and configurations even as these continue to expand. Servers should be tested in their "as shipped" or "out of the box" state, and target run time is around five hours, although this will vary with server generations.

The tool includes a number of elements shown in Appendix 4 Figure 1.



Appendix 4 Figure 1- Elements of the SERT tool

The SERT tool simulates a variety of common types of work via worklets. These are essentially software simulations of real working environments using "transactions" tailored to test discrete system components (e.g. processors, memory and storage) and subsystems (e.g. RAM and CPU). A range of worklets (with varying transaction types) is necessary to ensure platform neutrality, as performance of different server architectures will vary with different workloads. For ease of comparability, the worklet results are normalised against results for an arbitrarily selected baseline server model.

Worklets provide representative results but do not represent a particular application. They scale automatically with the available hardware, so that more hardware resource (e.g. increased processor/memory/disk capacity or additional processor/memory/disk module) results in an increase in the performance score. As worklets are transaction based, performance is indicated by the throughput in terms of the number of transactions completed per second.

The worklets can be grouped by the subsystems (workloads) they address, and can be adjusted to different loading levels as shown in Appendix 4 Figure 2.

Workload	Load Level	Worklet Name
CPU	100%, 75%, 50%, 25%	Compress
		CryptoAES
		LU
		SHA256
		SOR
		SORT
		XMLValidate
Memory	Flood: Full, Half	Flood
	Capacity: 4GB, 8GB, 16GB, 128GB, 256GB, 512GB, 1024GB	Capacity
Storage	100%, 50%	Random
		Sequential
Hybrid	100%, 87.5%, 75%, 62.5%, 50%, 37.5%,25%, 12.5%	SSJ
Idle	idle	Idle

Appendix 4 Figure 2- Summary of Workloads and Worklets used in SERT (Source: SPEC)

Network input / output worklets are not specifically included, due to i) the difficulty in validating external equipment which must be connected at the other end of the network ii) the low power consumption of the networking components relative to the whole server iii) results suggesting there are no significant differences in power utilization between 100% and 0% network utilization for current technology. If evidence available in future showed different behaviour justifying special consideration of network I/O, this could be handled via configuration of power/performance modifiers, i.e by applying an energy allowance to idle/active power for additional or faster network interfaces. In future versions, a worklet for very high power conditions (such as linpack) may be considered.

Likewise, worklets are not designed to specifically evaluate general purpose graphics processing units (GPGPUs) and other types of sub-processors.

2.3. Test conditions

It is important that the physical test environment is representative of typical user environments. The temperature range has an impact on the design and energy consumption of the server, particularly the internal cooling system such as the heatsink and fans. In particular, unusually low temperatures (below 20°C for a data centre) may result in artificially lower power demand (reduced fan operation) and improved performance during the tests.

The SERT tool must therefore be run within constrained environmental conditions, specified as follows:

Ambient temperature lower limit: 20°C[16]

[16] Comparing SERT operating ranges with ASHRAE 2011 Thermal Guideline Classes for data centres, A1 is the range 15 to 32°C, A2 is in the range 10 to 35 and A3 is in the range 5 to 40. Most servers are designed to operate within A2 conditions. The recommended operating conditions for all these is 18-27°C.

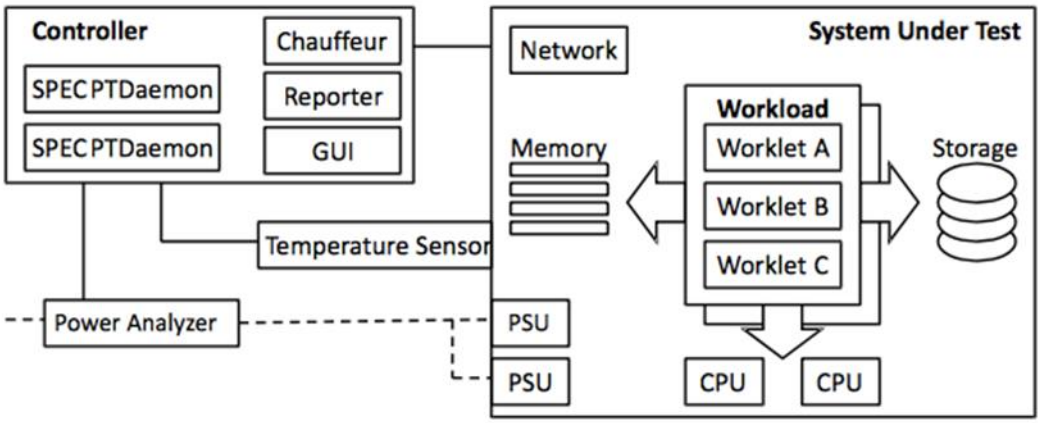
Ambient temperature upper limit: within documented operating specification of the SUT (but it is likely that servers will be tested as close to the lower limit as possible as this is where they perform most efficiently).

Elevation and Humidity: within documented operating specification of the SUT

No overt direction of air flow in the vicinity of the measured equipment in a way that would be inconsistent with normal data centre practices.

AC power supply (single or 3-phase). SERT is not compatible with servers using low voltage and 48V DC power supply.

Compliance with these conditions is validated as shown in Appendix 4 Figure 3 by the use of a temperature sensor[17] in the testing rig.



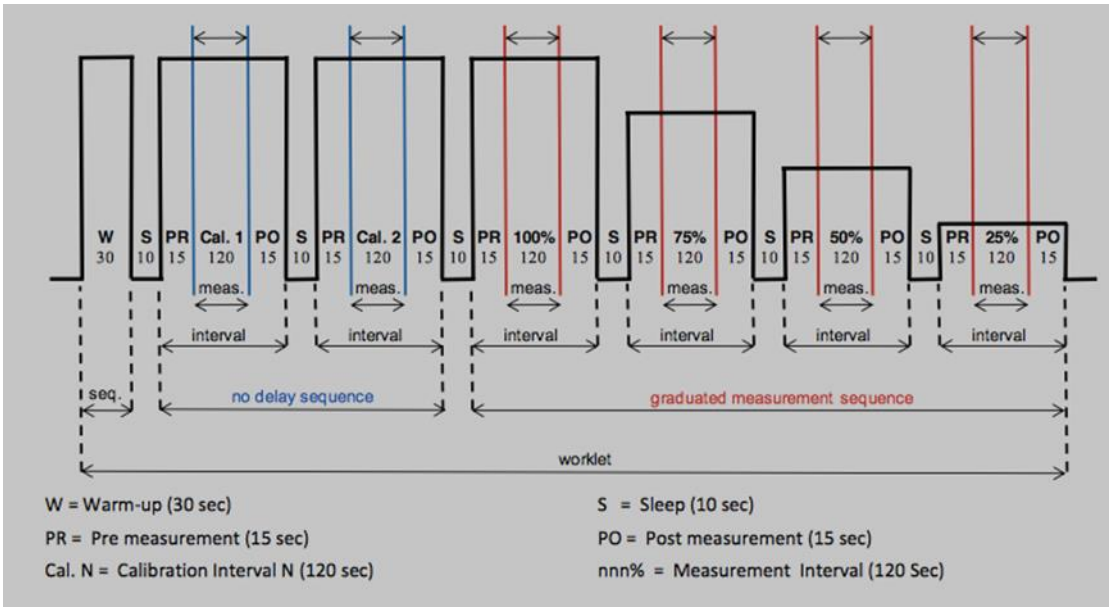
Appendix 4 Figure 3- Schematic for SERT controller and system under test (SPEC)

2.4. Results

Results are provided in both machine (XML) and human readable (HTML / TXT) forms, accompanied by summary and detail reports. Customised result reports can be configured if required. The tool includes some features to avoid favourable "gaming" of results - for example, divergence from standard settings (tuning parameters) is possible, but renders the output test results invalid.

Each worklet contains a number of testing intervals as shown in Appendix 4 Figure 4.

[17] Temperature must be measured no more than 50mm in front of (upwind of) the main airflow inlet of the System Under Test (SUT). The sensor must have an interface that allows its measurements to be read by the SERT harness. The reading rate supported by the sensor must be at least four samples per minute. Measurements must be reported by the sensor with an overall accuracy of +/- 0.5 degrees Celsius or better for the ranges measured during the SERT run.



Appendix 4 Figure 4 – Testing intervals typical to worklets

Worklets are designed to self-calibrate to the maximum loading level at the start of the test. The maximum loading level represents the maximum performance (throughput) the server under test is capable of achieving. The two calibration intervals are shown in blue in Appendix 4 Figure 4. After determining the 100% loading level, formal testing intervals can then be measured at the different required loading levels – in this case 100%, 75%, 50% and 25%.

2.4.1. Results by worklet:

Measurements and calculations for power:

For each interval / loading level the instantaneous power in Watts is measured then averaged.

This average power for each interval is summed to arrive at a [**sum of Power**] result for each worklet.

Measurements and calculations for performance:

For each interval the performance score (throughput) is measured in transactions/second. While the performance is predetermined by the loading level there is a very small deviation.

For each interval the performance score is normalised against (divided by) results from a baseline/reference machine for ease of comparability.

The normalised performance scores for each interval are summed to arrive at a total performance score [**sum of Normalised Performance**].

Calculating worklet efficiency score:

The efficiency score for each worklet is defined as:
1000 * [**sum of Normalised Performance**] / [**sum of Power (Watts)**] this is measured in transactions/Joule.

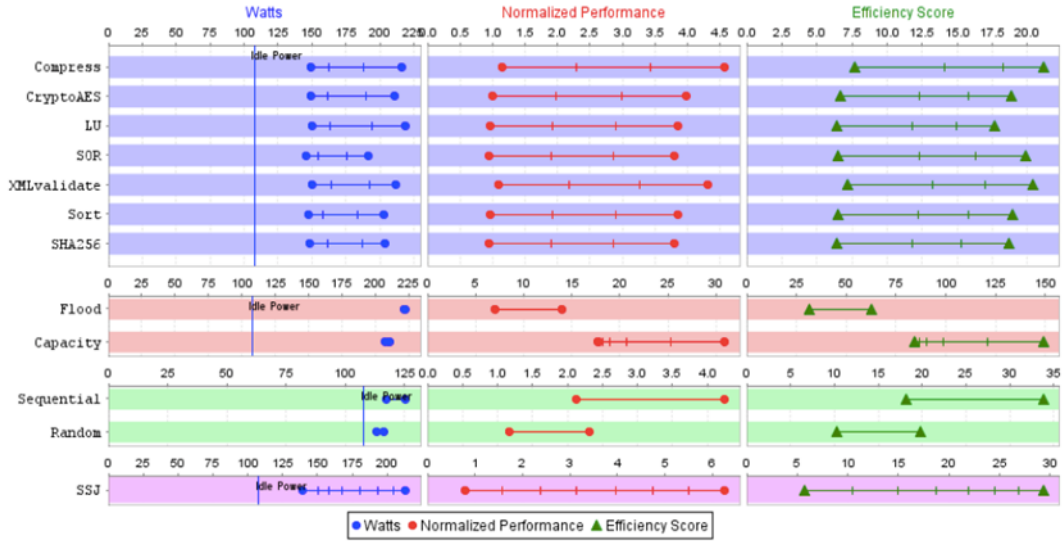
Efficiency for the Idle worklet is marked as not applicable (n/a) because the performance part is zero by definition. As such, idle power is not included in the per worklet efficiency score calculation.

The higher the worklet efficiency score, the higher the energy efficiency.

This approach to calculating worklet results remains the same in the majority of metric approaches that have been explored to date (and if alternative approaches are used they are mathematically equivalent). Note: SPEC discourages manufacturers from quoting numerical values for specific worklets for marketing purposes as taken in isolation these values can be misrepresentative.

The worklet efficiency scores can then be aggregated into a workload score (see Appendix 4 Figure 2 for worklet groupings into workloads). There are different approaches to doing this, which can be specified in each metric approach. Subsequently, the workload scores can then be combined together to arrive at a single number metric result that can enable an overall pass/fail conclusion in relation to the requirements of a particular policy.

Sample results are shown in Appendix 4 Figure 5. The ranges in the values represent the different loading levels tested – see Appendix 4 Figure 2 for relevant loading levels for each worklet.



Appendix 4 Figure 5 - Sample SERT results (SPEC)

2.5. Revisions and Licensing

2.5.1. License fee

The SERT software license must be purchased from SPEC. The license fee is $2,800[18], with a reduced fee for not-for-profit organisations of $900. No other reductions are available, and the license fee is not related to the size of the organisation ordering it.

2.5.2. Updates and Versions

At the time of writing of this paper, the SERT™ Tool is currently on version 1.1.1 (since Jan 29th, 2016).

Small changes such as necessary enhancements for usability / performance or new hardware capabilities that do not fundamentally change the results are included in "updates" (i.e. from 1.1.1 to 1.1.2). These are covered in the initial purchase price of a software license. There have been 4 updates in the last 2 years. Regardless of updates, test results from a specific version should still be broadly comparable, although there may be variations in scores between versions. For example, in the transition from SERT V1.1.0 or SERT V1.1.1, changes were made to rebalance the worklets and increase memory scores. Standard practice is that as soon as a new update is issued, the older software can no longer be used, but policy makers may prefer to request extended availability of previous updates in order to ensure minimum disruption to their initiatives.

A change in "version" (i.e. version 1.1.1 to version 2.0.0) would require a further license purchase. New SERT versions have a development cycle of around 4 years, and can remain in place for between 2 and 8 years. The next version could potentially be introduced between around 2017 and 2020. SPEC aims to halt support to previous versions of the tool, as soon as possible, but would work with policy makers to ensure this fitted with their timelines. SERT would not expect test results from different versions to be comparable.

