Universität Zürich[UZH]

Master thesis

For the attainment of a

Master of Advanced Studies in Real Estate

Data centre design standards and best practices for public research high performance computing centres

Author: Gilly
Ladina
Piazzora 3, 6805 Mezzovico

Submitted to: Dr. T. Guthknecht

Date of submission: 29.08.2016

Table of contents

List of abbreviations

ANSI	American National Standards Institute
ASHRAE	"American Society of Heating, Refrigerating and Air-Conditioning Engineers. Founded in 1894, ASHRAE is a global society advancing human well-being through sustainable technology for the built environment. The Society and its members focus on building systems, energy efficiency, indoor air quality, refrigeration and sustainability within the industry. Through research, standards writing, publishing and continuing education, ASHRAE shapes tomorrow's built environment today."[1]
ATM	automated teller machine
BICSI	Building Industry Consulting Service International, is the worldwide association for cabling design and installation professionals.[2]
CAGR	Compound Annual Growth Rate
CENELEC	European Committee for Electrotechnical Standardization
CMOS	Complementary metal–oxide–semiconductor (CMOS) is a technology for constructing integrated circuits used in current computers processors.
CRAC	Computer Room Air Conditioner
DCiE	Data Center infrastructure Efficiency. An energy efficiency metric proposed by The Green Grid. Reciprocal of PUE metric.
DOE	U.S. Department of Energy
FLOPS	Floating-point operations per second. Measure of computer performance.
HPC	High Performance Computer / Computing
HVAC	Heating Ventilation and Air Conditioning
ICT	Information and communications technology.
IDC	International Data Corporation
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

[1] https://www.ashrae.org/about-ashrae

[2] https://www.bicsi.org/double.aspx?l=1102&r=1104

kW	Kilowatt; power unit that corresponds to one thousand Watts.
LEED	Leadership in Energy and Environmental Design. Green building rating system.
MW	Megawatt; power unit that corresponds to 1 million Watts.
N + 1	Need plus One. System design with one unit in addition to requirement.
NITDR	The Networking and Information Technology Research and Development Program is the U.S.'s primary source of federally funded work on advanced information technologies in computing, networking, and software.[3]
PRACE RI	Partnership for Advanced Computing in Europe – Research Infrastructure. Also referred to as PRACE.
PUE	Power Usage Effectiveness. An energy efficiency metric proposed by The Green Grid. Reciprocal of DCiE.
RDHX	Rear-door heat exchanger
ROI	Return on investment
SaaS	Software as a Service
TCO	Total Cost of Ownership
TCP/IP	Transmission Control Protocol/Internet Protocol is the basic communication language or protocol of the Internet.
TIA	Telecommunications Industry Association
UPS	Uninterruptible Power Supply
USGBC	U.S. Green Building Council

[3] https://www.nitrd.gov

Glossary

Air-cooled	IT equipment where component are cooled by convection of air within the rack.
Cabinet	Frame for housing IT equipment. Also know as “rack”
Cloud	Also referred to as Cloud Computing. Describes the delivery of on-demand computing resources over the Internet.
Datacom	Used as intended by ASHRAE Technical Committee 9.9 as an umbrella term for the combined telecommunications and data centre industries
Exaflops	Compute performance of 10^{18} floating point operations per second.
Free cooling	An approach to lowering the air temperature in a building or data centre by using naturally cool air or water instead of mechanical refrigeration.
Liquid-cooled	IT equipment that is cooled using a liquid, such as water, refrigerant or dielectric, within the design control of the IT manufacturer.
Petaflops	Compute performance of 10^{15} floating point operations per second.
Rack	Frame for housing IT equipment. Also known as “cabinet”.
Teraflops	Compute performance of 10^{12} floating point operations per second.
The Green Grid	Global consortium of end-users, technology providers and government agencies intent on improving energy efficiency in data centres and enterprise computing ecosystems around the world.[4]
Tier level	Hierarchical classification of data centre infrastructures based on increasing levels of redundant capacity components and distribution paths.
Uptime Institute	“Uptime Institute is an unbiased advisory organization focused on improving the performance, efficiency, and reliability of business critical infrastructure through innovation, collaboration, and independent certifications.”[5]

[4] http://www.thegreengrid.org/en/about-the-green-grid.aspx
[5] https://uptimeinstitute.com/about-ui

List of tables

List of figures

Executive Summary

Public research high performance computing (HPC) centres form a key part of the infrastructure backbone enabling scientific discovery, driving innovation and ensuring the competitive advantage and security of nations. A cost effective, flexible and energy efficient data centre design will safeguard the longevity of capital expenditure and contain operational costs, thereby securing the majority of public funds allocated to these national endeavours for the purpose research.

By virtue of their mission HPC centres find themselves at the forefront of computing development, and as such, are thus the first to experience the disruptive changes in technology and requirements brought about by the pursuit of ever-greater compute performance. Their use case differs markedly from that of enterprise data centres that existing standards cater to.

Defining design criteria for a data centre with a life expectancy of several decades, which must be able to accommodate multiple generations of HPC systems is, therefore, a complex task that requires a combination of engineering expertise, great curiosity and educated guesses about future IT technology developments.

This work investigates to what extent existing standards are known and used by public research HPC data centres to guide their design decisions. It also provides an overview of existing standards and analyses those areas where they do not cover HPC requirements. Through interviews with HPC sites from three continents, design challenges and future-proofing strategies were collected, compared and finally the best practices discussed were compiled. The resulting document provides future managers of HPC data centre projects with a starting point for their design. The compilation of best practices should form a basis for the community on which to extend and build a shared body of knowledge and expertise.

1 Introduction

1.1 Motivation

This research stems from the author's personal experience in the building of the Swiss National Supercomputing Centre (CSCS, Lugano, Switzerland) between 2008 and 2012. During the course of planning, it became evident that the existing and expected requirements of high-performance computing (HPC) systems did not fit within the data centre design standards at the time. Combined with the expectations from the funding body, budget and time constraints, rapidly changing technologies and requirements as well as the comparatively short life cycles of HPC systems, the project appeared as a daunting task. In order to determine key design criteria for the project, the following was undertaken: 1) a detailed analysis of our business case, 2) a review of expected changes in HPC system requirements, and 3) site visits of selected public research HPC centres in Europe and the United States.

Comparisons with peer sites proved invaluable in that it showed the challenges CSCS was facing in terms of requirements were not unique within the HPC community. It also allowed us to integrate excellent ideas, avoid pitfalls in design, and gain insight into best practices at the time. As a result of CSCS's experience, this research is intended to offer a starting point for managers in charge of future construction projects of this type.

1.2 Problem definition

Public research HPC data centres are government funded national endeavours, and subject to public procurement laws. As such, the majority of public funds allocated must go toward enabling research, rather than funding buildings and infrastructure. Care must therefore be taken at the design stage to ensure the longevity of capital expenditure and optimised operating costs over the lifetime of the facility. The less funds spent on these items the more is available for the computer systems they house, thus, maximising the money spent on enabling the public-good service of these national endeavours.

Funding agencies, expect public research HPC data centres to achieve the usual lifetime of a building or infrastructure. This goal contrasts starkly with the expected three to five year lifecycle of supercomputers that is driven by the need to provide ever-greater compute performance to enable users to tackle complex questions of scientific and societal importance. The history of computing has seen a number of disruptive technological changes that have affected the requirements these systems make on the infrastructure and buildings that host them.

Being at the forefront of computing development, HPC is the first area to experience and have to negotiate these disruptive changes. Past changes have included, but are not limited to, the transition from vacuum tube based systems in the 1940s to solid state transistor in the 1960s, as well as the move from liquid to air-cooling in the 80s and the subsequent return of liquid-cooling in the first decade of this century. Future challenges lie in the pursuit of exascale[6] performance and the impending dusk of Complementary Metal-Oxide Semiconductor (CMOS) technology.

Defining a design criteria for a data centre against this backdrop is, therefore, a complex task. Whilst enterprise and HPC data centres differ in terms of their purpose and therefore also with regard to some of their design drivers, their data centres are, in both use cases, support infrastructures as opposed to the co-location industry, where the data centre *is* the revenue generator. A cost effective, flexible and energy efficient design will reduce the data centre costs and allow the enterprise or public research institution to focus their funds on their core business.

1.3 **Goal of the study**

The goal of this study is:

1) to provide an overview of current relevant standards and investigate where the key design criteria of public research HPC data centres differ from those found in enterprise. 2) This thesis examines what factors drive these differences. Interviews with representatives from leading public HPC data centres were conducted, whereby the limitations of existing standards and their applicability to current design criteria were investigated.

It is intended that the results of this study will provide future managers of construction projects with an overview of existing design standards and a compilation of best practices utilised by HPC data centres today.

---

[6] Appendix 1 provides an overview computer performances by magnitude

1.4 Hypothesis and research questions

This thesis grounded on the following hypotheses and research questions:

Hypothesis I:	“Current data centre design standards do not reflect the requirements of public research HPC data centres.”
Research question 1:	“To what extent are the existing data centre design standards known and applied within public research HPC data centres?”
Research question 2:	“Which design issues, common to public research HPC data centres, do the standards not cover?”
Research question 3:	“How do public research HPC data centres define design criteria where their requirements are not covered by the existing standards?”
Research question 4:	“What approaches to future proofing do public research HPC data centres apply?”
Hypothesis II:	“Public research HPC data centres share a common set of key design requirements that allow the definition of a design standard for this type of data centre.”
Research question 5:	“Is there sufficient similarity in the applied approaches to allow the definition of a standard for the design of public research HPC data centres?”

1.5 Delimitation of the topic

This study is aimed at management level decisions regarding key design criteria of public research data centres that will lay the corner stone for a construction project. It does not go into in-depth technical design details, nor does it discuss the IT side of the business. It is important to note that each centre is unique, built within different boundary conditions and, as such, this study cannot give deterministic solutions. Rather, it aims to discover and define common criteria and approaches to help guide future public research HPC data centres.

1.6 Research process and methodology

The literature review in chapter 2 provides an introduction to data centres, their historical development and economic importance. Existing data centre design standards are discussed and typical design drivers of enterprise data centres reviewed in order to provide a comparative basis for the design requirements encountered in public research HPC data centres. An overview of the history of supercomputing and its importance is also provided as well as how developments in HPC have affected the wider data centre industry. Also outlined are existing forums where the HPC community can exchange ideas and best practices.

Given the fast-moving and digital nature of this environment, the literature review comprises a number of citations from electronic industry publications, publicly available research and government reports, in addition to standards and books.

In order to investigate the hypotheses and answer the research questions, structured interviews were conducted with representatives from leading public research HPC data centres from the Department of Energy of the United States, the PRACE Research Infrastructure in Europe as well as selected sites from Asia and Australasia. A qualitative analysis of these interviews allows us to test the veracity of the hypotheses and related research questions.

1.7 Use of British and American spelling

This document is edited in British English. However, it contains a number of references to U.S. publications for which the original U.S. English spelling is used. This is particularly evident in the word data centre that is written data center or datacenter in U.S. English.

2 Literature review

2.1 An introduction to data centres

2.1.1 *Definition and historical background*

According to Webopedia "Data centers are physical or virtual infrastructures used by enterprises to house computer, server and networking systems and components for the company's information technology (IT) needs, which typically involve storing, processing and serving large amounts of mission-critical data to clients in a client/server architecture."[7] They have become the nerve centres of our digital economy.[8]

[7] http://www.webopedia.com/TERM/D/data-center.html

[8] ANIXTER 2007, p. 2

The term “data centre” was first used the 1960s, when large mainframe systems were housed in separate rooms due to a) the substantial amounts of space and controlled environmental conditions required and b) the need to keep them secure.