The EPA ENERGY STAR Version 2.1 specification references the "most current" SERT version, cross referencing to the most recently published memo located on the Enterprise Servers Specification Version 2.1 website[19]. In this way, the necessary edits to the ENERGY STAR server specification and test method are minimised. The EPA reviews each revision of SERT prior to requiring it for

[18] http://www.spec.org/order.html

[19] https://www.energystar.gov/products/spec/enterprise_servers_specification_version_2_1_pd

ENERGY STAR testing purposes. If a SERT update was not considered performance/energy neutral (i.e. to give comparable results with previous versions), it would not be approved for ENERGY STAR testing, but as the EPA works directly with SPEC on updates, no such issues have arisen to date.

2.5.3. Triggering of updates / new versions

New technologies are generally isolated to a small section of the market and are adopted for the additional performance improvement they provide either in terms of cost or power. When new technology appears on the market, it does not represent a typical configuration and is therefore unlikely to be fully covered in the current metric. However, as new technology becomes more common it may merit consideration.

No formal mechanism to handle new technologies in metrics and related tools has been identified. However, it would be possible to introduce one – for example by (either the programme institution or SPEC) monitoring the sales of common configurations in the market, and when new technology reaches at a particular sales or energy consumption threshold, the process to expand metrics to cover the new technology could be initiated

2.6. SERT test execution

SERT tests can be carried out directly by manufacturers or on manufacturer's behalves by an external testing laboratory.

In the case of the US ENERGY STAR label, prior to associating the label with any server product, it is necessary to obtain written certification of ENERGY STAR qualification from an EPA recognised Certification Body based on testing in an EPA recognised testing laboratory[20].

For EU ENERGY STAR registered products and for ecodesign conformity purposes, testing by certified bodies is not necessary and can be carried out directly by manufacturers, so the coverage of certified laboratories in Europe is lower. Further details on the practicality of SERT testing, including repeatability and measurement uncertainty can be found in Appendix 3.

2.7. SERT scalability

Scalability relates to the capability of a system, network, or process to handle a growing amount of work, or its potential to be enlarged in order to accommodate that growth. As more components (processors, memory, and disk storage) are added to the server, the power demand of the server will increase, but the performance reflected by the metric tool should also increase compared to a previous configuration or to other products with a lower configuration. Whilst performance can be scaled relatively easily based on the number of CPUs and their speed, for other components the complexity of scaling performance is substantially increased.

One of the main design goals of the SERT tool is that the system performance should scale in proportion to the system configuration. However, if additional resources are added to a server that is not able to use them effectively, there may be performance bottlenecks in other components, which could result in higher power consumption without a corresponding increase in performance.

The SERT tool has the following features in relation to scalability:

Socket coverage: Can be used on servers up to 8 sockets, although is currently only formally supported up to 4 sockets.

Worklet scaling: Integral to the design of SERT, worklets are able to scale with relevant capabilities and different server configurations (i.e. in terms of increased memory, different

SERT is unable to account for entirely new architectures. Predicting new architectures is not possible as it is only possible to base the tool on known parameters. However, if these became more widespread, new versions of the tool could be triggered to address these.

2.8. Referencing SERT in standards

Test methods need to be specified (either via transitional methods or via a harmonised standard) in order to demonstrate and check the compliance of products with Ecodesign and Energy Labelling requirements.

[20] A list of EPA-recognized laboratories and certification bodies can be found at https://www.energystar.gov/index.cfm?fuseaction=recognized_bodies_list.show_RCB_search_form

It is not common that formal harmonised standards would reference software such as rating tools, but this has been done in specific cases – for example, the multimedia (MPEG) compression standard includes a detailed technical specification and a reference software implementation, to avoid issues with bugs and to develop appropriate software.

Software specifically referenced in standards to support ecodesign regulations would need to be widely available to all parties who wish to use it. Furthermore, the software must be available over a reasonably long period of time to ensure that tests conducted at one point in time can be fairly compared against previous or future testing. The implementation of information requirements based upon the SERT tool in the US EPA ENERGY STAR demonstrate that it is logistically possible for an energy efficiency programme to reference a software tool.

The most suitable approach appears to be the definition of a software specification for the task, rather than a reference to a specific software version.

2.9. Applying SERT in policy

2.9.1. Data gathering

Based upon the way in which SPEC has interacted with the US EPA, SPEC recommends an initial 9 to 12 month data gathering phase (via a reporting requirement) be carried out by the programme institution. Once a "critical mass" of data has been gathered, metric and threshold development can begin.

2.9.2. Power modifiers

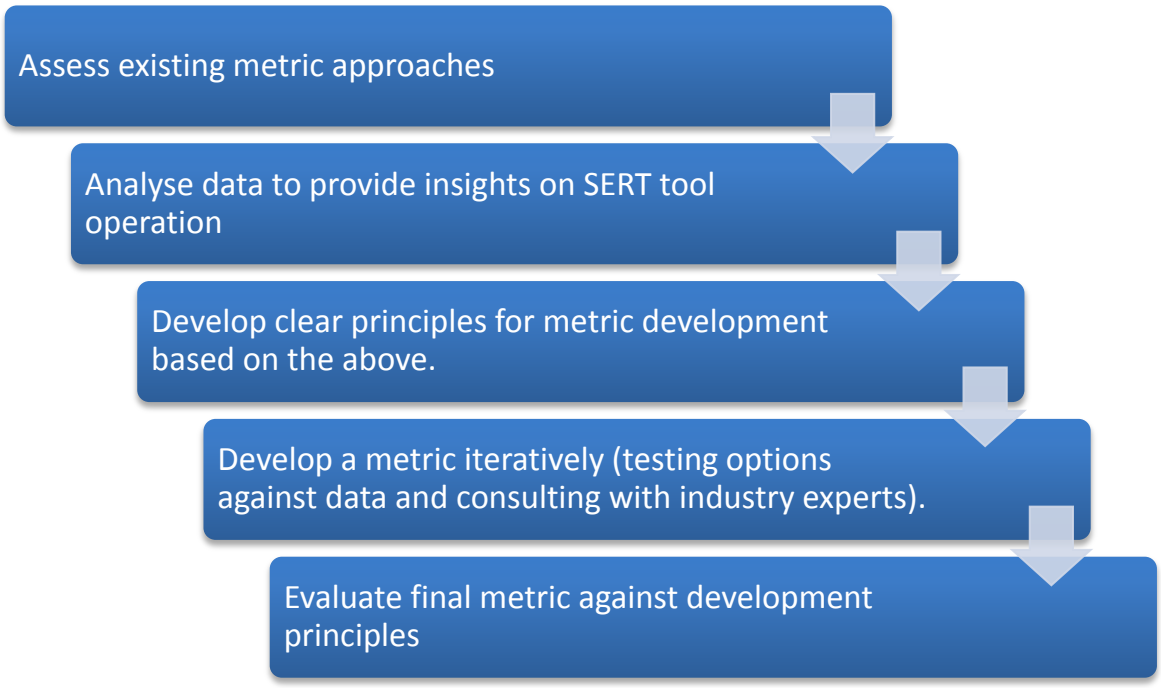
Servers may have optional features designed to increase the breadth of applications that require additional power, as well as redundant capacity in the power supply, cooling system, memory, storage or processing. SERT performance scores avoid over incentivising expandability as there may be efficient servers that do not have such expansion options (otherwise a false incentive could be created encourage unnecessary additional features to be added to servers to allow them to easily qualify). Likewise, the SERT tool does not adapt for redundancy (no measurements are taken under fault tolerant conditions when one of a redundant set of components is disabled). Therefore SPEC suggests that during the metric and threshold development phase, power/performance modifiers are used (sometimes referred to as adders) to account for items the SERT tool cannot measure or for which the performance cannot be determined. However, such allowances would need to be clearly supported by evidence and be proven essential to the metric implementation as they would add additional complexity.

2.9.3. Reporting

The SERT tool enables the user to input predetermined information describing the hardware and software, part of which can be discovered automatically by the tool. This information is included with the performance and power data in the SERT result file, which SPEC would expect to be sent directly to the programme institution. SPEC permit the publishing of results by programme institutions in formats differing from the original SERT file.

3. Study methodology for metric development

In order for the SERT tool to provide a basis to address server energy efficiency in policy initiatives, metrics based upon the output results must be developed. The metric development followed the methodology shown in Appendix 4 Figure 6.



Appendix 4 Figure 6 - Metric development methodology

4. Assessment of existing metric approaches

4.1. Existing activities on server metrics by policy makers

At the time of writing no policies have been identified that define policy requirements based upon metrics using the SERT tool. However, the following relevant international policy initiatives were identified:

Ireland: The Sustainable Energy Authority of Ireland's Triple E programme references the SPEC ssj_2008 test which only tests the CPU and RAM (similar to the hybrid worklet in the SERT tool). Outputs of power demand and performance rating at each load level enable three performance/power ratios to be developed (based on loadings of low (10-30%), mid (40-60%) and high (70-100%)). The specifications define minimum ratios for each utilisation level.

China: The Chinese standardization authority CNIS is currently working to define metrics based upon the SERT tool.

Korea: The Korean standardization body Kemco is currently working to define metrics based upon the SERT tool.

USA: The US Environmental Protection Agency (US EPA) launched the version 3.0 revision of their ENERGY STAR specification for servers in March 2016. The previous specification involved a data gathering exercise, whereby SERT test results had to be submitted for qualifying products. It is the intention that the version 3.0 specification would include a metric and requirements based upon the SERT tool.

As the timelines of these policies indicates, metrics based upon the SERT tool are still in progress, and no exemplary approach has yet been established. This implies that European development of a server metric based upon SERT could become influential in the international policy arena.

4.2. Existing activities on server metrics by industry

Metric analysis by industry to date has involved ITI, The Green Grid (TGG), SPEC and DIGITALEUROPE. Data used as a basis for analysis includes data gathered for the US ENERGY STAR server specification, through additional testing by The Green Grid, and by SPEC.

The TGG SERT Analysis Working Group (which includes Digital Europe and SPEC members) are in the process of evolving their proposals, and the project team has been in correspondence with industry in order to refine and finalise the proposals contained in this report. An analysis of the evolution of the industry proposals, in authors' understanding, is shown in Appendix 4 Table 1, against a selection of key criteria for effective metrics defined by the project team (and discussed in more detail in Appendix 4 section 8).

Appendix 4 Table 1 – Evolution of conceptual industry metrics

Approach	Power scaling / lower utilisation accounted for?	Appropriate weightings and averaging?	Intuitive	Comments
Stage 1: Worklet efficiency (combination of power and performance utilisation levels)				
Worklet efficiency calculated from the sum of the performance divided by the sum of the power at different utilisation levels	More emphasis on higher utilisations	Yes. Worklet performance is normalised	Yes. However, calculating the mean may be more intuitive	Standard SPEC approach
Worklet efficiency calculated from the peak performance divided by the sum of the power at different utilisation levels	More emphasis on higher utilisations	No. Different worklets have different numbers of utilisation levels which affects the result	Not as intuitive as SPEC approach but simplifies equations	This uses the same principle as the SPEC approach but may be less intuitive
Stage 2: Workload efficiency (worklet combination)				
Combination of worklet efficiencies using geometric mean and no weightings	(inherited from worklet efficiency approach)	Geo-mean is less mathematically representative of real life use	Yes	No weightings assumes workload is equally balanced around the worklets. High correlation between worklets means this is not critical
Combination of subset of worklets using geometric mean and no weightings	(inherited from worklet efficiency approach)	Geo-mean is less mathematically representative of real life use	Yes	Subset of worklets suggests other worklets are not needed to be representative of workload
Stage 3: Server efficiency				
Combination of workloads using geometric mean and weightings between the workloads to create single metric	(inherited from worklet efficiency approach)	Geo-mean is less mathematically representative of real life use. Weightings are appropriate	Yes	
Three use categories (compute intensive, memory intensive, storage intensive), with different weightings between the four workloads (CPU, Hybrid, Memory and Storage), focusing on higher utilisation levels.	(inherited from worklet efficiency approach)	Geo-mean is less mathematically representative of real life use	Potential for overlap between three use categories that could cause issues for policy approaches.	

Approach	Power scaling / lower utilisation accounted for?	Appropriate weightings and averaging?	Intuitive	Comments
Single metric based on aggregated performance over worklets divided by aggregated power over worklets and utilisation levels. Performance and figures are averaged using geo-mean and weighted. Power figures are averaged using arithmetic mean.	More emphasis on higher utilisations.	Geo-mean is less mathematically representative of real life use. May over emphasise memory power consumption by using arithmetic mean	Yes but less mathematicall y correct	This was the latest industry approach, but industry was continuing to investigate and develop metrics. It is distinct from the other two approaches but using the arithmetic mean to calculate power may place too much emphasis on the highest power consuming workloads, ie memory.

5. Analysis of data to provide insights on SERT tool operation

5.1. Evidence base

An evidence base was used for initial analysis. This data sets included the following:

- An ENERGY STAR data set with aggregate efficiency scores for each worklet based on the sum of performance results/ sum of power consumption.
- More detailed anonymised SERT test data provided by The Green Grid covering a subset of worklets for 84 managed 2 socket server configurations from various manufacturers.
- A full, detailed data set from SPEC containing the average power and performance data for each worklet at the 4 test points, 0, 25, 75 and 100% (i.e. the figures used to calculate the efficiency score of each worklet). This included:
  - ✓ 162 server product configurations, across a range of models from several generations from 2015 and before, measured with SERT V1.1.1
  - ✓ SSD storage results from older SERT versions (pre 1.0.0 and not in the same format, and without overall results for the other worklets).

  The full data set was the primary source used in the analysis and development of the metric. The ENERGY STAR data set was not used since the aggregate efficiency scores did not provide sufficient information. The subset of SERT worklets provided by The Green Grid was used to verify the final approach by comparison with the primary source.

5.2. Preliminary data analysis observations

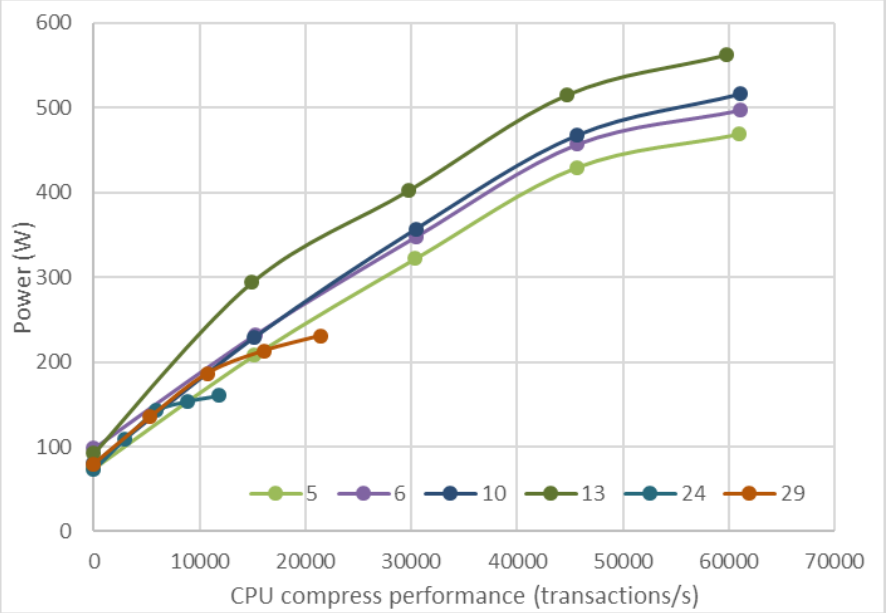
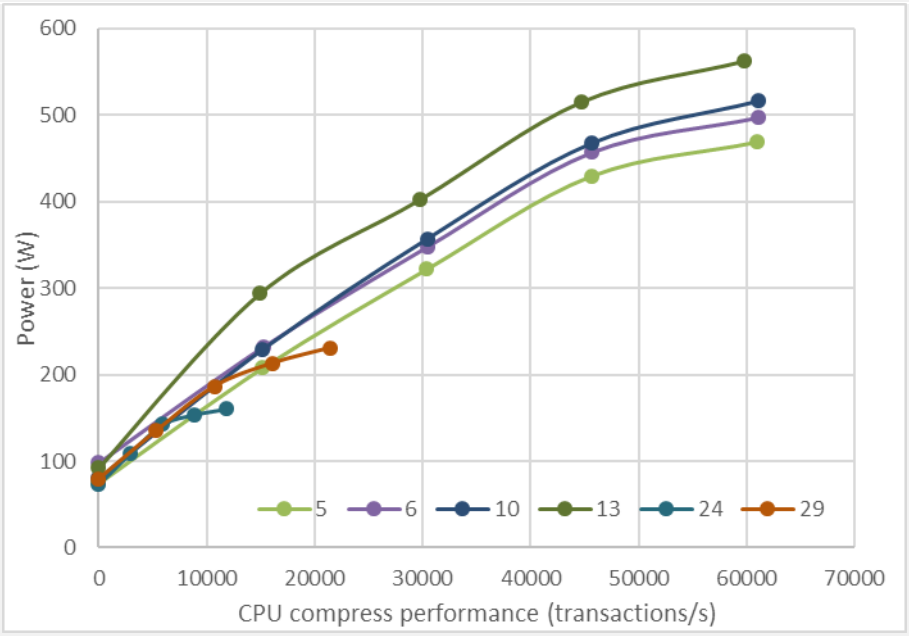
Using the SPEC data set, preliminary observations were drawn regarding the relation between the SERT worklets, and the power and performance results.

Results were analysed for thirteen servers, chosen to illustrate the effect of the varying individual components, CPU, RAM and number of HDDs while keeping the other components unchanged. These can be summarised by the six specifications of servers in Appendix 4 Figure 7.

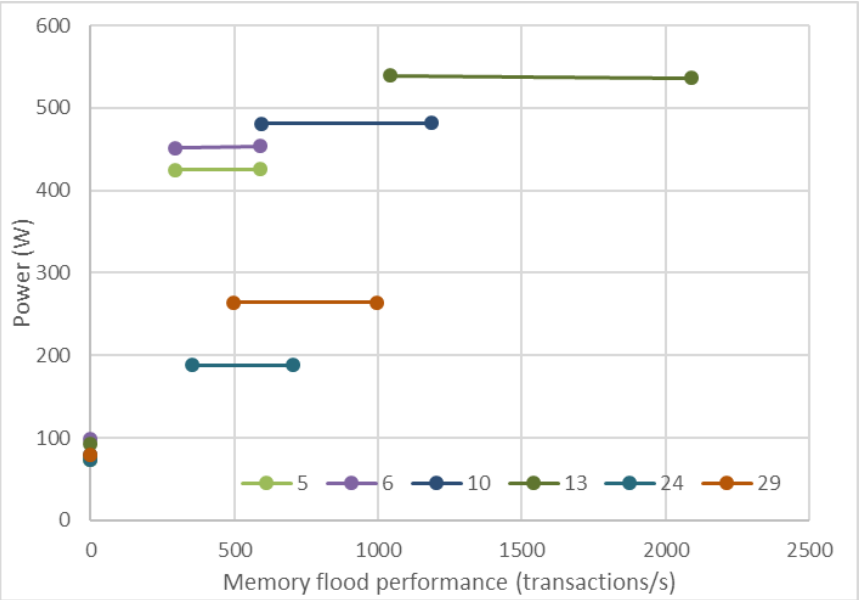
Sample number	CPU		Memory		Storage
	Number of CPUs x cores	Frequency (GHz)	DDR4 Modules / dimms (MB)	RAM (GB)	Number of HDDs
5	2 x 18	2300	8	64	1
6	2 x 18	2300	8	64	8
10	2 x 18	2300	16	256	1
13	2 x 18	2300	16	1024	1
24	2 x 6	1600	16	256	1
29	2 x 6	2400	16	256	1

Appendix 4 Figure 7 - Specifications of servers in preliminary analysis

Power was plotted against the tested performance for the various workloads, so that observations could be made regarding how the performance of the different server specifications varied in Appendix 4 Figure 8. (Note: The performance in each test is not normalised and due to the high correlation between worklets within a particular workload, only one worklet is shown).

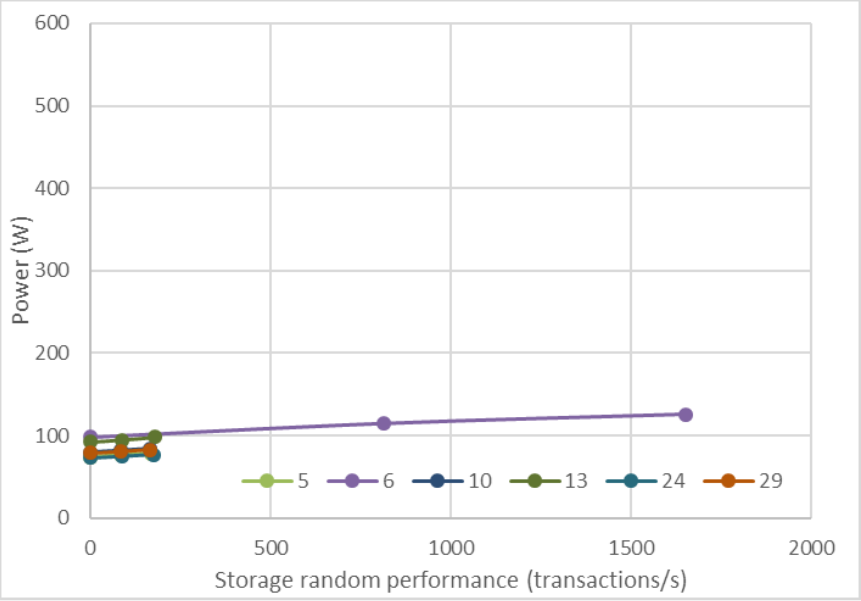
Workload	Observations on Plot of Power against Performance
CPU	 • Crossing point with the X axis shows the idle power demand. • Performance and power increase in proportion but not perfectly linearly (scale). • Data implies that impacts of RAM and storage due to power consumption are minimal because they contribute only a small proportion of overall power compared to the CPU. • Smaller CPUs are more efficient at lower performance but only at 75% load level and above, larger CPUs are overall more efficient at higher performance • Power consumed per unit performance is high relative to RAM, storage.
Hybrid	 • Performance and power increase in proportion (scale) very linearly except from 12.5% load and higher. • Data implies that very little of the performance score is influenced by RAM. • Smaller CPUs more efficient at lower performance and all load levels, larger CPUs at high performance - more variation in power consumption compared to CPU compress worklet.

Memory



- Doubling RAM doubles performance, but proportional power increase is very small compared to CPU and hybrid worklets.
- Power varies very little with load level - power predominantly influenced by CPU at 100% load rather than RAM.
- Performance is strongly influenced by CPU, not just the RAM
- Note: idle power is plotted but not connected to the other datapoints for visual clarity.

Storage



- Storage results are largely independent of RAM and CPU (except idle power)
- Doubling number of drives doubles performance
- Sample 6 has 8 drives, so shows eight times the performance but only 1.2 times the power relative to sample 5.
- Performance is largely independent of the storage device capacity
- Note: SSD performance is approximately 100x higher than HDD at slightly lower power, but no results were available at the time of analysis. However, HDD is often preferred for high capacity slower access applications, where it can be more efficient per GB.

Appendix 4 Figure 8 - Preliminary data analysis observations

5.3. Preliminary insights on combining workloads

Based upon the preliminary analysis above, the foundation approach to combining the workload scores into an overall efficiency metric should have the following key features:

CPU and hybrid workloads as key components: These workloads are the focus as they have a large influence on both power consumption and performance across CPU and memory worklets.

Memory workload included as a secondary component, weighted in relation to the CPU/hybrid workloads: Whilst doubling RAM doubles performance, the resultant relative power increase is very small. In addition, very few real world applications would show similar performance improvements with increasing RAM.

Storage workload not included: The storage workloads assess very different server characteristics - storage performance is almost completely independent of the RAM and CPU, and power consumption is only related to their idle power consumption. Therefore, combining the storage workload is not necessary, and may even negatively impact the metric, especially considering i) the 100-fold difference between SSD and HDD performance and efficiency scores ii) that the maximum SSD performance may not be attained in SERT during calibration iii) that newer PCIe SSDs are even higher performance iv) The SERT performance test is not designed to take into account the storage capacity which is a key criteria for selecting storage. Servers that are focused primarily on storage should be considered as storage and addressed via a storage-specific metric.

6. Principles for metric development

The broad objective of a metric is to provide an indicator of the energy efficiency and energy consumption of a particular server model and configuration under representative 'normal' use conditions. Ideally, this would meet the requirements in Appendix 4 Table 2(combined requirements as defined by project study team and industry):

Appendix 4 Table 2 Metric requirements

Metric Requirement	Explanation	Design considerations
Include consideration of power scaling / efficiency at lower utilisation	Use conditions for servers are varied, and optimal efficiency is very dependent on configuration - therefore it is preferable to optimise approaches toward the maximum potential for savings, whilst remaining agnostic and scalable from a user/market perspective. Many large scale data centres now operate very efficiently due to improvements in environmental management, combined with more energy efficient data centre equipment and improvements such as virtualisation[21]. However, smaller-scale data centres (i.e. server rooms for SME[22] workloads) are less likely to be optimally configured. In smaller data centres, users tend to have the least technical ability to select efficient products and are less likely to have specialist workloads such as high performance / supercomputing type work. Server loads tend to be lower - in many data centres, utilisation[23] levels may be between 5 and 12%[24]. Smaller server rooms represent 49% of the total electricity used by all data centres in the US, and it can be expected that the proportions are similar in the EU[25].	Metric should take into account both low and high utilisation, but account for power scaling in some way. Introduce a "dynamic range" factor to account for the degree of power scaling can assist, based on one of the following for example: • Max power : idle power • Max performance : idle power • Max efficiency score : idle power

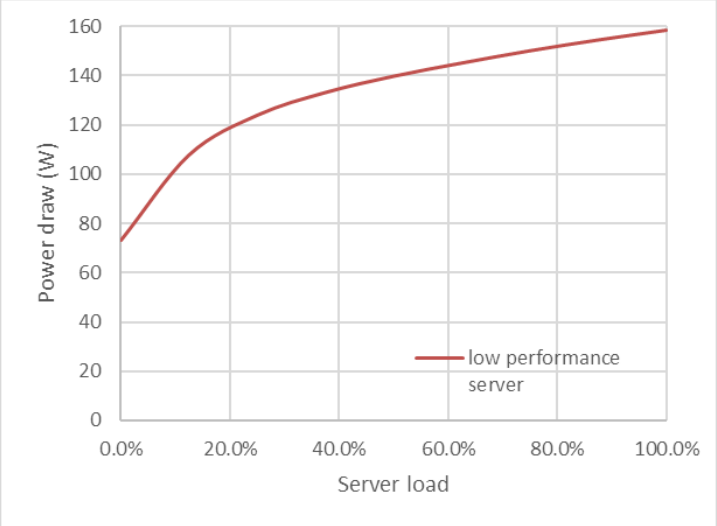
[21] server virtualisation is the running of multiple applications (virtual servers) on a single physical host server. Therefore, instead of many servers operating at low utilisation levels, virtualisation combines the processing power of many servers onto fewer servers operating at higher total utilisation rates.