In the 1980s air-cooled microcomputers emerged, alongside the large mainframe systems, allowing computers to enter offices and be installed anywhere. The Internet boom in the 1990s, combined with growing complexity of the IT environment, led to the new generation servers once again being installed in dedicated computer rooms. During this boom, the concept of a data centre was perfected and began to proliferate.[9,10]

With the exception of a brief collapse immediately following the burst of the dotcom bubble, the growth of data centres initiated by the Internet boom continues to this day.[11]

2.1.2 *Convergence of telecommunications and data processing*

Data centres, as we know them today, were brought about by the convergence of the telecommunications and computer industries. Until the 1970s telecommunication had, for most of its history, ensured the transmission of communication that relied on wired networks, whilst the computers that emerged in the 1940s and 1950s were primarily used to store information and execute complex computations. The invention of the Internet Protocol (TCP/IP) in the 1970s connected these two worlds, spurring explosive growth that led to the dotcom bubble at the end of the last millennium and forming the basis for today’s digital economy. Telecommunications became wireless and the decreasing cost of compute power made it accessible to the masses, thus further giving rise to the need for compute and storage facilities.[12,13]

2.1.3 *Different types of data centre*

As data centres have progressively been integrated into various industries they have had to adapt to different business cases. While sources are divided on the differences and commonalities of data centres and resulting categories, for ease of reference the

---

[9] Bartels, A. 2011, no page reference
[10] Nutt, A. 2008, no page reference
[11] Donnelly, C. 2016, no page reference
[12] Beaty, D. L. 2013, p. 74 - 78
[13] Ceruzzi 2003, p. xi and p. 1 - 21

work of M. Nispel, *Extreme Networks*[14], which is based on reports by market analysts Dell'Oro, is utilised in this thesis.[15] According to Nispel (2013), the market is divided into the following segments:

- Web 2.0 Software as a Service (SaaS) data centres (e.g. Google, Facebook)
- Service provider data centres, including
    - cloud data centres (e.g. Amazon)
    - co-location data centres (e.g. Interxion)
- Enterprise data centres
- Niche markets such as:
    - High frequency trading (HFT) data centres
    - Research data centres for high performance computing (HPC)

	Web 2.0 SaaS	Service Provider	Enterprise	HPC[16]
Degree of standardization	Very high	Medium	Low	Medium
Number of applications supported	Very low	Very high	Medium	High
Focus on cost	Very high	High	Medium	Medium
Focus on value to the business	Very high	High	Very high	Very high
Workload scale	Very high	Very high	Low to medium	Very high
Multitenancy	Very high	Very high	Very low	Very low
Physical scale	Very large	Very large	Small	Medium

Table 1: Comparison of the three main data centre market segments based on Nispel (2013)[17]

The table above is a graphical translation of Nispel's comparison of the first three segments, based on common attributes. As he does not discuss the two niche markets in his comparison, due to the lack of commonalities they have with the three main segments, the "HPC" column has been added by the author, based on her experience.

2.2 The private data centre industry

As the digital economy, the Internet of Things (IoT) and online entertainment spread, the amount of data traffic generated by end users and businesses continues to increase,

[14] Nispel 2013, no page reference
[15] http://www.delloro.com
[16] Based on the author's experience.
[17] Nispel 2013, no page reference

making data centres increasingly important to the livelihood of a business and driving the need for data centre space and services.[18,19] Most data centres around the world are built and operated by private industry. Web 2.0 data centres, for example, provide software as a service to paying customers. Service provider data centres specialise in hosting outsourced enterprise compute services with co-location facilities hosting servers and services from multiple enterprises. Enterprise or corporate data centres are owned and operated by the company that uses the compute infrastructure and services housed therein.[20] The origin and main target audience for existing data centre design standards stem from this last category; they have a different set of motivations and challenges than publicly funded HPC data centre initiatives. In this chapter we analyse how these private facilities form the backbone of the digital economy and are fuelling real estate investment, as well as challenges that the industry is experiencing with the existing design standards. Use case and design drivers for enterprise data centres are also examined.

2.2.1 *An engine for the digital economy that is fuelling real estate investment*

Whilst for businesses, data centres represent an increasing cost item, the increasing need for facilities of this kind has been a boon for the construction sector. Although investors shunned the industry following the dotcom bubble, the high-sustained growth since the market recovery in 2007-2008 and the industry's proven ability to weather recessions, as well as healthy investment returns, are attracting them back.[21] In the following we take a look at these key infrastructures from a business cost perspective as well as from an investor perspective.

Market research company International Data Corporation (IDC), cited in the New York times, estimates that there are over three million data centres worldwide supporting the digital economy. Koomey, who co-authored the widely known U.S Environmental Protection Agency (EPA) Report to Congress on Server and Data Centre Energy Efficiency in 2007 and went on to author the follow-up report in 2011 commissioned by the New York Times, estimates data centres in the U.S. alone consumed the equivalent of two per cent of the total electricity usage of the U.S. in 2010. [22,23,24]

---

[18] Restivo 2015, slide nr. 7
[19] Donnelly 2016, no page reference
[20] Miller 2012, no page reference
[21] Ibid
[22] Glanz 2012, no page reference

Whilst data centres have taken on an increasingly central role, an industry review by McKinsey in 2008 found their cost was also increasing to the point of having an adverse impact on profitability. Efforts to tackle inefficiencies and reduce costs were being hampered by silo-mentality decision-making and fragmentation of data centre cost reporting. To address this issue McKinsey introduced the Corporate Average Datacenter Efficiency (CADE) measure[25] in collaboration with the Uptime Institute in 2008. [26]

During the banking crisis in 2008 many enterprises moved away from building their own data centres and toward outsourced solutions to meet capacity needs. Slow-downs in GDP growth around the world are putting pressure on enterprises profits and cyber security has become paramount, whilst the overall complexity of running a data centre has increased, driving businesses to further consolidate.[27] This is expected to lead to a shift in data centre investment from enterprise space toward outsourcing or colocation space leading to the average size of enterprise data centres shrinking, whilst third-party providers are increasing their average footprint. [28,29]

The DatacentreDynamics Intelligence Report on data centre evolution from October 2015 put expected data centre investment growth for the year at $184.4 billion, which represents a 10.2% increase on 2014. The study also expects to see data centre investment increase by a Compound Annual Growth Rate (CAGR) of 9.2% between 2014 and 2020 for an equivalent of $283.4 billion. The main construction activity is occurring in new construction projects and consolidation projects, which has led to the number of data centres worldwide tripling between 2007 and 2014.[30]

2.2.2 *Use case and design drivers for enterprise data centres*

In order to be able to compare the requirements of HPC to those of the enterprise data centre, this chapter analyses the use case of enterprise data centres and how this impacts their design drivers.

---

23 Koomey 2011
24 U.S. Environmental Protection Agency 2007
25 Please refer to appendix 2 for compute formula for the CADE and further data centre metrics.
26 Kaplan/ Forrest/ Kindler 2008
27 Jones Lang LaSalle 2016, p. 1 - 2
28 Uptime Institute 2014, no page reference
29 Jones Lang LaSalle 2011, p. 1 - 3
30 Donnelly 2016, no page reference

In today's digital economy, data centres are the place where an enterprise processes its business transactions, hosts their central IT services such as e-mail, financial records and websites, as well as processing and safeguarding the company's intellectual property. An enterprise data centre can thus be likened to the brain of a company. It is relied upon to allow the company to connect and communicate with the rest of the world, store information, power research and development and support central administrative processes. The availability– also known as uptime - of this "brain" is therefore key to a company's success. Depending on what service a company offers, availability can be more or less critical. For a data centre that hosts servers for the purpose of processing Automatic Teller Machine (ATM) transactions any downtime will result in significant financial losses. An enterprise that relies on their data centre to support their daily business processes but not their transactions may suffer loss of productivity but incur less damage than the ATM transaction processing site. As higher availability of a data centre corresponds to higher capital expenditure and operating costs, the correct matching between data centre availability and business case is of paramount importance. [31]

According to the Uptime Institute, a data centre's uptime performance is a result of its design topology, its robustness and operability along with factors such as site selection, construction implementation management and staffing. The higher the level of uptime performance required, the more robust a data centre will need to be. In order to achieve this robustness, the redundancy of components and distribution paths will need to be increased. The Uptime Institute data centre classification system defines four hierarchical levels of robustness that are referred to as Tiers. Higher availability may also be reached by using multiple geographically distributed data centres.[32]

A comparison of Tier levels, that will be further discussed in chapter 0, and their respective costs conducted by Anixter shows that designing to the next higher Tier level can add fifty per cent to the construction cost of a data centre. The steepest increment in price and performance benefit is to be found between Tier levels II and III. Anixter further notes that Tier IV (highest level of robustness) data centres are very rare today, as their expense is hard to justify.[33]

---

[31] Alger 2005, p. 5 – 11, and 22 - 25
[32] Turner et al. 2008, p. 2 - 12
[33] Anixter 2007, p. 21

The end user's expectation that digital services be instantaneously available at all times puts huge pressure on data centres to avoid downtime at all costs. This has led to data centre operators overprovisioning both building infrastructure and IT hardware, subsequently leading to very low server utilisation rates of 6 – 12% and a fear of disruptions that makes the industry very risk averse when it comes to adopting new technologies aimed at improving usage and energy efficiency.[34] The Uptime Institute industry survey in 2014 shows that the industry made good progress in improving their Power Usage Efficiencies (PUE)[35] between 2007 and 2011 but subsequent gains have been smaller and increasingly expensive. Efforts so far have been focused mostly on improving building infrastructure efficiencies. At this point it is recognised that further improvement can only be achieved if the issue of poor IT utilisation is tackled. Currently the incentive for this to happen is low as most companies still allocate the power bill for the data centre to the facilities or real estate budgets rather than the IT departments. However, the latest Uptime Institute survey shows that large organisations have been more aggressive in pursuing energy efficiency, adopting new technologies and out-sourcing less critical loads. This is especially true for the financial services companies who's profitability is directly correlated with their ability to manage their IT infrastructure efficiently.[36]

Traditionally hardware racks[37] in enterprise data centres have power densities of 4 – 5 kW and data centre infrastructures are built to accommodate an average density of around 1 – 2 kW/m^2 (100 – 200 watts/sq. ft.) as up until now it has been cheaper to keep densities low at the expense of a larger footprint.[38]

2.2.3 *Data centre industry challenges with current Tier classification*

Although existing data centre standards are globally recognised and applied within private industry, they are recently beginning to come under pressure due to their inability to accommodate innovation with novel technologies or emerging areas of interest, such as energy efficiency. Latest figures suggest progress on the energy efficiency front, but data centres face strong criticism as large power consumers. Innovative data

---

[34] Glanz 2012, p. no page reference
[35] The PUE expresses the ratio of total energy used by a computer data centre facility to the energy delivered to computing equipment. Please refer to appendix 2 for the compute formula for PUE and it's reciprocal DCiE metric.
[36] Uptime Institute 2014, no page reference
[37] Frame used to house IT equipment.
[38] Mitchell, R. L. 2010, no page reference

centres see potential in further improving the record of sustainability in the industry by incorporating more alternative energy sources, but come into conflict with existing regulations. For example, these innovative designs do not fit within the existing four Tier levels that the current standards are built on. In essence, innovation may potentially be punished. A data centre design may reflect societal and political aspirations of energy efficiency, but fall short in terms of certification and accreditation within the requirements of the standards. This is an issue for businesses, especially in heavily regulated industries where Tier certification is a requirement.