[22] Small and medium-sized enterprises (SMEs) are defined in the EU recommendation 2003/361. The main factors determining whether an enterprise is an SME are: staff headcount and. either turnover or balance sheet total.

[23] Utilisation: the fraction of total computing resources engaged in useful work.

[24] Data centre efficiency assessment, NRDC

[25] Data centre efficiency assessment, NRDC

Metric Requirement	Explanation	Design considerations
	 **Server power characteristics**[26] Server utilisation is an issue due to the absence of power scaling with load level particularly for lower performance servers. This means that a server may use close to 50% of the power at full capacity when it is idling (not carrying out useful work) – see the figure above. Therefore, a metric which takes into account both low and high utilisation (e.g. by weighting results to this loading) could have the greatest informative and energy saving impact under real, common operating conditions. Alternatively, as performance at low load levels can be similar to idle, another option could be to determine a means of factoring in consideration of the idle power overhead or scaling with load level of a server within an active mode metric. The current ENERGY STAR metric is based on the idle power.	
Correct energy efficiency ranking	A metric should be tested against real data to ensure that it ranks servers by their energy efficiency in a representative way, keeping in mind the following questions: • Does the metric favour one size of server over another? • Does the metric agree with well understood results of energy efficiency comparisons?	Test metric out against hypothetical scenarios and real data
Appropriate workload weightings	The weighting of each workload in the metric should ideally be in proportion to the way in which real world servers perform similar work. A sample of a number of real world applications can enable assessment of the way in which the most important characteristics of a volume server (CPU and memory capabilities) affect its operation. An analysis of data provided by industry stakeholders, suggests that a weighting of CPU to memory worklets of somewhere between 60:40 or 70:30 is appropriate for a general server profile[27]. The storage workloads assess very different server characteristics, and incorporation of these worklets into a single metric could negatively impact the metric.	Use a 60:40 or 70:30 weighting of CPU to memory workloads. Give storage workload 0 weighting.

[26] Based on real data for hybrid performance of a 2015 server from the preliminary data analysis.

[27] The weighting determines the optimal server configuration between CPU performance and RAM performance/capacity. Analysis of the final metric shows optimal configurations in line with expectations, further supporting a weighting within this region.

Metric Requirement	Explanation	Design considerations
Appropriate worklet averaging	A metric can define alternative approaches to average worklet results into workloads. Ideally the end result should reflect how real-world performance changes with different configurations and types of server. The three main approaches to calculating averages of worklet results are detailed below: • **Arithmetic mean:** This is the most common approach for averaging. It is equivalent to assuming the time spent under each worklet is the same. Since this is not true under different configurations, this approach is not appropriate and would result in the highest performing worklet heavily dominating the metric result. For a series of n numbers, x_1, x_2,…,x_n. The arithmetic mean can be written as: $$ArithmeticMean = \frac{x_1 + x_2 + \cdots + x_n}{n}$$ • **Geometric mean:** A geometric mean is generally used to combine terms that have different scales in order to prevent the mean from being dominated by the largest item in the list. It is analogous to considering the total server performance as the cube formed by the axes of memory, CPU and storage performance. However, this analogy does not fully apply since the worklet power and performance measure the whole server already and the CPU and memory performance are interlinked causing some elements to be overcounted. The benefit of geomean compared to arithmetic mean is that high and low performing worklets have some influence on the result, although, the highest performing worklet will still dominate the overall result. This is the current proposed Industry approach. For a series of n numbers, x_1, x_2,…, x_n. The geometric mean can be written as: $$GeometricMean = \sqrt[n]{x_1 \times x_2 \times \ldots \times x_n}$$ • **Harmonic mean:** This is used to calculate the average rate given a fixed known output value. For example, the average speed of a car travelling at different speeds for fixed and known distances is calculated using the harmonic mean. This could be considered analogous to the performance of a server under different worklets. Since it is known that the performance under each worklet will differ based on the configuration, a fair comparison can be made under the same, fixed workloads. The harmonic mean is strongly influenced by the slowest rate. This means that it is impossible to reach an arbitrarily high average by targeting one worklet only and prevents gaming of results via configuration – for example by increasing memory. The harmonic mean favours a balanced configuration since addressing performance bottleneck will have the biggest impact on the average performance. For a series of n numbers, x_1, x_2,…,x_n. The harmonic mean can be written as: $$HarmonicMean = n \times \left(\frac{1}{x_1} + \frac{1}{x_2} + \cdots \frac{1}{x_n}\right)^{-1}$$	Use the harmonic mean to average worklet results to workload level
Intuitive	The metric should be understandable for non-experts likely	• Avoid different categories within

Metric Requirement	Explanation	Design considerations
	to be dealing with server energy efficiency measurements.	metric. • Simplify equations where this does not have an adverse impact.
Technology neutral	There are wide variations in the use of servers. For example, a virtualised server may operate 24 hours a day compared against an enterprise server which could sit in idle for a large proportion of time. In addition, different server technologies will have different features – for example the extra circuits contained in resilient servers. A recent report by The Green Grid highlighted that SERT efficiency scores and power consumption do indeed differ for particular technologies, e.g. HDD vs SSD, and resilient servers. Whilst a metric needs to account for different utilisation levels – for example, lower utilisation for resilient servers - ideally the metric itself would be neutral, with no integrated allowances for different segments (e.g. tower, rack, managed, resilient). This would ensure that policy makers could define product categories as considered necessary for their initiatives. A neutral metric approach can enable greater transparency, easier analysis/interpretation of results and greater longevity (market segments and technologies change over time but comparison is still necessary across all server types).	Avoid integrated allowances for different server types.
Interoperable	Interoperability relates to the ability of a product to work with other systems or products without needing to be specially adapted. Energy efficiency standards, metrics and policy need to avoid any negative impacts on interoperability. For example, metrics that emphasise power management at the cost of network availability and response times.	The metric analysis includes balancing the benefits of the increasing efficiency at all utilisation levels and interoperability.
Appropriately accounting for the influence of ambient temperature	Fan speeds and server power consumption increase at higher temperatures. For data centre operators, finding the optimal balance between internal server fan cooling and data centre cooling is necessary to optimise efficiency. The SERT tool does not measure how power consumption varies at different inlet temperatures and utilisation, and there is no standardised reporting of test results at different temperatures by manufacturers. It simply specifies a temperature range for testing and requires that the inlet air temperature is measured. It can be assumed that the most favourable temperature will be used for testing purposes and declared in the testing report, and therefore no correction for temperature is considered necessary to SERT results. Temperature performance is difficult to include within energy efficiency / performance metrics due to the additional testing requirements and integrating the additional information into a single metric while remaining informative and intuitive. Whilst SERT is designed for one inlet temperature, there is potential for tests to be carried out multiple times at different temperatures. This could provide useful information to data centre operators, but would have an impact on testing costs, especially as precisely controlled temperature environments for testing would be required.	Avoid inclusion of temperature as a factor in the metric and do not attempt to include any ambient temperature correction.

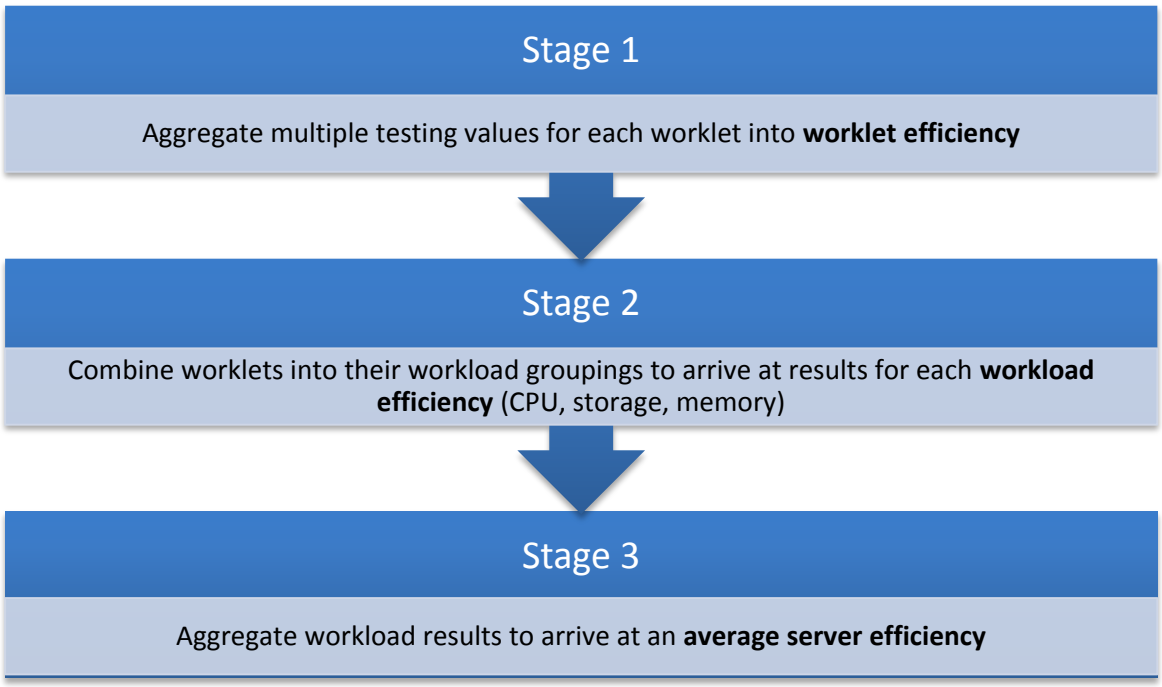
Metric Requirement	Explanation	Design considerations
Network / system level scalable	The SERT tool is only designed to assess the capability of one server in isolation, and cannot provide a systems perspective on how the server would perform if 5 were in use in tandem. Metric results would ideally be able to make effective comparisons between differently performing products. For example, comparing the efficiency between the following options for the same application: o Solution A: One high performance server with high efficiency and relatively high idle power o Solution B: Three lower performance severs with similar efficiency and lower idle power. While the smaller servers in the above example seem individually more efficient, when the idle powers are combined, total idle power is much higher and therefore this configuration is more inefficient. The behaviour of the metric in this type of situation is analysed in Appendix 4 Annex2.	Assess system scalability of metric by testing results using a "deployed power" approach to ensure that the metric favours servers which result in lower datacentre deployment power.
Avoids negative market influence	It is important that metrics are designed: • To avoid unintended consequences such as over-specified servers that operate in-use at lowered utilisation, and therefore lower efficiency. • To appropriately incentivise design changes which will decrease actual energy efficiency. • To avoid design changes resulting in a decrease in end user energy efficiency having a counter intuitive increase in the SERT score. • To carefully consider the profound market influence a metric could have in defining what is considered an "optimal configuration".	Test metric out against hypothetical scenarios and real data. Ensure that configuration to test is representative.

7. Final Metric

7.1. Average Server Efficiency Metric

The metric was developed iteratively, testing options against data and in consultation with industry experts. The final proposal is a metric is based on a 60:40 weighting, but using the harmonic rather than geometric mean. Whilst the metric relates to the calculation of average efficiency, calculations of indicators for server average performance and average power consumption may also be useful for policy makers and are included in the Appendix 4 section 7.2.

The way in which the metric is derived from the SERT testing results can be broken down into three stages as shown in Appendix 4 Figure 9.



Appendix 4 Figure 9 – Stages in determining average server efficiency from SERT results

7.1.1. Stage 1 : Calculating the efficiency for each worklet

At the worklet level, the efficiency calculation is identical to the SPEC approach:

$$worklet\ efficiency = \frac{\sum performance\ at\ each\ utilisation\ level}{\sum power\ at\ each\ utilisation\ level}$$

Different approaches have been explored, such as using the peak power measured, but the above approach is considered the most intuitive. This is because each worklet has a different number of utilisation levels, and this would need to be taken into account before aggregating the worklets to calculate the workload efficiency.

7.1.2. Stage 2: Calculating the efficiency for each workload

For each workload, all the associated worklets are combined using the harmonic mean of the worklet efficiency results (see Appendix 4 Figure 2) to calculate the workload efficiency.

$$workload\ efficiency = no.worklets \times \left\{\sum \frac{1}{worklet\ efficiency}\right\}^{-1}$$

7.1.3. Stage 3: Combining workload components into an average server efficiency metric

The server average efficiency metric is calculated as the harmonic mean of the workload results, weighted based upon typical server work ratio of 60:40 CPU to memory[28].

Two options for the server metric are specified below – with or without the dynamic range component (ratio of idle to max power). Dynamic range makes the metric more suitable for applications where there is lower utilisation, increasing the impact of idle mode. For example, there may be long idle periods at night in many situations, including for geo-located cloud services. Therefore, the inclusion of dynamic range provides the most representative results for single server deployments and multiple server deployments spending more than 15% of the time in idle (typically low performance servers intended for use in server rooms and small data centres). The other more general metric without dynamic range is better suited to large-scale deployments where utilisation is constant and higher, i.e. for servers in highly efficient data centres. As level of utilisation is not always clear when the product is placed upon the market, provision of values for both metrics provides the most informative insight.