Large players such as Google and Facebook's Open Compute Project[39] and The Green Grid[40], are advocating the need for innovation in data centre technology. A discussion has recently emerged in the industry requesting the Tier classification system be reviewed to take technology innovations of the past two decades into account, making it possible to reward those who pursue sustainability. The new classification model proposed for discussion would allocate a resilience score for each of the following five areas: energy source, electrical system, mechanical system, network topology and IT. The total design classification would be comprised of the overall resiliency score, the sustainability score (graded A – F) and the efficiency score based on PUE.[41,42]

2.3 **Public research HPC**

Although significantly less in number than those in the private sector, public research HPC centres form a key part of the infrastructure backbone enabling scientific discovery. Discussed in this chapter are the definition, history and evolution of high-performance computing, as well as its role in driving innovation and its growing importance in securing not only competitive advantage but also the security of nations. The impact of technological changes of supercomputers is examined, as are the trickle -down effects from this domain to the enterprise data centre segment. To round off the discussion we look at the discussion forums that have formed in recent years around the topic of HPC data centres.

---

[39] http://www.opencompute.org/about/
[40] http://www.thegreengrid.org
[41] Bilderbeek/ Coors 2016, p. 2 - 5
[42] Coors 2016, no page reference

2.3.1 *Definition and historical background of High Performance Computing*

The term “high performance computer” or “supercomputer” is attributed to the fastest computers at any given moment in time. These systems are used to pursue challenging scientific problems in computational science that could otherwise not be tackled.[43] The speed of these computers is measured by the number of floating point operations per second (FLOPS) they are able to perform.[44] Mannheim Supercomputer Statistics published information about supercomputers installed around the world based on information provided by manufacturers from 1986 to 1993.

This statistic was replaced in 1993 by the Top500 list,[45] which ranks commercially available supercomputers worldwide twice a year based on the Linpack benchmark.[46] This list is self-reported and does not contain all systems, as some operators elect not to list their systems. Figure 2 above shows the evolution of computer performance si nce the introduction of the ranking as well as a projection until 2020.

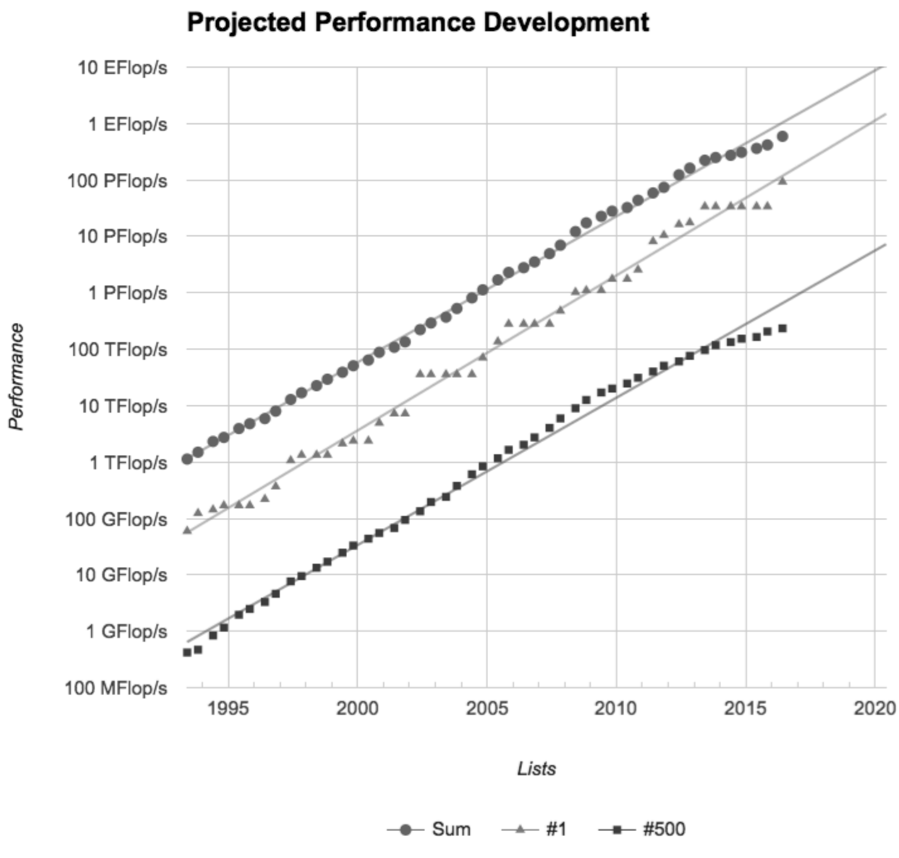

Figure 1: Speed of the World’s fastest and 500th-fastest Supercomputer, by year, 1993 – 2015, top500.org

[43] Vetter 2013, p. 3 - 11
[44] Kaufmann/ Smarr 1993, p. 32 - 34
[45] https://www.top500.org/project/
[46] https://www.top500.org/project/linpack/

Since 2007 it has also been possible to enter supercomputers on the Green500 list, which ranks systems based on energy efficiency. This latter list started as a reaction to increasing awareness surrounding energy consumption by supercomputers.[47,48] The genesis of modern supercomputing can be found in World War II, which led to massive government funding to develop existing ideas for mechanical brains into usable systems in order to support the war effort. Results of these endeavours can be found in the Colossus machine at Bletchley Park (UK) and the Manhattan Project in the U.S. After the war, advances made in computing technology were transferred to peacetime uses. For example, whilst the *Electronic Numerical Integrator and Computer (ENIAC)*, installed at the University of Pennsylvania in 1946, was built for military and scientific purposes, the *UNIVAC* (1950) was the first commercial computer that could be programmed for different applications. Government agencies were amongst the first customers for this system and its competitor, the *IBM 701*.[49]

During the 1950s and 1960s, mainframe computers were adopted by the business world and used for electronic data processing. Parallel to this, Control Data Corporation (CDC) entered the market in 1957, pursuing the scientific segment. Sources vary whether the term supercomputer was first used for the *CDC 6600* of 1964 or the *Cray 1* of 1976. Both computers were, however, designed by Seymour Cray, who founded his own company in 1972.[50,51]

Until the 1980s, computers had consisted of a small number of processors, but the 1990s saw the advent of supercomputers with thousands of processors. Computing, heretofore largely dominated by the U.S., began to spread around the globe. Three Japanese computer manufacturers, for example, entered the market during the 1990s, sparking a global competition. The *Intel ASCI Red* supercomputer at Sandia National Laboratories (SNL) was the first computer to exceed one teraflop of performance in 1996 and ten years later the *IBM Roadrunner* supercomputer, installed at Los Alamos National Laboratory (LANL), inaugurated the era of petaflop computing. In 2004 NEC's *Earth Simulator* was the first Japanese system to top the Top500 list. The 21st century has seen China make a first appearance amongst the 500 fastest computers in the world in 2003, from where it has climbed steadily to claim the top position in No-

---

[47] www.green500.org
[48] www.top500.org
[49] Ceruzzi 2003, p. 35 – 37, 46 – 53
[50] Ibid, p. 54 - 64
[51] Murray 1997

vember 2010 with the *Tianhe-1A* system and again in June 2016 with the *Sunway* system.[52, 53]

The introduction of supercomputers gave scientists a new tool for discovery, establishing simulation as the third pillar of science alongside theory and experiment. The sheer speed of supercomputers enabled scientists to create simulations of the natural world and create graphic renderings of the results, thereby extending the boundaries of research from the smallest known parts that make up our world to the discovery of the universe. Simulation also enables examination of problems that would otherwise be beyond reach. For example, the introduction of simulation has allowed great advances in areas such as climate research, chemical discovery, product design, engineering, pharmaceuticals and even financial services.[54]

2.3.2 *An engine for research and innovation*

Since World War II U.S. Department of Energy (DOE) laboratories have continued to push the boundaries of discovery and engineering to solve challenges of societal and scientific importance. Today, scientific discovery, commercial innovation and national security rely on HPC, making this technology central to a nation's ability to compete globally. Supercomputing is key to the continued exploration of existing energy sources and facilitating the discovery of new ones. It enables scientists to better understand our climate and predict severe events, renders cars and airplanes safer and more energy efficient, contributes to the advancement of medicine, expedites the development of consumer products, and has revolutionised the financial services and entertainment industry. The importance of this technology led the U.S. to formulate the "High-Performance Computing Act"[55,56] in 1991 with the aim to ensure the U.S. maintains its leadership in high-performance computing. Federal funding for this effort is coordinated through the Networking and Information Technology Research and Development (NITDR) Program.[57] Between the years 2000 and 2015 the annual expenditure of this programme grew from $1500 million to $4378 million and funding re-

---

[52] Vetter 2013, p. 4 - 5
[53] Fosdick 1996, p. 1 - 27
[54] Kaufmann/ Smarr 1993, p. x - xi
[55] https://www.nitrd.gov/congressional/laws/102-194.pdf
[56] United States Congress House Committee on Science 2003, p. 7
[57] https://www.nitrd.gov

quests for 2017 total $4542 million.[58] By purchasing the highest-end computers and making them available to a wide range of researchers, the U.S. government indirectly supports research in the computer industry. More recently, other nations have realised the importance of HPC in ensuring global competitiveness and have joined the race to achieve exascale performance. This increased competition saw the U.S. launch of National Strategic Computing Initiative[59] in 2015 aimed at maintaining U.S. leadership in the development and application of supercomputers.[60,61]

The co-existence of computer manufacturers and high-end scientific users in the U.S. creates a self-enforcing eco-system. Scientists, for example, formulate challenges in their fields of study and the computer industry develops systems to support this work. This approach enables scientific discovery and technological development, thereby giving science and businesses the competitive edge and first-mover advantage. In 2015, the broader computer-manufacturing sector of the U.S. employed approximately one million people with average salaries two and a half times higher than the U.S. national average.[62] The sector also generates a net trade surplus due to the fact that most supercomputers are built with processors produced by U.S. companies.

The ability to perform advanced modelling, simulations and data analytics has made HPC indispensable for advanced manufacturing businesses. It reduces cost and accelerates the speed of research and development by diminishing or even eliminating the need for prototyping and testing. According to the market research company IDC, every dollar invested in HPC generates $515 in revenue and $43 cost savings or increased revenue in the U.S..[63] A 2014 study by the U.S. Council on Competitiveness found seventy-six per cent of enterprises believe that “Increasing performance of computational models is a matter of competitive survival”.[64] In Europe, IDC found every Euro invested in HPC generates a revenue increase of €876 and €69 of additional profits.[65] Consequently, making HPC accessible and affordable to a broad base of businesses is

---

[58] NITRD Supplement to the President’s FY 2017 Budget 2016, p. 6 - 13
[59] White House 2015, p. no page reference
[60] Joseph/ Dekate/ Conway 2014, p. 1 - 6
[61] See Bartels 2011, no page reference
[62] U.S Bureau of Labor Statistics 2015, as quoted in Ezell/ Atkinson 2016, p 13
[63] Joseph/ Conway/ Sorensen 2014, slide 2, 12
[64] The Council on Competitiveness 2014, p 18
[65] European Commission 2015, p. 4

one of the key challenges to ensuring that maximum benefits can be reaped from this technology.[66]

Currently, U.S. companies generate approximately two thirds of global HPC revenues. Sixty-nine out of the first one hundred machines listed in the November 2015 Top500 list were produced by American companies. 40.9% of HPC systems were deployed in the U.S. with a further 33.9% going to the EMEA region, 18.2% to Asia-Pacific and 6% to Japan.[67]

In 2004 Japan topped the list of the fastest computers in the world with the NEC *Earth Simulator*. This feat was repeated in June 2012 with the *K* Computer. In a bid to regain its leadership position Japan has now launched its *Flagship 2020* programme, which will inject $1 billion of investment, with the aim of developing and delivering a system one hundred times more powerful than the current *K* computer. The system is expected to help tackle challenges in health, environment, energy, industry and science through a co-design effort in developing both hardware and software.[68]

Similarly, in 2012 the European Union articulated a plan to achieve exascale performance within the same timeframe as the U.S., Japan and China. To this end, it has increased funding for research and development systems in HPC and launched the European Technology Platform on High-Performance Computing (ETP4HPC). This €700 million investment is augmented by €400 million of mostly in-kind contributions via the PRACE consortium.[69] Despite the advances in funding and declaration of goals, the formulation of a coordinated and cohesive plan of how to achieve it has yet to emerge. IDC estimates a further € 1 billion in funding would be required.[70] The recently published Horizon 2020[71] funding programme announced the creation of a consortium of hardware and software developers aiming to build an exascale prototype by 2018. It is hoped this will increase the development of indigenous HPC technologies.[72]

HPC leadership is a national priority in China with a focus on increased self-production. India, South Korea and Russia have been identified by IDC as the other

---

[66] The Council on Competitiveness 2015, p. 4
[67] Ezell/ Atkinson 2016, p. 30 - 34
[68] http://www.aics.riken.jp/fs2020p/en/
[69] PRACE 2016, p. 27 - 30
[70] European Commission 2015, p. 4
[71] http://www.exanest.eu
[72] Saarinen 2016, no page reference

main competitors in the exascale race. The table in appendix 3 provides an overview of investments in national HPC strategies for each of the above named countries.[73]

2.3.3 *Evolution of HPC and impact on data centre infrastructure*

The short life-cycle of supercomputers combined with rapid changes in density and cooling technology test the flexibility and adaptability of the building infrastructures hosting them to the limit. In the following, the impact of changing transistor technology and density on cooling technology is illustrated as an example. The earliest computers were primarily air-cooled and their equipment filled entire rooms. With increasing compute power came the need to control the environmental conditions of these rooms to keep the systems cool. IBM began liquid-cooling their mainframe computers as early as 1964 in order to take advantage of the vastly superior thermal performance of liquid in removing heat from the increasingly dense IT equipment, and the vast majority of large-system servers continued to use liquid cooling into the 1980s. As the figure below illustrates, the move from bipolar processors to complementary metal oxide semiconductor (CMOS) processors in the early 1990s drastically reduced the

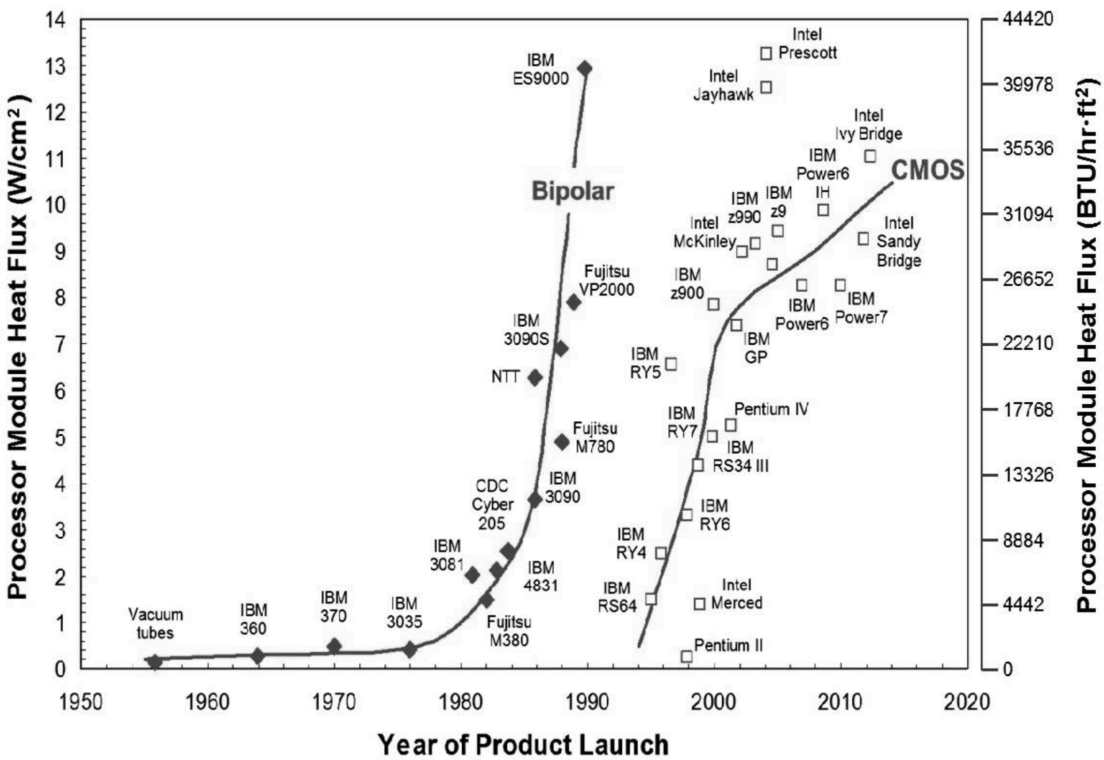

Figure 2: Evolution of processor module level heat flux in high-end servers. ASHRAE 2016, p. 27

[73] Ezell/ Atkinson 2016, p. 22 and 39 - 41

heat density of these systems, which once again made air-cooling the most cost-effective cooling method.

However, the continuing increase of power and packaging density since the advent of CMOS technology has once again made the efficient removal of heat from IT equipment a critical issue. 2008 saw IBM introduce the first large liquid-cooled server based on CMOS technology and numerous other vendors have since followed suit for their high-performance systems.[74]

Whilst in the 1970s there was an urgent demand for both processing power and infrastructure capacities to develop, the advent of CMOS and lower densities took the pressure off the infrastructure until the turn of the century. Today infrastructure capacities are once again reaching their limits due to the high power densities of the processors.[75]

The International Roadmap for Semiconductors (ITRS) published in 2015 anticipates that transistor scaling will probably become economically unviable after 2021 as the physical properties of the materials used make a further reduction of transistor dimensions practically impossible. The advent of 3D scaling is likely to enable continued increases in semiconductor performance until 2030. Currently, one of the main challenges is the requirement to increase performance at constant power density, due to thermal constraints. The U.S. aims to achieve exascale performance with a system that consumes less than 20MW of power. As a comparison, the current ninety-three petaflops[76] number one *Sunway* system from China registered an average power consumption of 15.37 MW during the Linpack benchmark test. Hosting a 20MW system will require significant upgrades to the facilities that will host them.[77,78,79]

At the same time, supercomputing is preparing for a future beyond CMOS and discussions revolving around new technologies such as nanoscale devices, quantum computing, superconducting computing and neuromorphic computing are on-going. What challenges these new technologies will pose for existing data centre infrastructures is yet unknown.[80]

---

[74] ASHRAE 2016, p. 25 - 28
[75] Northwest University 2012, slides 2 - 3
[76] equivalent to 0.093 exaflops
[77] ITRS 2015, 31 - 32
[78] Trader 2016, no page reference
[79] Bernhardt 2010, p. 3
[80] http://beyondcmos.ornl.gov

2.3.4 *Impact of HPC developments on enterprise data centres*

In the private industry, recent data centres have been designed to support 1 – 2 kW/m^2 (100 – 200 watts/sq. ft.). Experts estimate that with careful optimisations, these infrastructures should be able to support racks of up to 25kW. This is important, as IT vendors have recently been promoting high-density racks to market segments other than traditional HPC in an attempt to increase their potential customer base. With most service providers in the industry charging for data centre usage based on space, there is an incentive for customers to move toward higher densities. Rising costs for electromechanical plants are, however, prompting operators to start charging customers based on power use instead. Whilst some operators see both a strategic and competitive advantage in using higher densities, others question their economic sense. Going beyond power densities of 25kW per rack would require the retrofit of liquid cooling capacities alongside the traditional air-cooling systems. Although doing so can significantly decrease the cost of air cooling and add the potential to exploit free-cooling technologies, it requires substantial adaptation of the building infrastructure as well as adding more weight on the raised floor.[81] A study of the total cost of ownership (TCO) of liquid-cooling found that the economic case for retrofitting an existing site to liquid-cooling is positive for sites that are energy inefficient or have high electricity costs. Higher server utilisation rates can also improve ROI. In comparison, building a new data centre with liquid-cooling has almost immediate payback.[82,83]

In addition to a change in cooling technology, a move toward higher densities in enterprise data centres would also require an increase to 7,5 – 17,5kW per m^2 (700 – 1650 watt per sq. ft.) along with a move to 3-phase 480-volt power distribution to the racks and much higher breaker ratings than the commonly used 16 and 32 Ampere (A). Not all HPC racks conform to the same standard form factor used in all data centres and typically weigh significantly more than these due to their dense packaging.[84]

Although the debate as to whether high-density and liquid-cooled computing will become mainstream is still on-going, understanding HPC requirements will allow enterprise data centres to anticipate the requirements that the future may bring. The crux lies in trying to future-proof a data centre design to mitigate the risk of being unable to

---

[81] Mitchell 2010, no page reference
[82] Demetriou/ Kamath/ Mahaney 2015, p. 5 - 7
[83] The Data Center Journal 2016, no page reference
[84] McCarthy 2011, no page reference

adapt to high-density, without having to execute expensive retrofits and at the same time avoid overprovisioning. One proposed strategy to achieve this is to build a mixed-density data centre.[85]

2.3.5 *Forums that discuss best practices in HPC*

Whilst the private data centre industry has had dedicated discussion forums since the advent of standards, exchanges between public research HPC centres on data centre design and infrastructure issues have only really caught on in recent years. Various workshops such as the "Department Of Energy (DOE) HPC Best Practices" workshop (2007 to 2011)[86] and its continuation from 2013 in the format of "The HPC Operations Review",[87] the "European HPC Centre Infrastructure Workshop"[88] series that has been running since 2008 and the annual Energy Efficiency HPC Working Group (EE HPC WG) workshops that have taken place since 2010 bear testimony to this.[89] These forums provide a great opportunity for sites to learn from the experiences of their peers.

2.4 **Data centre standards**

2.4.1 *The definition and purpose of a standard*

According to the British Standards Institution, a standard is "… an agreed way of doing something. … Standards are the distilled wisdom of people with expertise in their subject matter and who know the needs of the organizations they represent…. Standards are knowledge. They are powerful tools that can help drive innovation and increase productivity. They can make organizations more successful and people's everyday lives easier, safer and healthier…. The point of a standard is to provide a reliable basis for people to share the same expectations about a product or service."[90]
The International Organization for Standardization (ISO) defines a standard as "…a document that provides requirements, specifications, guidelines or characteristics that

[85] The Data Center Journal 2016, no page reference
[86] http://www.nersc.gov/events/hpc-workshops/
[87] http://science.energy.gov/ascr/community-resources/workshops-and-conferences/hpc-operations-review-and-best-practices-workshops/
[88] https://www.lrz.de/services/termine/infrastructure2016/
[89] https://eehpcwg.llnl.gov see "Conferences"
[90] www.bsigroup.com/en-GB/standards/Information -about-standards/what-is-a-standard/

can be used consistently to ensure that materials, products, processes and services are fit for their purpose.”[91]

2.4.2 *Relevant standards for data centres*

A number of professional bodies and organisations provide standards specific to sub-systems that are also found in data centres (e.g. power, cooling, ventilation, etc.). For the purpose of this study we will discuss only those standards that are explicitly aimed at the data centre as a whole and address all of the involved disciplines.