High utilisation

$$server\ efficiency = \left\{\frac{0.6}{CPU\ eff} + \frac{0.4}{memory\ eff}\right\}^{-1}$$

Low utilisation

$$server\ efficiency = \underbrace{\left(\frac{idle\ power}{\max power} + 1.5\right)}_{Dynamic\ range} \times \left\{\frac{0.6}{CPU\ eff} + \frac{0.4}{memory\ eff}\right\}^{-1}$$

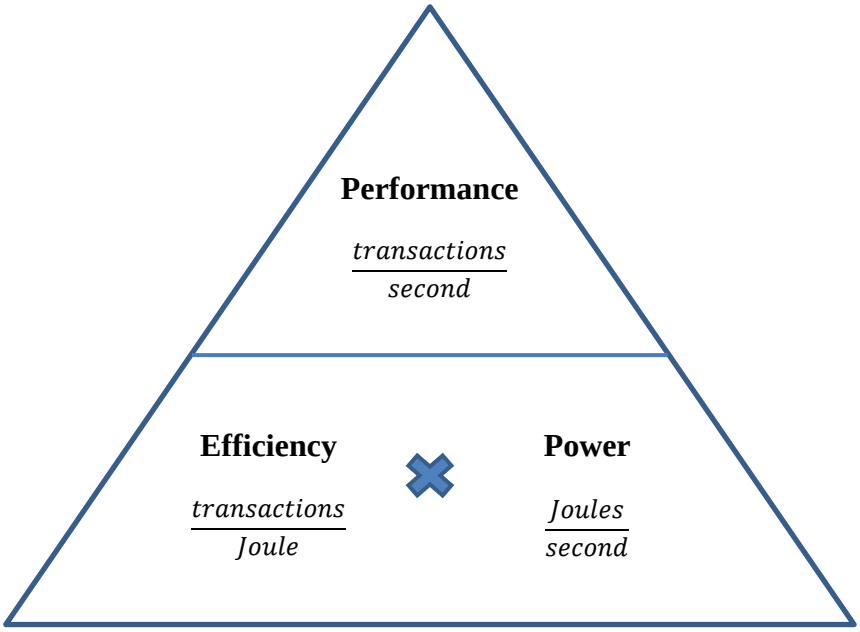
The dynamic range has a factor of 1.5 added to it in order to ensure that it has the appropriate weighting in the calculation to represent and compare single servers of different performances (annex 1) and the proportion of time spent in idle and active modes in a deployment of multiple servers (this was determined through the analysis detailed in annex 2).

[28] For comparison purposes, this weighting is considered constant. However, for informational purposes, calculations to more closely represent the specifics of real life applications can be made easily by changing the weightings – for example to consider a CPU intensive ratio of 85:15 or a memory intensive ratio of 40:60.

7.1.4. Supplemental indicators for average server power consumption and performance

Similar to the server efficiency metric calculation, indicators for average server power consumption and average server performance can be calculated. These enable calculation of the number of servers required for a specific deployment, as well as more accurate estimates of the power consumed by the servers under different workloads, such as the CPU intensive and memory intensive workloads. It is expected this will be valuable to data centre operators with very well defined requirements.

There interrelation between power, performance and efficiency terms is shown in Appendix 4 Figure 10.



Appendix 4 Figure 10 - Power, performance and efficiency relationship

7.1.4.1. Stage 1 : Data collection at the worklet level

As well as the worklet efficiency, the following information can be useful for input to the next stage:

- Worklet power sum: This is the sum of the power consumption at each utilisation level of each worklet. It is abbreviated to *worklet power sum*
- Worklet performance sum: This is the sum of the performance at each utilisation level of each worklet. It is abbreviated to *worklet perf sum*
- Number of utilisation levels per worklet (or number of worklet power measurements)

7.1.4.2. Stage 2: Calculating individual workload power consumption and performance

At the workload level, the performance and power consumption can also be calculated. For each workload, all the associated worklets are combined (see Appendix 4 Figure 2).

Workload Performance :

The worklet performance sum is divided by the number of utilisation levels to give the average (arithmetic mean) across all utilisation levels to account for the different number of utilisation levels in different worklets. The amount of time spent at different utilisation levels is assumed to be the same and is not influenced by the configuration, therefore, the arithmetic mean is used.

$$workload\ perf = harmonic\ mean\ \frac{worklet\ perf\ sum}{no.utilisation\ levels}$$

$$workload\ perf = no.worklets \times \left\{\sum \frac{no.utilisation\ levels}{worklet\ perf\ sum}\right\}^{-1}$$

Workload Power Consumption:

The workload power depends on the power consumed by each worklet and the relative time spent executing each worklet. Each worklet power must therefore be weighted by the time spent when calculating the average workload power. Since the time spent on each worklet is inverse to the performance, the worklet power consumption is weighted by the inverse of the worklet performance (see Appendix 4 Figure 10).

In addition, the worklet power sum and performance sum are divided by the number of utilisation levels to give the average (arithmetic mean) across all utilisation levels to account for the different number of utilisation levels in different worklets. The amount of time spent at different utilisation levels is assumed to be the same and is not influenced by the configuration, therefore, the arithmetic mean is used.

$$workload\ power = inverse\ performance\ weighted\ arithmetic\ mean\ of\ \frac{worklet\ power\ sum}{no.utilisation\ levels}$$

$$workload\ power = \left\{\sum \frac{no.worklet\ utilisation\ lvl}{worklet\ perf\ sum}\right\}^{-1} \times \sum \frac{worklet\ power\ sum}{worklet\ perf\ sum}$$

7.1.4.3. Combining workload components into the server performance and power consumption indicators

The server performance and power indicators are therefore calculated using the harmonic mean and the weightings in line with the efficiency calculation as follows:

Average server performance:

$$server\ perf = weighted\ harmonic\ mean\ of\ workload\ performance$$

$$server\ perf = \left\{\frac{0.6}{CPU\ perf} + \frac{0.4}{memory\ perf}\right\}^{-1}$$

Average server power consumption:

$$server\ power = inverse\ workload\ and\ performance\ weighted\ arithmetic\ mean\ of\ workload\ power$$

$$server\ power = \left\{\frac{0.6}{CPU\ perf} + \frac{0.4}{memory\ perf}\right\}^{-1} \times \left\{\frac{0.6 \times CPU\ power}{CPU\ perf} + \frac{0.4 \times memory\ power}{memory\ perf}\right\}$$

8. Server Efficiency Metric evaluation against development principles

8.1. Applying server efficiency metric to real server data

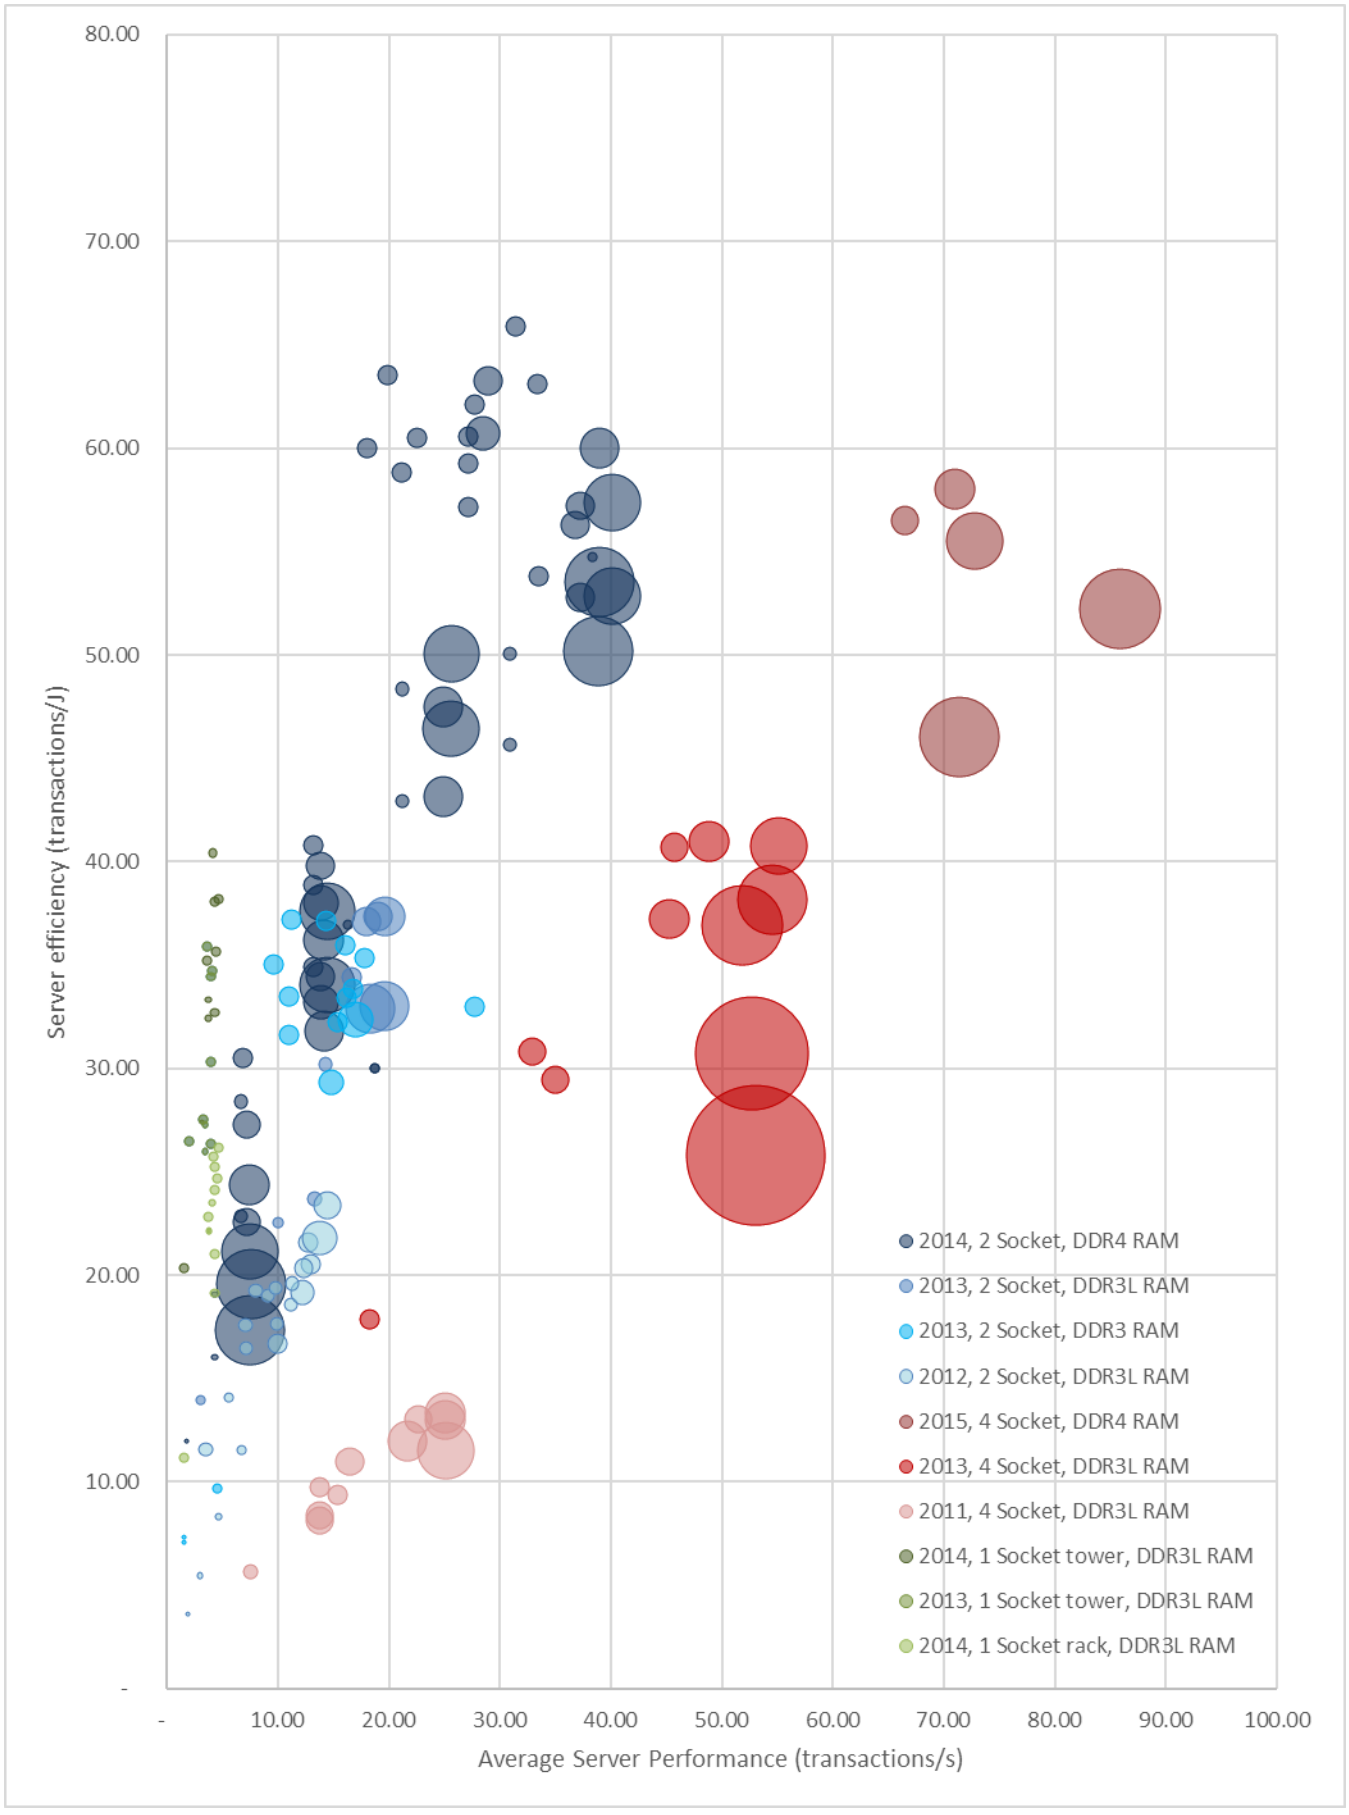
Metric results were calculated for the data points from the detailed SPEC data set (Appendix 4 Figure 11). The efficiency calculated from the efficiency metric (in the case of servers with low utilisation) is plotted on the y-axis and the average server performance, as defined in Appendix 4 Section 7.1.4.3**.** above on the x-axis. In addition, the size of the bubbles indicates the RAM in each configuration. This is used to verify the optimal CPU:RAM weighting by indicating where efficiency starts to fall as RAM increases above the optimal level. The data points are grouped by colour to indicate the different types of server, i.e. number of sockets and form factor, and shade to indicate the approximate server generation.