The following section aims to provide an overview and insight into the most relevant and widely recognised data centre standards. The earliest data centre standard dates back to the late 1990s.[92] An increasingly digitalized economy means data centres are increasingly critical to the businesses they support thus generating the need for standards and practices to ensure the integrity and functionality of the equipment therein.[93] The main driver and focus of these standards is therefore the availability and reliability of data centre infrastructure. Availability is also commonly referred to as “uptime”, whilst loss of availability is referred to as “downtime”.

For this work we will be reviewing the following data centre standards:

- Uptime Institute Data Center Tier Classification and Performance Standard
- ANSI/ TIA-942 Telecommunications Infrastructure Standard for Data Centers
- ANSI/ BICSI 002-2014, Data Center Design & Implementation Best Practices
- EN 50600 Series
- The American Society for Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) Datacom Series

Each standard is covered succinctly hereafter. As a number of the ASHRAE Datacom series specifically address HPC environments a more detailed overview of these has been provided in appendix four for reference.

Uptime Institute Data Center Tier Classification and Performance Standard

The Uptime Institute Tier Standard classification was first published in the late 1990s and played an instrumental role in standardising data centre design and construction. It classifies data centres into four hierarchical Tiers based on the availability of the physical infra-

---

[91] http://www.iso.org/iso/home/standards.htm
[92] Donnelly 2016, no page reference
[93] ASHRAE 2007, p. 3 - 5

structure. It is possible to have a data centre Tier level certified by the Uptime Institute. The Tier level is determined by the subsystem with the lowest level of availability.[94]

As per the Uptime Institutes specifications, the Tier Classification and Performance Standard provides the basis to compare different sites based on their infrastructure design topology and inherent functionality, capacities and cost. It is based on the following two documents that should be used in conjunction in order to ensure that the expected performances are achieved.

- The Tier Standard Topology document describes various possible topology solutions for each classification as well as a set of performance tests to determine that the end product complies with the defined requirements and chosen Tier classification. It describes requirements for power and cooling equipment with respect to their level of redundancy as well as their ability to allow planned maintenance and sustain unplanned outages without impacting the critical IT load.[95,96]
- The Tier Standard Operational Sustainability document outlines a methodology to align the facility management of a data centre with its Tier level so as to achieve the expected performance and availability levels. It describes behaviours and risk management procedures that should be integrated during the planning stage of the data centre and carried through into operation.[97]

The Tier classification of a site impacts its expected availability from an IT user perspective.

- *Tier I* has single distribution paths with N active components. This setup is not concurrently maintainable or fault tolerant and provides an average availability of 99.67% or 28.8 hours downtime per year.
- *Tier II* has single distribution paths but N+1 active components. It is not concurrently maintainable or fault tolerant but achieves an average availability of 99.75% or 22 hours of down time per year.
- *Tier III* has two distribution paths - one active and one alternative – and N+1 active components. This setup is concurrently maintainable but not fault toler-

---

[94] Donnelly 2016, no page reference
[95] Turner et al. 2008, p. 3
[96] Uptime Institute 2012, p. 1 - 3
[97] Uptime Institute 2013, p. 1 - 2

ant and has an average availability of 99.98% or 1.6 hours of downtime per year.

- *Tier IV* has two simultaneously active distribution paths and is both concurrently maintainable and fault tolerant. This setup achieves an average availability of 99.99% or 0.8 hours of downtime per year.

2.4.3 *ANSI/TIA-942 Telecommunications Infrastructure Standard for Data Centers*

The Telecommunications Industry Association (TIA) standard 942 on telecommunications infrastructure standards for data centres was first published in 2005 and updated in 2013. It is accredited by the American National Standard Institute (ANSI). As its name indicates, this standard is firmly rooted in the telecommunications industry.[98]
The scope of this standard is to specify the minimum requirements for data centre telecommunications infrastructure.
The standard provides the requirements for the design and installation of a data centre, computer room or extension of an existing setup. It provides designers with the input necessary to optimally integrate telecommunications requirements into their design at the earliest possible stage. The standard discusses data centre telecommunication spaces and topologies, cabling systems, cabling pathways and data centre redundancy from a telecommunications point of view. Within these topics, the standard also discusses the requirements that these pose on the surrounding building and technical infrastructure. The ANSI/TIA-942 Tier levels are congruent with those of the Uptime Institute.

2.4.4 *ANSI/BISCI 002-2014, Data Center Design & Implementation Best Practices*

Recognised by the American National Standards Institute the Building Industry Consulting Service International (BICSI) 002-2014 data centre design standard was first published in June 2010 with a revision in March 2011 and publication of the current edition in December 2014. BICSI is the worldwide association for cabling design and installation professionals.[99]

---

[98] Telecommunications Industry Association 2014, p. vii - xi
[99] BICSI 2014, p. 1 - 2

This standard defines both mandatory and advisory criteria for the design of data centres. It covers the topics of site selection, space planning, architectural requirements, electrical systems, mechanical systems, fire protection, security telecommunications cabling, infrastructure, pathways and spaces, and information technology and provides useful decision trees for a number of topics. The annexes provide informative content on topics such as the design process, reliability and availability, the alignment of data centre services reliability with application and system architecture, data centre services outsourcing models, multi-data centre architecture and examples of testing documentation.[100]

2.4.5 *EN 50600 Information Technology - Data centre facilities and infrastructures*

The European EN 50600 series "Information Technology - Data centre facilities and infrastructures" currently comprises seven standards that were approved by the European Committee for Electrotechnical Standardization (CENELEC) between 2012 and 2016. The English version is published by the British Standards Institution (BSI).
The following provides a brief overview of the topics covered by each publication.

- EN 50600-1: General concepts - defines common aspects such as terminology, parameters, reference models and general design principles that form the basis for the EN 50600 standard series.[101]
- EN 50500-2: Building construction - addresses the construction of buildings and other structures that may accommodate a data centre.[102]
- EN 50600-2-2: Power distribution - addresses power supplies to and within the data centre.[103]
- EN 50600-2-3: Environmental control - addresses the topic of environmental control.[104]
- EN 50600-2-4: Telecommunications cabling infrastructure - addresses the topic of telecommunications cabling infrastructure.[105]
- EN 50600-2-5: Security systems - addresses the topic of physical security.[106]

---

[100] Ibid
[101] BSI 2012, p. 5 - 7
[102] BSI 2014a, p. 4 - 7
[103] BSI 2014b, p. 4 - 7
[104] BSI 2014c, p. 4 - 7
[105] BSI 2015, p. 5 - 7

- EN 50600-2-6: Management and operational information - addresses the information, documentation, management and operational processes necessary to ensure the safe and energy efficient operation of a data centre.[107]

2.4.6 *The American Society for Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) Datacom Series*

The American Society for Heating, Refrigeration and Air-Conditioning Engineers (ASHRAE) dates back to 1894 and contributes to shaping built environments through the publication of research, standards and education on the topics of building systems, energy efficiency, indoor air quality, refrigeration and sustainability.[108]

The Datacom Series is authored by the ASHRAE Technical Committee 9.9 [109] that was founded in 2003 to address concerns around growing equipment power density in data centres and the increasingly mission critical nature of these facilities. The work of TC 9.9 covers mission critical facilities, data centres, technology spaces and electronic equipment. TC 9.9 collaborates with professional telecom organisations such as TIA and BICSI and continues to collaborate internationally to harmonise industry approaches. The Datacom Series is comprised of thirteen publications; a brief overview of each of them is provided here.

- *Thermal Guidelines for Data Processing Environments* first appeared in 2004 and the current fourth edition was published in 2015. It provides comprehensive information regarding temperature and humidity requirements for IT equipment and defines 5 envelopes for air-cooled environments for which vendors test the functionality of their equipment. This standard has been expanded to comprise definitions of environmental classes for liquid-cooling.[110]
- *Datacom Equipment Power Trends and Cooling Applications* was first published in 2005 and updated in 2012. It discusses the implications of trends in IT power density and cooling technology on the design of a data centre.[111]

---

[106] BSI 2016, p. 4 - 7
[107] BSI 2014d, p. 5 - 7
[108] https://www.ashrae.org/about-ashrae
[109] https://tc0909.ashraetcs.org/about.php
[110] ASHRAE 2015a, p. xi – xii, 1 – 8
[111] ASHRAE 2012, p. ix, 1 - 14

- *Design Considerations for Datacom Equipment Centres* first appeared in 2005. The current and second edition was published in 2009. It covers the basic design criteria for data centre facilities including Heating Ventilation and Air Conditioning (HVAC) loads, cooling systems, air distribution and liquid cooling, as well as providing information on ancillary spaces, contamination, acoustics, structural and seismic design, fire suppression, commissioning, availability, redundancy, and energy efficiency.[112]
- *Liquid Cooling Guidelines for Datacom Equipment Centres* was first published in 2006 and the current second edition appeared in 2014. This provides guidelines for liquid cooling strategies and distribution topologies. IT equipment vendors frequently refer to the water quality requirements described in this publication.[113]
- *Structural and Vibration Guidelines for Datacom Equipment Centers* addresses the increasing requirements of datacom facilities in terms of structure and vibration performance. This guideline provides requirements to keep a datacom facility up and running even during and after more extreme natural or man-made events and goes beyond those set out by building codes.[114]
- *Best Practices for Datacom Facility Energy Efficiency* was first published in 2008. The second and current edition appeared in 2009. This guideline provides detailed information on how to minimise life-cycle cost of a data centre and maximise energy efficiency by applying sustainable design approaches.[115]
- *High Density Data Centres – Case Studies and Best Practices* was published in 2008 and discusses seven ventilation schemes that are frequently applied in the industry.[116]
- *Particulate and Gaseous Contamination in Datacom Environments* was first published in 2009. The current and most recent edition appeared in 2013. It discusses monitoring, prevention and control of particulate and gaseous contaminations in data centres.[117]

---

[112] ASHRAE 2009a, p. 1 - 4
[113] ASHRAE 2014a, p. 1 - 8
[114] ASHRAE 2007, p. 1 - 8
[115] ASHRAE 2008a, p. 1 - 11
[116] ASHRAE 2008b, p. 1 - 6
[117] ASHRAE 2014b, p. v – vi, 1 - 11

- *Real-Time Energy Consumption Measurements in Data Centres* was published in 2010 with the aim to reduce annual energy consumption by implementing real-time energy efficiency measures that would allow data centres to measure and improve their energy consumption. It discusses energy efficiency metrics PUE and Data Center infrastructure Efficiency (DCiE) formulated by The Green Grid in 2007.[118]
- *Green Tips for Data Centers*, published in 2011, discusses techniques for optimizing energy efficiency and carbon footprint of existing data centres that can be easily implemented at relatively low cost.[119]
- *PUE™: A Comprehensive Examination of the Metric* was published in 2013. Its aim is to provide executives with a high level understanding of the concepts surrounding the PUE metric, at the same time as providing those implementing and reporting data centre metrics with in-depth application knowledge and resources.[120]
- *Server Efficiency – Metrics for Computer Servers and Storage* was published in 2015. It provides an in-depth description of the tools available to quantify the energy consumption of IT equipment, with the aim of providing managers with the information needed to relate product requirements to their specific environments. It introduces the metric of performance per watt for IT equipment.[121]
- *IT Equipment Design Impact on Data Center Solutions* was published in 2016. It provides detailed information about cooling design and thermal management and trends for IT equipment and their impact data on centre operations.[122]

2.4.7 ***Conclusion regarding reviewed standards and their applicability to public research HPC data centres***

In conclusion to the overview of these data centre standards we find that they each have their individual focuses and provide different levels of detail.
The TIA-942-A standard has a strong focus on telecommunication and cabling infrastructure.