In terms of performance, there is a clear distinction between server types and generations. As expected, 1 socket servers have the lowest performance, as shown by the near horizontal line of green datapoints, and 4 sockets (red) have the highest performance. Server performance approximately doubles between 2-4 socket servers as expected since they use similar CPUs. Server performance has also increased substantially across generations for 2 socket and 4 socket servers. The latest generation of two socket servers show the widest range of performance, in line with wider range of configurations available and tested.

Efficiency has improved very rapidly over the past 4-5 server generations and the efficiency improvement made from one generation to the next is generally more significant than the efficiency variation between server configurations in the same generation. Overall, a wide range of efficiencies is observed, with the most efficient server from 2015 over ten times the efficiency of a 2011/12 server. Even in the latest generation of 2 socket servers, the most efficient is five times the efficiency of the least. Since this is the largest volume market by sales, this suggests there is still significant scope to improve efficiency.

Energy efficiency correlates with the average server (computation) performance within the same generation. This supports the accuracy of the metric at a product level.

In conclusion, the efficiency metric is able to distinguish the efficiency of servers of different types and configurations and in accordance with expected behaviour.



Appendix 4 Figure 11 - Efficiency / Performance comparison for analysis data set

8.2. Evaluating metric in relation to requirements

Appendix 4 Table 3 shows how the proposed metric performs in relation to the key metric requirements defined earlier.

Appendix 4 Table 3 - Evaluation of proposed metric

Metric Requirement	Evaluation of final metric
Correct energy efficiency ranking	✓ Ranking of theoretical servers at the extremes is intuitive. See Appendix 4 Annex 1 for details of the analysis to support this conclusion.
Include consideration of power scaling / efficiency at lower utilisation	✓ The metric accounts appropriate for efficiency at lower utilisation through the inclusion of the dynamic range. See Appendix 4 Annex 2 for details of the analysis to support this conclusion.
Appropriate workload weightings	✓ The metric adopts workload weightings that are consistent with data provided by industry. See Appendix 4 section 8.1
Network / system level scalable	✓ The metric allows for the influence of server configuration and components (excluding storage) on performance to be taken into account and can be applied to a wide and growing range of server types supported by SERT including 1-4 socket volume servers, and blades. Whilst it does not consider the wider data centre efficiency, it does at least use dynamic range to ensure that idle power overheads of each server are factored in to the final result. See appendix 4 annex 2 for details of the analysis to support this conclusion.
Appropriate worklet averaging	✓ This approach has been justified from a mathematical foundation, see Appendix 4 section 6
Intuitive	✓ The metric avoids unnecessary categorisation and simplifies equations on the basis of mathematical principles.
Technology neutral	✓ The metric is not technology specific, enabling policy makers to apply it to different server types as considered appropriate.
Interoperable	N/A
Appropriately accounting for the influence of ambient temperature	N/A
Avoids negative market influence	✓ The metric does not prefer smaller servers or over penalise HDD vs SSD.

In conclusion, this proposal provides a robust metric, meeting the key design requirements. It has been developed from a solid mathematical foundation and tested against real server data. The metric has potential to be applied in policy and industry-led initiatives.

The (optional) inclusion of the dynamic range component of the metric is important in order to encourage the reduction of idle power especially since this does not vary or scale as performance reduces, and there still exist widespread inefficiencies at low loading.

9. Guidance on products to test

A metric should be supported by a determination of the types of products that it can be applied to. SERT can be used effectively and consistently across 1-4 socket rack and tower servers, blades, and resilient servers, including new to market server specifications and different architectures (x86, POWER, ARM) but it is not suitable for every type of server and computer architecture. SERT is not endorsed to run on 8-socket servers, supercomputers, mainframe servers, clusters of multiple servers, and MIPS architecture. In addition, it does not test the performance of graphics card, field programmable gate array (FPGA) and other specialised hardware. Therefore, the scope of any policy requirements needs to be clearly defined.

As servers are highly customisable it is necessary to determine which particular configurations should be considered representative to test. ENERGY STAR defines "product families" and performance categories so that the highest consuming product configuration in each category can be tested to ensure that if this configuration is qualified, it is assumed that all products within this family will comply with requirements. ENERGY STAR v2 defines four categories, minimum power, low performance, typical and high performance.

Analysis shows that the proposed metric can be applied to all configurations defined by ENERGY STAR. While the metric is not restricted by configuration, server models can have very many permutations of different CPUs, RAM, HDDs, I/O devices and even PSUs. Testing every possible configuration is therefore clearly not feasible.

The ideal solution would fulfil the following criteria:

- Allow the buyer to compare the efficiency between different models and specific configurations before purchase.
- Enable monitoring and verification activities (i.e. ensure that declared configurations are available for purchase)...
- Not entail excessive costs or resources for the manufacturer.

The current recommended solution is that manufacturers test and declare at least two configurations, representing low, and high performance. These should align with the configurations advertised on the manufacturers website to assist consumer insight and enforcement activities. Since the high and low performance configurations have different uses and different efficiency levels, this defines a boundary range of efficiencies within which the other configurations in the product family will perform. The minimum and maximum power are not recommended since these are not representative of real life, useful configurations.

In line with industry recommendations, the following configuration restrictions could be applied:

- High performance configuration to include two SSDs and low performance to include two HDD. Since the storage worklets are not included in the metric, standardising this improves comparability.
- Installed memory optimised for performance and efficiency in combination with the installed CPU. This ensures an optimal configuration is provided and marketed.
- No additional I/O cards, or computational cards should be included. This is because the metric does not measure performance of these devices, which are only configured if there is a specific need for them.

Note: While these are fewer than the options defined in the ENERGY STAR version 2 specification, these configurations are being considered in the development of the next (version 3) specification.

In addition, the possibility of developing a model to estimate the efficiency of other configurations could be investigated with manufacturers. Power calculators are already available from a number of manufacturers that allow in depth configuration of servers. The following points should be discussed:

- Configuration boundaries – what exotic and atypical configurations can justifiably be excluded?
- What additional data is required/available, such as component level power data?
- What accuracy range is possible/desirable?

Appendix 4 Annexes

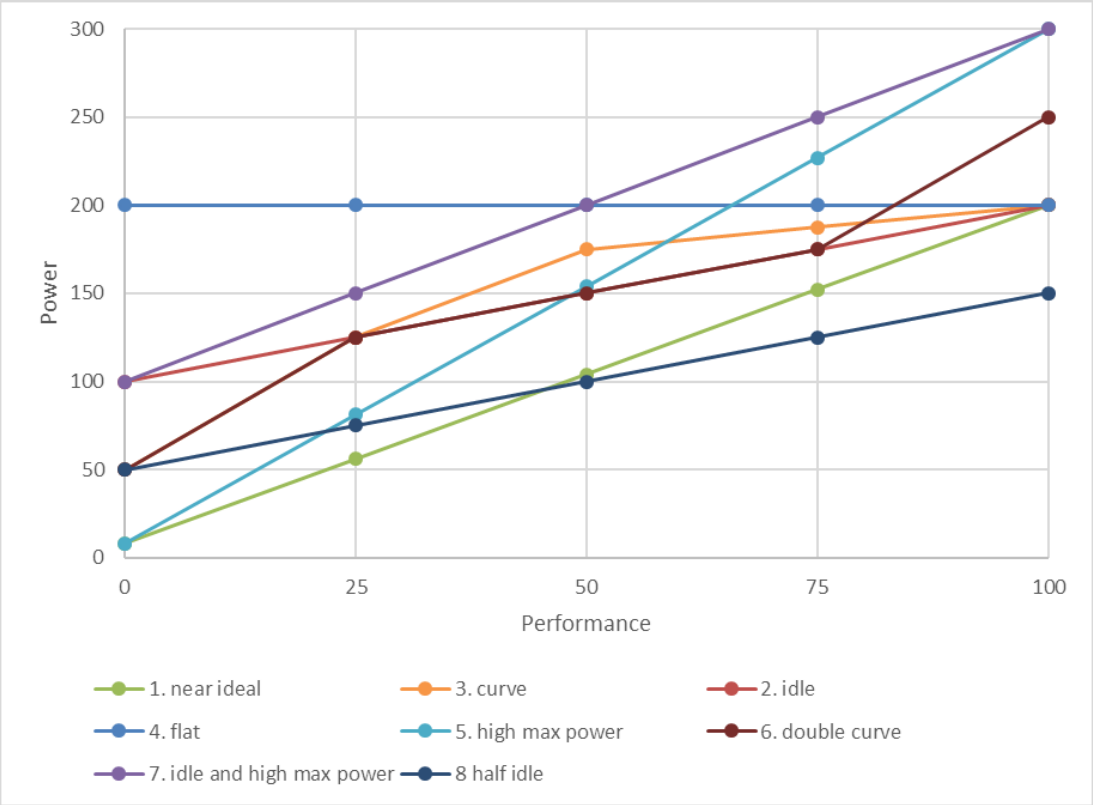
Appendix 4 Annex 1: Testing metric for achievement of appropriate energy efficiency ranking of products

The metric was assessed when applied to a range of hypothetical efficiency curves to ensure that it behaved as expected and was able to respond to future server developments. Hypothetical (boundary) curves (see Appendix 4 Figure 12) were used to represent a wide range of possible server configurations. This was preferred to evaluating the metric against current data, as currently all idle power levels are very similar. Therefore later evolutions in idle power could run the risk of making it inapplicable unless extreme boundaries were assessed.

Appendix 4 Table 4 shows the hypothetical efficiency curves used in the assessment, with arbitrary power / performance values.

Appendix 4 Table 4- Hypothetical efficiency curves considered

Hypothetical curve	Line type	Idle power	Max power
1. Near ideal	Linear	5W	200W
2. Idle	Linear with non zero idle	100W	200W
3. Curve	log curve	50W	200W
4. Flat	Linear (flat),	200W	200W
5. High max power	Linear (similar to “near ideal” but higher max power)	5W	300W
6. Double curve	Inverse S curve	50W	250W
7. Idle and high max power	Linear with non zero idle and high max power	100W	300W
8. Half idle	Linear with non zero idle and low max power	50W	150W



Appendix 4 Figure 12 - Hypothetical efficiency curves modelled

The "near ideal" efficiency curve, representing perfect scaling between performance and power, was used as the baseline by which to normalise the results for the other hypothetical performance lines for comparability.

Appendix 4 Table 5 shows that the proposed metric results in an intuitive ranking of the different hypothetical servers:

Appendix 4 Table 5 - Relative performance of hypothetical curves

	1 ideal	2 idle	3 curve	4 flat	5 high max power	6 double curve	7 idle and high max power	8 half idle
SERT	100%	79%	74%	64%	67%	73%	57%	114%
peak:idle	100%	9%	18%	4%	100%	18%	9%	18%
Proposed metric: sert:(idle/max) DR "+1.5"	100%	61%	66%	39%	68%	66%	48%	96%

The emphasis on higher utilisation efficiency is shown in the SERT metric by the relatively small drop change in the metric as the idle power increases (4 flat curve), and the larger drop as the peak power increases. The peak idle metric is calculated by dividing the peak performance by the idle power. This tends to show the opposite effect, and a drop in idle power causes a very significant drop in the metric. Furthermore, since the peak power is not taken into account, any change of peak power has no effect on the metric result (5 high max power) and this approach is therefore considered unsuitable.