---

[118] ASHRAE 2009b, p. 3 - 13
[119] ASHRAE 2011, p. ix-x, 1 - 9
[120] ASHRAE 2013, p. 1 - 8
[121] ASHRAE 2015b, p. xv – xviii, 1 - 13
[122] ASHRAE 2016, p. ix – xi, 1

The BICSI/ANSI 002-2014 provides extensive coverage of topics relevant to data centre construction. The standard also suggests decision trees to assist in determining some of the key design criteria.

The 50600 series publications are kept very succinct and provide a general overview of the relevant topics whilst referencing the respective standards that will provide greater detail. They provide guidance regarding electricity and physical security but barely mention mechanical infrastructure.

The Uptime Institute guides provide guidance for both topologies and management practices to attain a desired availability level for a data centre. They do not specifically discuss energy efficiency.

The ASHRAE Datacom Series is the only standard to specifically address the requirements of high-density environments and HPC discussed in chapters 2.3.3 and 2.3.4, as well as having a strong focus on metrics and energy efficiency. Liquid cooling is covered extensively and this is the only standard that covers particulate and gaseous contamination.

Based on this review, the author concludes, that as far as the fundamentals of design process, site selection, space planning and maintenance are concerned the interested reader would be well served with the BICSI/ANSI 002-2014 standard. Networking and telecommunications people will find the TIA-942-A helpful, whilst the ASHRAE Datacom series provides in-depth and complementary information on a number of key topics for data centre operators. Interested readers from the HPC domain will find the ASHREA publication *Thermal Guidelines for Data Processing Environments* specifically addresses the high-density environment of high performance computing. It is important to note that although the publication *High Density Data Centres – Case Studies and Best Practices* also specifically addresses the HPC environment, this book was published prior to the return of liquid cooling and therefore does not cover this topic.

3 Methodology

This chapter describes how, based on the literature review, questions were formulated and interviews conducted to collect data in order to be able to answer the research questions and test the hypotheses set out in chapter 1.4.

3.1 **Choice of method**

The research was conducted in person via structured phone interviews with persons in charge of HPC data centre design and/or operation. The idea of an online survey was discarded after a sample survey during the European HPC Centre Infrastructure Workshop showed that members of the target audience were less prepared to respond to such an online survey but willing to cover the same material in an interview. The interview questions were tested with four staff members of the Swiss National Supercomputing Centre (CSCS).

3.2 **Survey design**

Based on the hypotheses formulated in chapter 1.4 the survey is structured in five parts. The first part comprises questions aimed at determining to what extent existing data centre standards were known to the interview partners and whether or not they had been used to inform the design criteria for their construction or extension projects.
Part two focuses on understanding in what areas the requirements for public research HPC data centres differ from those of enterprise data centres and for what reasons. This was accomplished by comparing the values for a range of attributes found in the data centre industry with the values for these same attributes found at the public research HPC sites interviewed.
Part three is aimed at discovering which design criteria the interviewees found particularly challenging to determine and what strategies and thought processes they adopted to overcome these challenges. Part four is aimed at understanding what future proofing strategies the interviewees had applied to protect their capital investment. The last part provided interviewees with the opportunity to bring up any topic or issue they would have liked to discuss that had not been covered by the questions in parts one to four. The full survey can be found in appendix five.

3.3 **Selection of sample**

The HPC community has, since 1993, recorded and ranked the fastest general-purpose supercomputers from around the world on the Top500 list.[123] Since then this list has been published twice a year, in June and November, and records details such as vendor, location of the computer, its size, and performance when running the Linpack

[123] See top500.org/lists/2015/11/

benchmark.[124] Under this test the system will usually reach maximum power draw thus also acting as a stress test for the building infrastructure supporting the system.

For this study, we are interested in sites that have a sufficiently significant proportion of HPC racks in their data centre for these systems to drive their building requirements. As the different sites vary substantially in terms of size of the centre as well as size of the system, there is no clear-cut watershed. For the scope of this study the systems and sites appearing in the first one hundred positions of the Top500 list of November 2015 were selected, as these systems are sufficiently large to substantially impact and drive the sites' requirement in terms of electricity, cooling and footprint.

These top one hundred systems are hosted at seventy-five different sites, with some sites hosting more than one. Of these sites twenty-nine will not be considered for this study because they are not destined to use by public research (i.e. commercial, classified) or are dedicated to the purpose of weather prediction.

Commercial and classified sites will not be pursued as they fall outside the scope of public research. Sites that run weather prediction systems will not be considered for this study as they are subject to service level agreements that significantly impact their building requirements and design, usually requiring the site to conform to one of the higher Tier levels. The remaining forty-six sites comprise the following seven DOE laboratories in the USA:[125]

- Argonne National Laboratory (ARL)
- Lawrence Berkeley National Laboratory (LBNL)
- Lawrence Livermore National Laboratory (LLNL)
- Los Alamos National Laboratory (LANL)
- Oak Ridge National Laboratory (ORNL)
- Pacific Northwest National Laboratory (PNNL)
- Sandia National Laboratory (SNL)

and a further sixteen European sites that are part of the PRACE Research Infrastructure[126] comprising:

- Barcelona Supercomputing Center (BSC), Spain
- CEA Très Grand Centre de Calcul (TGCC), France
- Centre Informatique National de l'Enseignement Supérieur (CINES), France

---

[124] See http://www.top500.org/project/linpack/
[125] http://energy.gov/about-national-labs
[126] http://www.prace-ri.eu

- Centre National de la Recherche Scientifique – Institut du Développement et des Ressources en Informatique Scientifique (CNRS – IDRIS), France
- Center for Scientific Computing (CSC), Finland
- Centro di supercalcolo, Consorzio di università (CINECA), Italy
- Forschungszentrum Jülich (FZJ), Germany
- Höchstleistugnsrechenzentrum Stuttgart (HLRS), Germany
- IT4Innovations National Supercomputing Centre, Czech Republic
- Leibniz Rechenzentrum (LRZ), Germany
- Poznan Supercomputing and Networking Centre (PSNC), Poland
- Royal Institute of Technology (KTH), Sweden
- Science and Technology Facilities Council – Daresbury Laboratory (STFC), United Kingdom
- SURFsara, The Netherlands
- Swiss National Supercomputing Centre (CSCS), Switzerland
- University of Edinburgh (EPCC), United Kingdom.

In addition to these, the author reached out to the following four sites: National Renewable Energy Laboratory[127] in the U.S., Pawsey Supercomputing Centre[128] and the National Computational Infrastructure[129] in Australia, as well as RIKEN[130] in Japan. The sample thus comprises twenty-seven public research HPC sites from the top one hundred positions of the Top500 list from the USA, Europe and Asia-Pacific. Appendix six gives a detailed overview of the rank occupied by the systems hosted at these sites on the November 2015 Top500 list.

3.4 Data analysis method

3.4.1 *Familiarity with and application of existing data centre standards*

The answers provided in part one were checked for mentions of data centre standards known to the interviewee and coded as follows:

0 = not known 1 = known

[127] http://www.nrel.gov, DOE laboratory that is not listed on the Top500 list.
[128] https://www.pawsey.org.au
[129] http://nci.org.au
[130] http://www.riken.jp/en/

Regarding the use of the standard in informing design decisions, the author noted that a number of interviewees were unsure as to whether the standards had been applied by their design teams.
For this reason answers were coded as follows:
0 = not used 1 = used 2 = possibly used by design team
This data in conjunction with chapter 2.4.7 allowed us to answer research question 1.

3.4.2 *Comparison of design attributes between data centre industry and HPC*

The answers provided in the second part of the interviews were compiled and compared to the values observed in the wider data centre industry. Where the values between the two types of data centre differed noticeably the data was analysed more closely within the sample. This was the case for the load per cabinet, the raised floor height and ratings, use of Uninterruptible Power Supply (UPS) power, cooling technology and the planning and construction duration.
For the analysis answers were either compared directly based on the value or coded as follows:

Building type:	1 = tenant	2 = stand alone		
Staffing shifts:	1 = 1 shift	2 = 1 -2 shifts	3 =24/7	4 = 1 shift with On-call service
Equipment on UPS:	0 = none	1 = all	2 = critical loads only	
Cooling medium:[131]	1 = air	2 = liquid	3 = hybrid	

The answers provided regarding the reasons for the differences found between the data centre industry and public research HPC sites were extracted and compared in order to provide a compilation thereof.
This analysis provided the basis to ascertain where public research HPC centres differ from the data centre industry and compare these findings to the topics covered by data centre standards in order to find disparities.
This data allowed us to answer research question two.

[131] Based on the medium that removes the heat at its source. In the context of this work "hybrid" is used for situations where the heat is removed at the source by air but subsequently cooled by water within the data centre.

3.4.3 *Challenging design criteria and approaches to defining them*

Part three of the interviews was analysed by extracting a list of all the design criteria that were mentioned as having been challenging to define. Interviewees sometimes mentioned such challenges in other parts of the interview, therefore all parts were checked and any design challenges mentioned were included in the analysis.

The challenges were grouped into topics and a list compiled. The interview transcripts where then analysed for mentions of approaches or thought processes applied by the interviewees to overcome these challenges and the findings summarised.

This data allowed us to answer research question three.

3.4.4 *Future-proofing – challenges and strategies*

The fourth part of the interview was analysed in a similar fashion to part three. A list of future-proofing challenges and strategies was compiled from the interview transcripts and the findings summarised. This data allowed us to answer research question four.

3.4.5 *Scope for definition of a design standard for HPC centres*

For this part of the analysis, we compared the challenges and strategies interviewees identified in parts three and four of the interview. This data formed the basis to answering research question five.

3.5 Responses, sample size

For the survey, site representatives in charge of HPC data centre design and/or operation were contacted. Out of the twenty-six sites contacted, eighteen were available for interview and two reported data relative to two data centres belonging to the same site. The interviews were conducted by phone or videoconference, whereby the conversations were recorded with permission and later transcribed by the author. Once in written form, the interview was fact checked with the interviewee to ensure accuracy and veracity of their statements.

3.6 About the interviewees and their sites

3.6.1 *Geographic distribution*

Interviewees from the eighteen sites pertained to:

- eight DOE laboratories from the U.S.
- nine PRACE sites from Europe
- one site from Asia

In the analysis we refer to these eighteen *sites* that represent twenty *data centres*.

3.6.2 *Key parameters of the data centres in the sample*

The following section provides information about the sites in the sample as a background for the analysis and results. Of the twenty data centres from eighteen different sites, eleven were built prior to the return of liquid cooling around 2008. The remaining nine were built after this threshold.

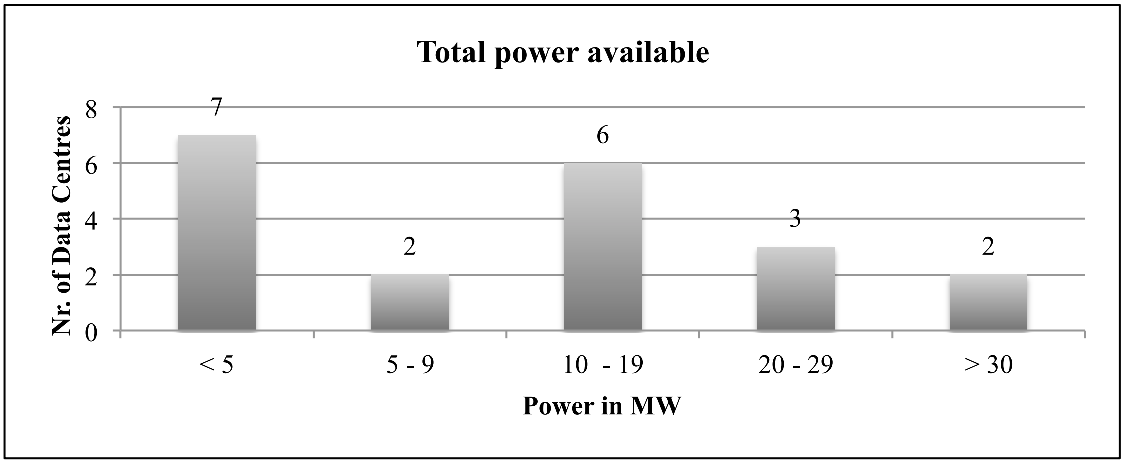

Figure 3: Total power available in MW

The power available to a site is one of the parameters that defines its ability to host compute power. Figure 3 shows how the sites in the sample are distributed in terms of power available to the data centre. Six of the seven sites that have less than 5MW available are in Europe, whilst of the five sites with 20MW or more of available power four are in the US and one in Japan.

The variety of machine room size within the sample is shown in Figure 4 below:

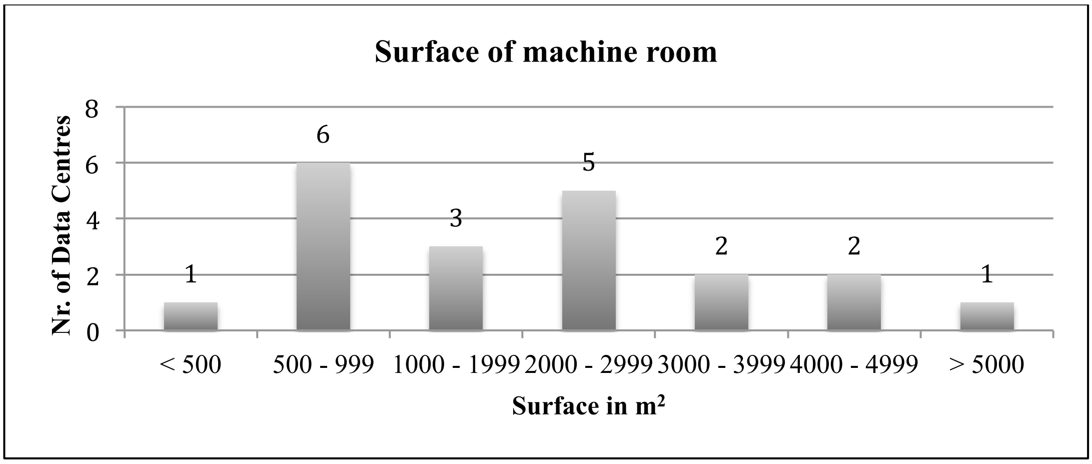

Figure 4: Surface of machine room in m^2

Four of the smaller rooms in the first two categories are in Europe and three in the U.S., whilst all three machine rooms in the two largest categories are in the U.S.

Sixteen of the data centres are stand-alone buildings dedicated to the HPC operation. In some cases these buildings include offices directly related to this activity. Three have machine rooms located within larger building complexes that house activities beyond those related to the HPC operation. One site's answers pertain to space rented within a co-location facility.

4 Findings

This chapter reviews the findings from the interviews in relation to the hypotheses and research questions formulated in chapter 1.4. We will first discuss the findings for the research questions before testing the relevant hypothesis.

4.1 Research question 1 – familiarity with and application of standards

The interviewees' answers show that the best-known standard in this is the ASHRAE Datacom Series, which is familiar to all but one site. With twelve sites reporting that they use it to inform their design decisions it is also the most used standard.

The Uptime Institute standard (indicated in Figure 5 by UI) is the second-best known standard amongst the sites but is only applied by four of the eleven sites who are familiar with it. Three sites reported that this standard does not fit HPC environments, as the redundant distributions required for the higher Tier levels are not affordable for

HPC power loads. Two sites also mentioned that they find the classification useful as it provides a common language, although they do not design to meet a Tier level.

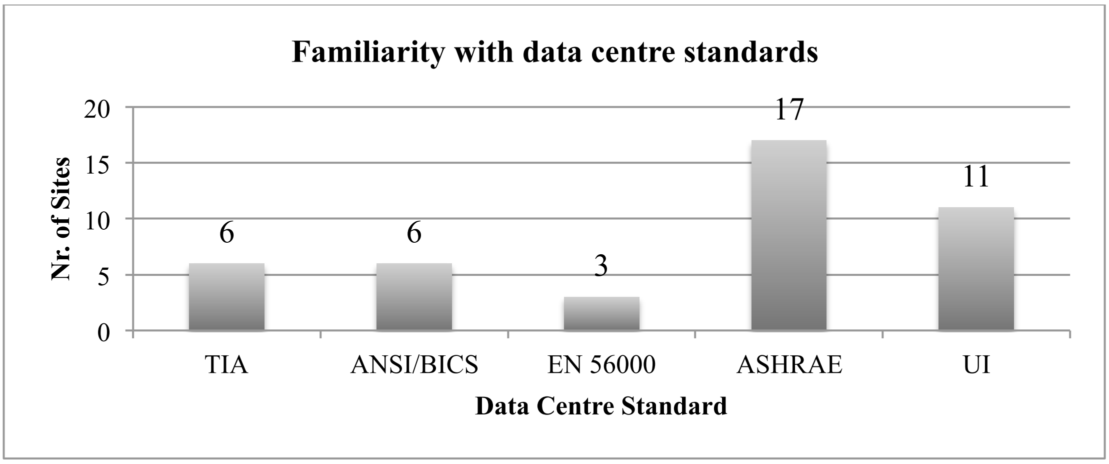

Figure 5: Familiarity with data centre standards

The EN 50600 Series is the least known and least applied standard. This may, in part, be due to the standard being unknown outside Europe and having only recently been published. The sites that use this standard are all located in Europe and have recently been involved with infrastructure upgrade projects.

The TIA-942-A standard is familiar to and used by one third of the sites that are equally distributed between Europe and the USA.

One third of the sites – two of these outside the U.S. - are familiar with the ANSI/BICSI 002-2014 standard, but only four sites used it to inform their design decisions. Those familiar with the standard noted that it was a very good standard for enterprise data centres. Whilst these sites mentioned that it is not directly applicable to HPC, all but one of these sites agreed it is a useful basis to build on and diverge from when the standard does not cover HPC requirements.

In the case of three of the sites, the interviewee did not have precise information about whether any of the standards were used by their design team.

Local building codes, fire and seismic regulations were also mentioned by some sites but are not discussed here as these apply for all sites. Fire protection was mentioned as an important issue by six sites due to the amount of power present in facilities of this kind.

In addition to the above, interviewees mentioned having based their design decisions on the sources reported in Table 2 (below). Each source is accompanied by an indication of how many sites mentioned it.

	Mentions
IT vendor requirements	11
National Fire Protection Association (U.S.)	6
Internal design criteria catalogue	1
U.S. GBC LEED[132]	1
DOE High Performance and Sustainable Buildings Guidance (U.S)[133]	1
European Code of Conduct[134]	1

Table 2: Additional sources that informed design decisions

In answer to research question 1, we observe that with the exception of the ASHRAE Datacom Series, existing data centre standards are not very well known or widely applied by public research HPC data centres.

4.2 Research question 2 – where standards do not cover HPC requirements

The interview responses to part two of the interview show, that beyond their mission, HPC centres differ markedly from the wider data centre industry[135] with regard to the seven attributes listed in Table 2 below.

Attribute	Data centre industry	Public research HPC sites in sample	
Load per cabinet	1 – 15kW	Min. 20 kW	Max. 100 kW
Raised floor height	Up to 1.1 m	Min. 0.4 m	Max. 6 m
Raised floor rating	Up to 1220 kg/m^2	Min. 980 kg/m^2	Max. 3410 kg/m^2
Equipment on UPS	All	13 sites: critical equipment only 6 sites: all equipment 1 site: no equipment	
Utility voltage	Tiers I + II low tension Tiers III + IV: medium tension	14 sites: medium tension 4 sites: low tension	
Cooling technology	Air	16 sites: liquid 16 sites: hybrid 14 sites: air[136]	
Time to build[137]	Up to 30 months	Min. 24 months	Max. 180 months

Table 3: Attributes that show significant differences between industry and HPC

[132] U.S. Green Building Council Leadership in Energy and Environmental Design standard
[133] http://www.wbdg.org/pdfs/hpsb_guidance.pdf
[134] Aimed at improving the energy efficiency of data centres. http://iet.jrc.ec.europa.eu/energyefficiency/ict-codes-conduct/data-centres-energy-efficiency
[135] Please refer to appendix 7 for the full list of attributes observed by the Uptime Institute in the wider data centre industry.
[136] All sites use a combination of cooling technologies.
[137] includes time spent securing funding and support for project

Table 3 shows the characteristic values for each attribute within the data centre industry and the minima and maxima for each attribute for the HPC sites.

The following section discusses these attributes and, where applicable, looks at where the current data centre standards do not cover the requirements identified for public research HPC sites. The topics of utility power and time to build are not covered by standards but discussed hereafter because they are areas in which there is a noticeable difference between the industry and HPC sites.

4.2.1 *Load per cabinet*

In contrast to the load per cabinet observed in the wider data centre industry, the HPC sites report significantly higher loads for compute cabinets as well as for networking and data storage cabinets. Figure 6 shows the maximum reported load densities per rack reported by the interviewees.

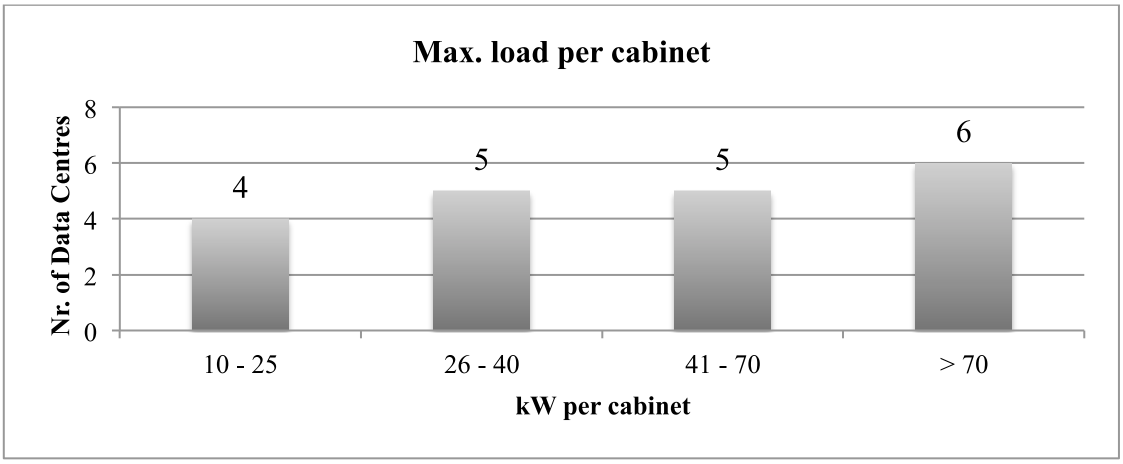

Figure 6: Maximum load per cabinet in kW

As we saw in chapter 2.3.2 all the sites in the sample have a mission to provide the compute capacity needed to pursue scientific grand challenges and advance the leading edge of scientific research.

When asked what, in their opinion, drove the compute density observed in HPC the interviewees named the following reasons:

- Latency and bandwidth sensitivity drives the need for shorter cables and thus more tightly packed systems.
- Need to pack ever greater compute capacity into a given footprint, energy and cooling envelope.

Given the nature of the scientific problems tackled by HPC, there is a need for the systems to be tightly coupled, making cable lengths important, especially for latency sensitive problems. The limitations of copper as well as its cost lead to the necessity of keeping cable lengths short.

Two sites noted that although the growth in compute capacity was desirable, the increasing density of the systems made them hard to accommodate.