Altering the SERT metric by dividing by the ratio of the idle power/maximum power ("dynamic range") increases the impact of low utilisation efficiency and idle power on the overall metric. The dynamic range is weighted by a "+x" figure to change the impact. In this case, the weighting uses "+1.5". This results in the flat power curve efficiency metric dropping to 39% of the near ideal

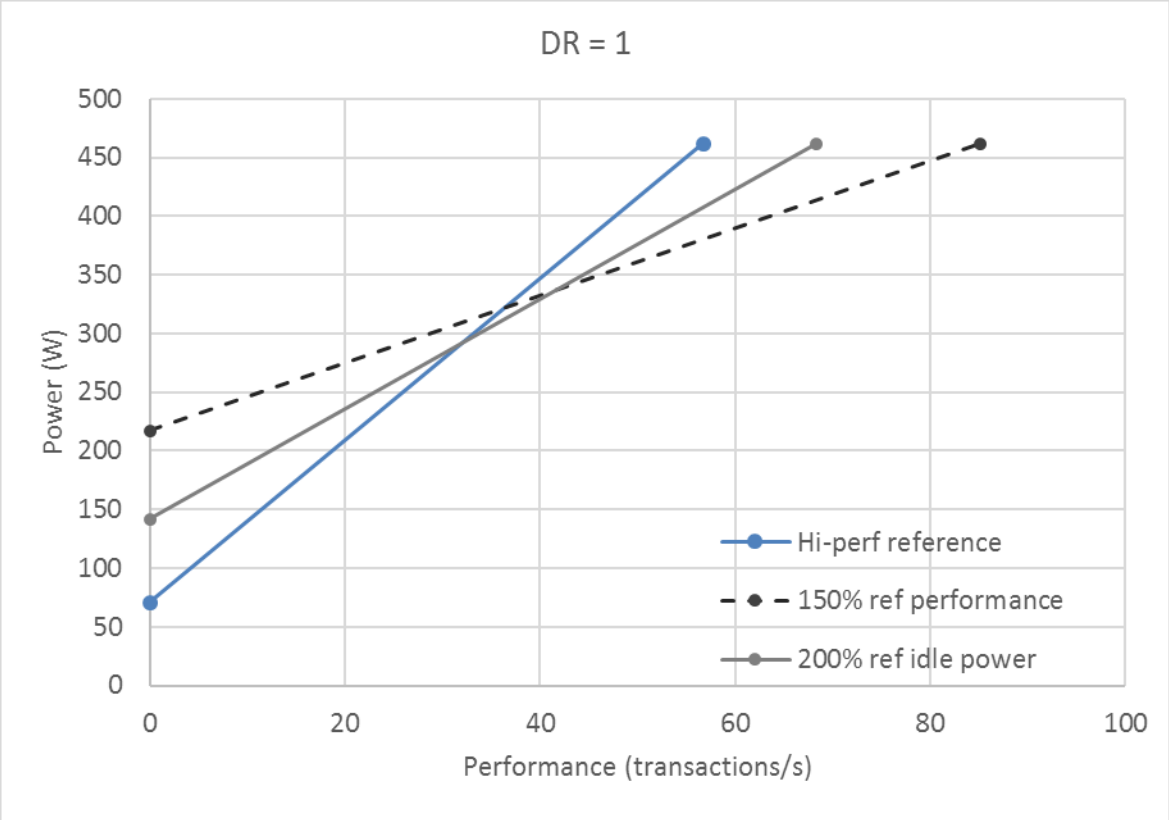
scenario which would be expected when power consumption is higher over all utilisation levels and is over twice as high at any utilisation level below 50%. It also takes into account the impact of raising peak power (5 high max power). Finally, it shows an intuitive balance when idle increases while peak power reduces by the same amount, giving almost the same efficiency as the ideal scenario.

Investigating dynamic range (idle/max ratio) between different performance servers

The previous analysis is based on servers with equal performance to understand the metric behaviour with different power profiles. In addition, it is important to understand how the metric will perform between servers of different performance. It is expected that the maximum performance of servers will continue to increase and that this will be achieved while maintaining the similar maximum power since current data centres and servers have reached their power and cooling design limits.

One of the impacts of this design strategy is that idle power is expected to increase incrementally. This will create a trade-off between increasing performance and increased idle power. Appendix 4 Figure 13 shows the performance curve of a current high performance server in blue and two hypothetical servers in grey, one with 50% higher performance and one with double the idle power. To simplify the comparison, the efficiency curves are assumed to be completely linear. At lower utilisation, the reference server consumes less power and is therefore more efficient. The difference in power then narrows until the point at which they meet, after which the hypothetical server is more efficient. Determining which server is more efficient in real life therefore depends on the average utilisation level and where the crossover occurs. As a result, it becomes more critical that utilisation levels remain high for such future products and the metric should reflect this. It is important to remember the average utilisation level must also include any time spent in idle.

This section analyses how the metric behaves in this scenario, over a variety of idle and performance increases to determine where this trade off point lies according to the metric.

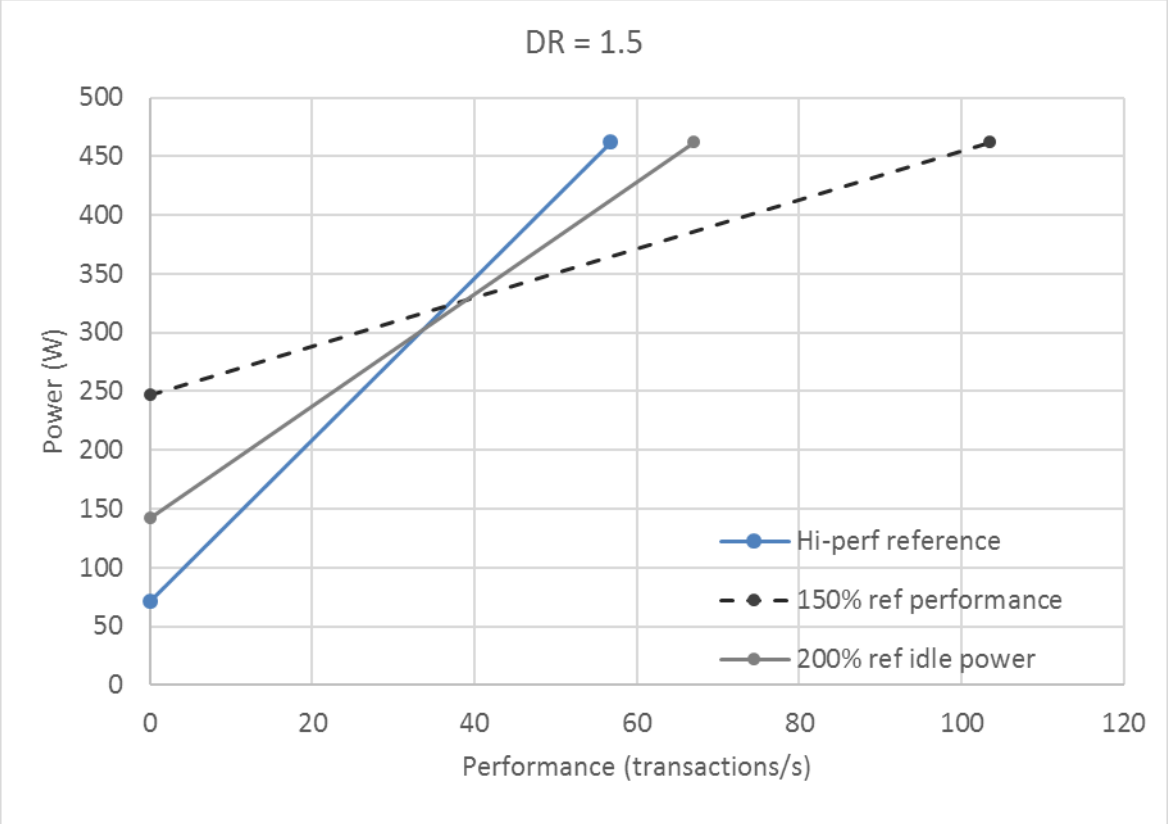


Appendix 4 Figure 13 - Power-performance curve of servers with equal efficiency under metric with dynamic range weighting of +1

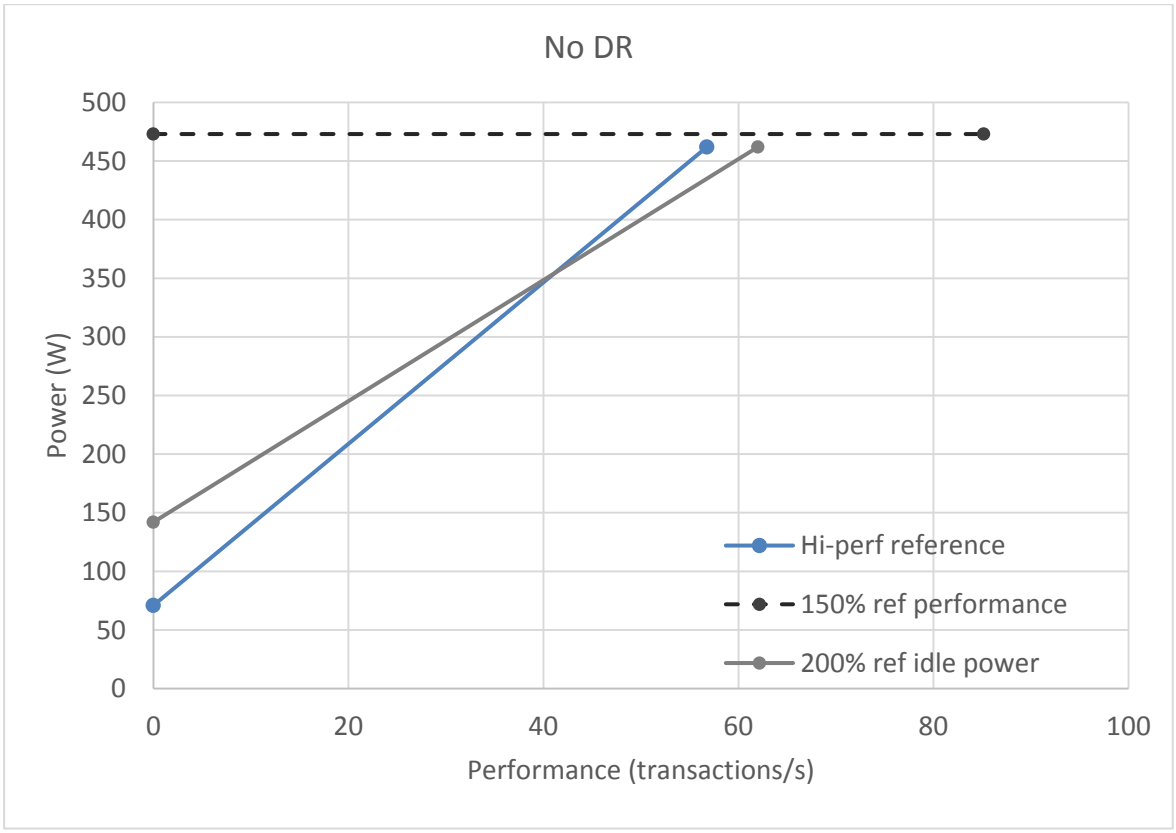
Appendix 4 Figure 13, Appendix 4 Figure 14 and Appendix 4 Figure 15 show three efficiency curves calculated to give the same efficiency results based on the metric formula under different idle/max weightings. Appendix 4 Figure 14 uses '+1.5' the same weighting used in the hypothetical curve

analysis, Appendix 4 Figure 13 reduces the number and to '+1.0' and therefore increases the weighting of the idle/max. Finally, Appendix 4 Figure 15 has no idle/max ratio applied.

As expected, increasing the weighting of the idle/max ratio results in the crossover point falling to lower utilisation levels. Without the idle/max ratio, a server with 50% higher performance and the same efficiency would have a flat power consumption over 470W which is above the maximum power of the reference curve. This highlights the high utilisation level assumed by the SERT efficiency score. With the idle/max ratio of "+1", a doubling of idle power with just a 20% increase in performance results in the same efficiency and cross-over utilisation level of approximately 50%.



Appendix 4 Figure 14- Power-performance curve of servers with equal efficiency under dynamic range weighting of +1.5



Appendix 4 Figure 15 - Power-performance curve of servers with equal efficiency with no dynamic range weighting

The conclusion of this evaluation was that the SERT approach, combined with an idle/max factor of '+1' or '+1.5' provided the most representative differentiation between individual products with different performance levels by increasing the impact of efficiency at lower utilisation levels in the efficiency metric.

However, this analysis is only valid for comparing individual servers of relatively similar performance. As the difference between the performance increases, it is no longer becomes valid because multiple servers are needed to match performance of the single high performance server. This is not accounted for and therefore a deployment approach which also calculates the number of servers is required.

Appendix 4 Annex 2 : Deployed power approach and testing metric for system scalability under different utilisations and workloads using real data

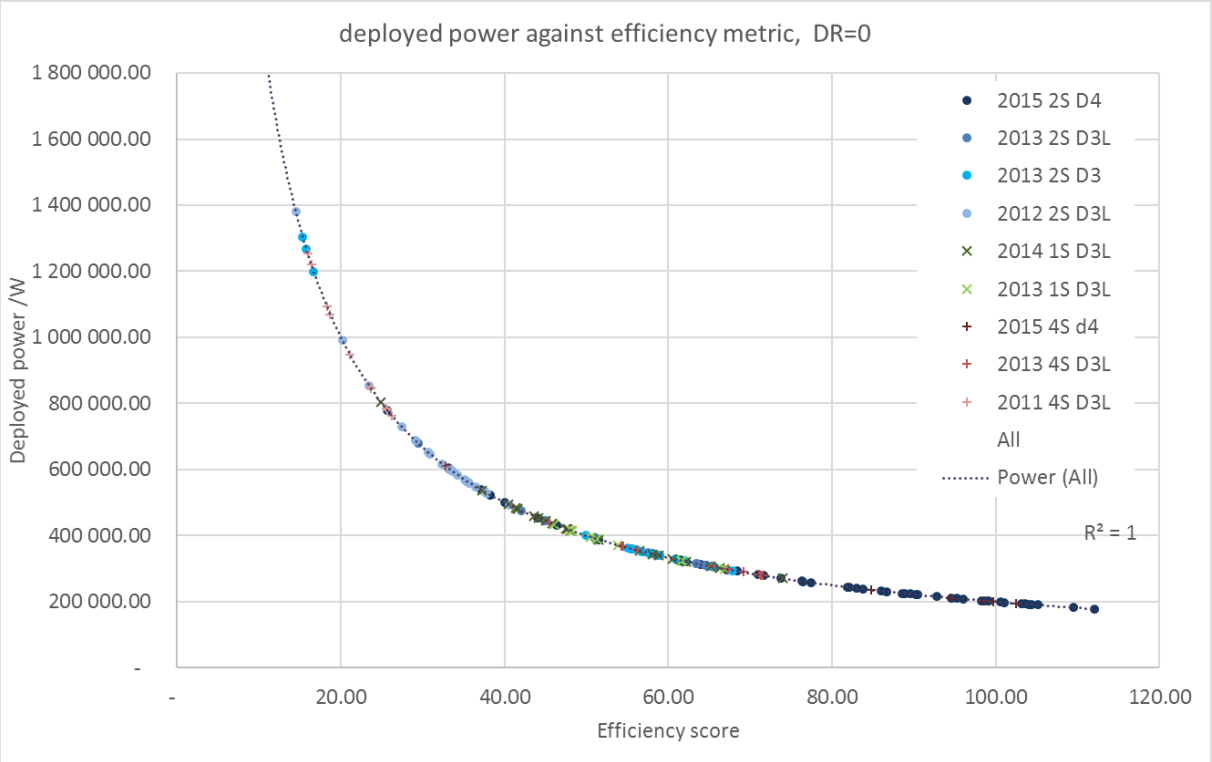
The deployed power approach is designed to take into account the efficiency at system or data-centre level. In this situation, servers are not purchased as individual items but in larger quantities with the intention of using them together in a larger system to fulfil a larger volume of work.

The methodology for calculating the deployed power can be described as:

1. Determine the total performance required from the system
2. Calculate the number of servers required based on the individual server performance
3. Calculate the total power based on the server power and number of servers.

Steps 2 and 3 can then be repeated for various servers and configurations to find the lowest power, and therefore most efficient system. The accuracy of the efficiency metric can then be determined by the correlation between the deployed power and the metric result.

Due to the metric development approach, comparing the deployed power directly against the efficiency metric will give a perfect correlation Appendix 4 Figure 16. This only confirms the formulas for efficiency, average power and average performance are mathematically coherent. Therefore, comparisons at different utilisation levels and against different metrics are used to understand how the correlation changes and the validity of the metric under different scenarios.



Appendix 4 Figure 16 - Average deployed power plotted against efficiency metric

Performance of metric at different utilisations levels

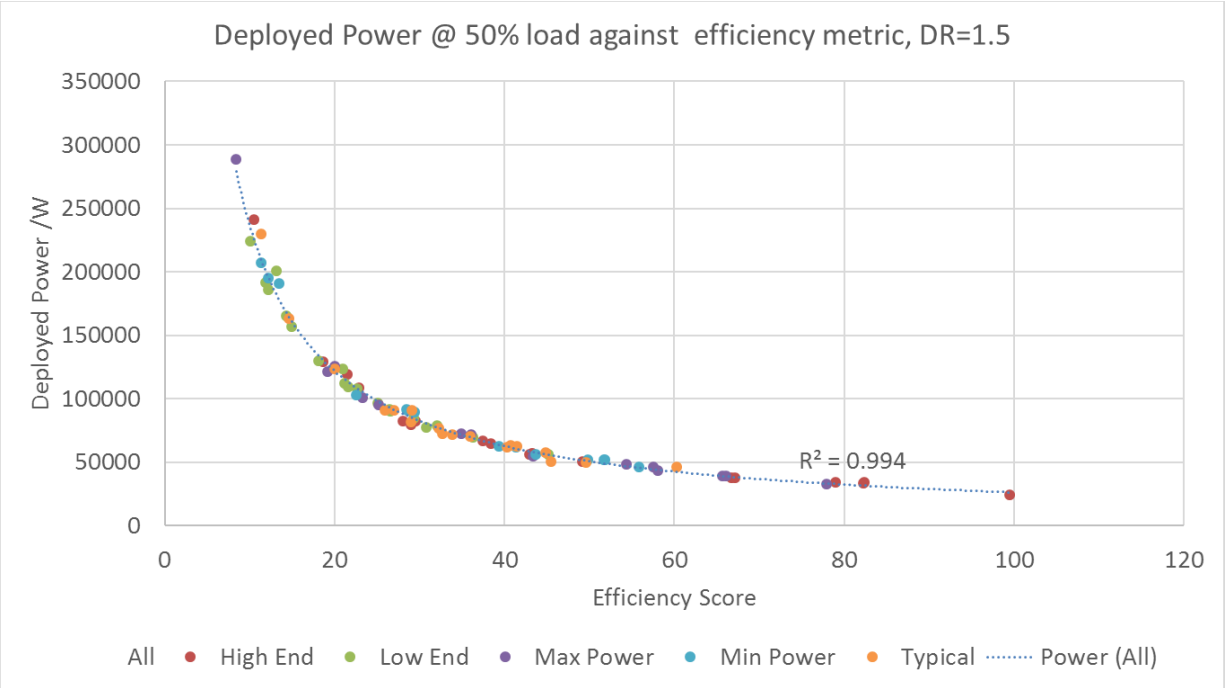
Instead of using only the average power, the power consumption at different utilisation levels can be calculated either from a fixed utilisation level, taken from the SERT test data, or the combination of average power and idle power. Appendix 4 Table 6 shows the level of correlation between the deployed power at different utilisations levels against the efficiency metric with different idle/max ratios. Based on the analysis, the DR increases the correlation with lower utilisation levels when calculated as a mix of idle and utilisation levels. At fixed utilisation, the DR causes a drop in correlation.

Based on the SPEC dataset across all types of server, this analysis suggest that a DR="+1.5" gives a better overall match when time in idle is over 15% (approximately 55% utilisation). The overall correlation under fixed utilisation is very high for both metric options, with No DR being slightly better.

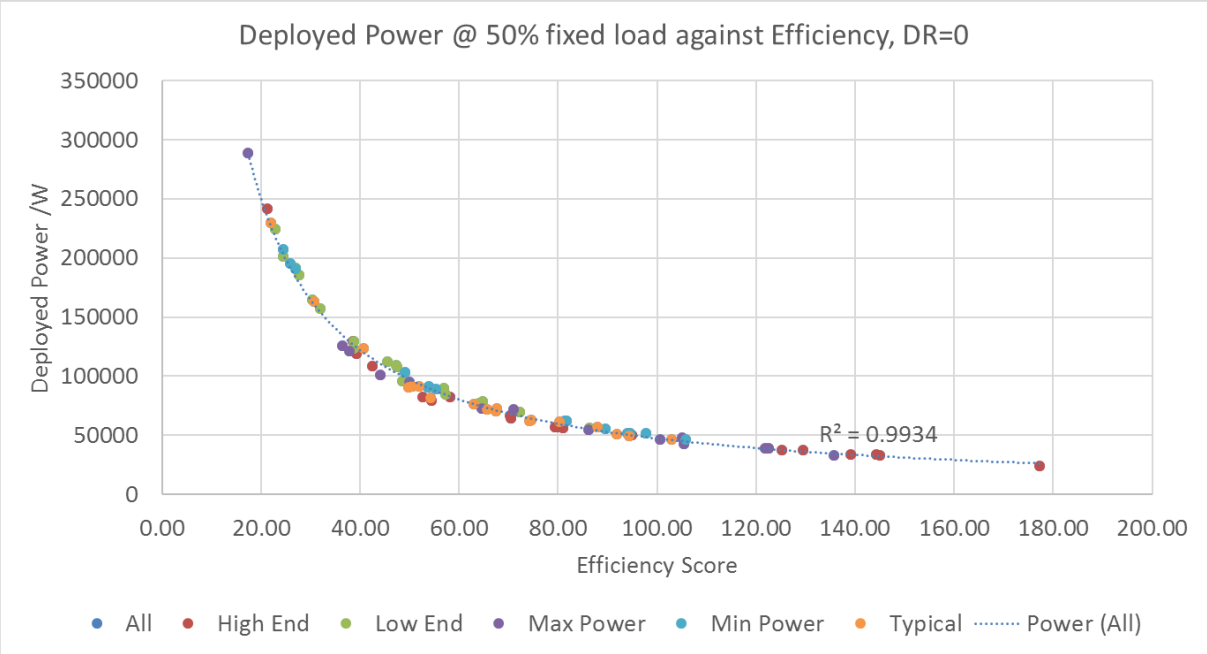
Appendix 4 Table 6 - Metric correlation with deployed power using balanced workload under various utilisation levels, with and without dynamic range

Balanced workload	DR=1.5	No DR
50% fixed utilisation	0.9903	0.9953
25% fixed utilisation	0.9913	0.9940
Sum of utilisation levels (approx. 64% average utilisation)	0.9916	1.0000
50% average utilisation (22% idle, 78% sum utilisation)	0.9979	0.9975
25% average utilisation (60% time idle, 40% sum utilisation)	0.9916	0.9702
Idle	0.9177	0.8674

Appendix 4 Figure 17 and Appendix 4 Figure 18 show examples of the same analysis based on the 2 socket managed servers data provided by The Green Grid. This shows a slightly higher level of correlation for DR=1.5 but slightly lower for no DR. Overall it supports the accuracy of the metric against the deployed power but suggests the appropriateness of DR or no DR is not absolutely clear.



Appendix 4 Figure 17 - Deployed power at 50% fixed utilisation against efficiency metric with dynamic range weighting 1.5, based on The Green Grid data set



Appendix 4 Figure 18 - Deployed power at 50% fixed utilisation against efficiency metric without dynamic range weighting, based on The Green Grid data set

In conclusion, both metrics show high correlations and good scalability under a variety of utilisation levels. There is no perfect guidance but if idle time is expected to be over 15%, then DR=1.5 metric is recommended.

Performance of metric under different workloads

The efficiency metric is calculated using a 60:40 weighting of CPU:RAM based on an average server workload. However, under some applications, the workload weighting will be different and this will affect the server configurations performance and power. This section analyses the accuracy and applicability of the metric by comparing the deployed power and performance under two different workloads, CPU intensive and memory intensive.

The workloads are based on industry proposals and defined as:

- CPU intensive 85:15 CPU:memory weighting
- Memory intensive 40:60 CPU:memory weighting

Recalculating the performance and power consumption using the CPU intensive (see Appendix 4 Table 7) and memory intensive workloads (see Appendix 4 Table 8) give significantly different deployment power. Comparing this against the efficiency metric, based on the balanced metric, shows the correlations drops below to 0.90 and below in the majority of cases for the CPU intensive workload and below 0.95 for the memory intensive workload. The idle/max ratio causes the correlation to rise for CPU intensive workloads, but drop for memory intensive workloads.

Appendix 4 Table 7 - Metric correlation with deployed power using CPU intensive workload under various utilisation levels, with and without dynamic range

CPU intensive workload	DR=1.5	No DR
50% fixed utilisation	0.8764	0.8478
25% fixed utilisation	0.9140	0.8793
Sum of utilisation levels (approx. 64% average utilisation)	0.9060	0.8832

50% average utilisation (22% idle, 78% sum utilisation)	0.9397	0.8707
25% average utilisation (60% time idle, 40% sum utilisation)	0.9019	0.8243
Idle	0.7880	0.8243

Appendix 4 Table 8 - Metric correlation with deployed power using memory intensive workload under various utilisation levels, with and without dynamic range

Memory intensive workload	DR=1.5	No DR
50% fixed utilisation	0.9410	0.9623
25% fixed utilisation	0.9249	0.9447
Sum of utilisation levels (approx. 64% average utilisation)	0.9365	0.9593
50% average utilisation (22% idle, 78% sum utilisation)	0.9523	0.9677
25% average utilisation (60% time idle, 40% sum utilisation)	0.9752	0.9714
Idle	0.9493	0.9135

In conclusion, under different workloads, the efficiency metric provides general guidance. However, as this metric was developed within the context of the Ecodesign directive, which focuses at the product rather than system level, it is not fully scalable to data centre level. Therefore, for information purposes it could be recommended that where the data centre operator is aware of the specific workload, the deployment power is calculated.

HOW TO OBTAIN EU PUBLICATIONS

Free publications:

- one copy:
  via EU Bookshop (http://bookshop.europa.eu);

- more than one copy or posters/maps:
  from the European Union's representations (http://ec.europa.eu/represent_en.htm);
  from the delegations in non-EU countries
  (http://eeas.europa.eu/delegations/index_en.htm);
  by contacting the Europe Direct service (http://europa.eu/europedirect/index_en.htm)
  or calling 00 800 6 7 8 9 10 11 (freephone number from anywhere in the EU) (*).

  (*) The information given is free, as are most calls (though some operators, phone boxes or hotels may charge you).

Priced publications:

- via EU Bookshop (http://bookshop.europa.eu).

Priced subscriptions:

- via one of the sales agents of the Publications Office of the European Union
  (http://publications.europa.eu/others/agents/index_en.htm).

ET-02-16-987-EN-N



Publications Office

doi: 10.2873/782959