As we saw in chapter 2.4.7 high-density environments are only covered by the ASHRAE Datacom series. The other standards do not cover this requirement.

4.2.2 *Raised floor height*

Raised floors in the wider data centre industry typically range from 0.3 - 1.1m (12 – 42 inches) and are rated for 415 – 732 kg/m^2 (85 – 150+ psf.).

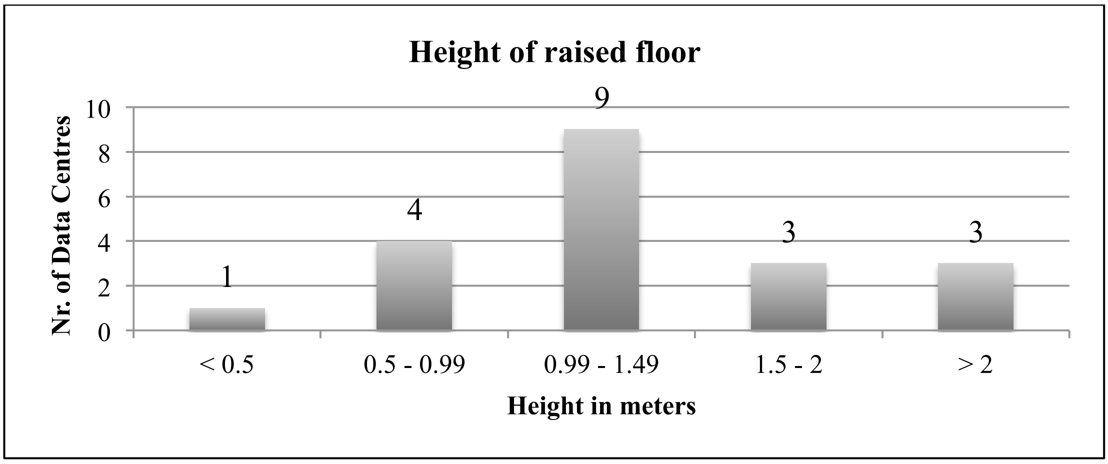

Figure 7: Height of raised floor in metres

As seen in Figure 7 above, the responses to the interview show that for public research sites the raised floor height varies between less than 0.5m (19.7 inches) to over 2m (78.7 inches). The site with the shallowest raised floor reported a height of 0.4m (15.7 inches) and three sites have significantly higher ones of respectively 5m (16.4 ft.) and 6m (19.7 ft.). The exact height for the third site was not indicated, but consists of an entire storey. In two of these cases, this significantly deeper raised floor was a specific design requirement, whilst for the third site it is a result of the data centre needing to be integrated into a larger building and therefore having to match the depths of adjacent spaces. Half the sites have a raised floor height between 0.99 and 1.49m (39.4 – 58.7 inches).

In enterprise data centres, the height of the raised floor has to be sufficient to provide adequate airflow to ensure uniform cooling. As seen in Table 3 (in Chapter 4.2), most HPC sites use some form of liquid cooling alongside traditional air-cooling. The piping for this is frequently accommodated within the raised floor. This drives the need for more depth in order to ensure proper airflow around the pipes. HPC sites therefore tend to have a preference for deeper raised floors. The preference for the 0.99 – 1.49m (39.4 – 58.7 inches) bracket in the sample may be related to the limitations that the various building codes[138] put on raised floor heights although this was only explicitly mentioned in one of the interviews.

At one end of the spectrum, we have the three sites with the deepest raised floors that report the design has proven beneficial in reducing the necessary space for the mechanical plant, shortening connections and facilitating maintenance procedures and installation of new machines. At the other end of the spectrum, one site is already planning to move to slab on grade[139] and a further two are advocating for this in order to avoid of load rating concerns. On the downside, the complexity of running communication cabling, power, fire detection and extinguishing systems, cooling and overhead lighting was mentioned. Sites are still divided about whether having electricity and cooling in the same space is something they feel comfortable with or not.

4.2.3 *Raised floor load rating*

When we look at floor ratings for static loads (See Figure 8 below) all but one of the sites interviewed have ratings that exceed those normally found in Tier IV data centres. There is a concentration in the 1500 – 1999 kg/m^2 (310 – 410 psf.) bracket of the sample, but nine sites have ratings in excess of 2000 kg/m^2 (410 psf.).

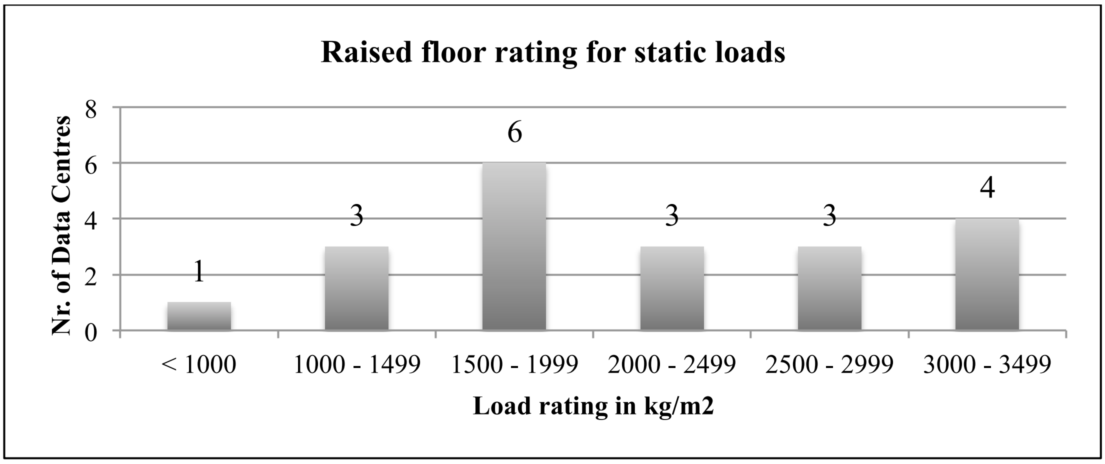

Figure 8: Raised floor rating for static loads (kg/m^2)

Eleven of the sites also specified a point load rating (kN or lbs.).

---

[138] e.g. regulations regarding permit required confined spaces in the U.S.

[139] Approach where systems are sited directly on the cement floor.

Due to their packaging density, HPC systems are significantly heavier than the average rack of equipment in an enterprise data centre. To compound this, a number of these racks are built to stand on four casters thereby concentrating the full weight of the rack in four relatively small points. Four of the sites interviewed report having had to reinforce their raised floors at considerable cost and effort to be able to accommodate new systems. One site had to reinforce the cement slab supporting the raised floor. The sites on the West Coast of the U.S. and in Japan face the additional challenge of having to contend with seismic risks.

The standards indicate the following ratings for raised floors:

	TIA[140]	ANSI/BISCI[141]	EN 50600[142]	ASHRAE[143]	UI
Static load	1221 kg/m^2	1221 kg/m^2	1221 kg/m^2	25% of point load	N/A
Point load	N/A	6.6 kN	5 kN	4.4 – 5.5 kN	N/A

Table 4: Raised floor ratings from standards

Based on a comparison in the above table and the attributes seen in the sample (Figure 8), the standards do not cater to the requirements of HPC sites with regard to raised floor ratings.

4.2.4 *Equipment on UPS*

As we saw in chapter 2.2.2, business continuity and uptime are crucial to the wider data centre industry. Therefore, all equipment usually has redundant power feeds and is protected by UPS and this is what the standards advocate for.

Due to the size of the loads seen at public research HPC sites, UPS backup is costly and for this reason thirteen of the sites interviewed only put their critical loads on UPS and one site does not provide any UPS at all as shown in Figure 9.

Six European sites still provide UPS backup for all equipment. In two cases this was linked to the requirement to ride through brownouts.[144] The point was made by one site that when equipment that is liquid cooled is on UPS it is important for the corresponding liquid cooling supply to also be on UPS in order to avoid overheating of the equipment due to a brownout.

[140] Telecommunications Industry Association 2014, p. 23
[141] BICSI 2014, p. 93
[142] BSI 2014, p. 22
[143] ASHRAE 2009a, p. 101 - 102
[144] A brownout is a drop in voltage in an electrical power supply system.

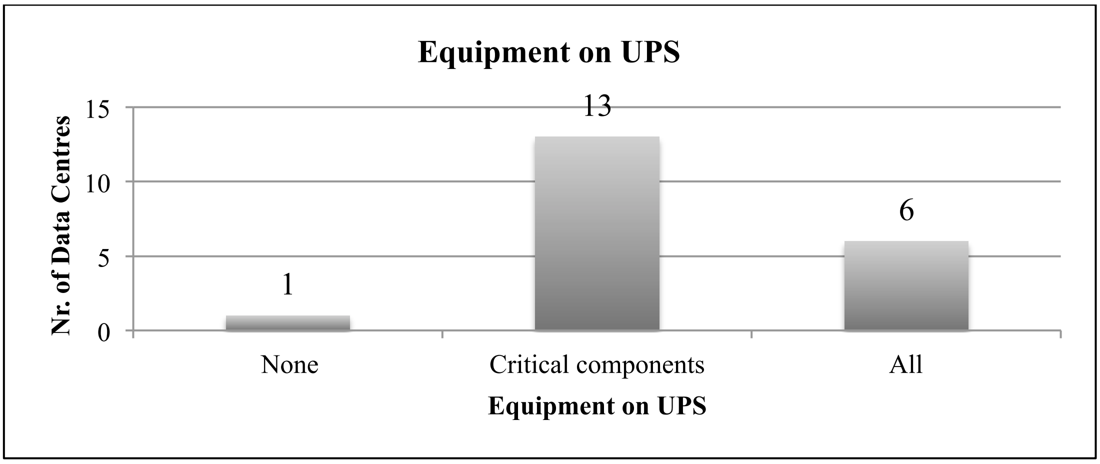

Figure 9:Equipment on UPS

Four U.S. sites report concerns related to potential degradations of the power grid that may increasingly cause blackouts.[145] This concern has led a number of the U.S. sites to consider various technologies from co-generation to modular nuclear reactors, which would allow them to be energy self-sufficient. So far, none of the sites has implemented such a solution due to unfavourable ROI valuations.

4.2.5 *Utility voltage*

In the wider data centre industry, lower Tier sites tend to be supplied by low-voltage, whilst higher rated sites are more likely to have medium-voltage. As indicated by the ANSI/BISCI standard, the utility company will usually determine incoming supply voltage based on the load required by the site.[146] Standards, therefore, do not provide explicit guidance as to utility voltage. As previously shown in Figure 3: Total power available in MW (see Chapter 3.6.2), eleven of the sites interviewed have a current capacity greater than 10MW and an additional three draw more than 5MW. This is probably a key reason why fourteen of the sites in the sample are supplied with medium-voltage, with only four being connected to low-voltage.

More recently, changes in processor technology, combined with the higher power density of systems, have led to HPC sites experiencing increasingly large load swings of multiple MW in very short timeframes. These can prove challenging for the infrastructure as well as for the utility company and nine sites in the sample raised this topic as one of the challenges they are facing.

[145] A blackout refers to a total loss of power.
[146] BICSI 2014, p. 117

4.2.6 *Cooling technology*

Within the wider data centre industry, the main cooling medium used is still air, as this remains the most economic solution for low-density environments. As seen in section 4.2.1 that is not what we are dealing with here. Due to the load densities observed in HPC, vendors began re-introducing liquid cooling in 2008 and this approach has since become widespread for HPC systems. Because HPC environments host other less dense systems beyond the HPC machines, they end up needing to provide for several different cooling technologies to accommodate the different systems. Low-density systems continue to be air cooled, but have in some cases been located in enclosures that are liquid cooled or equipped with rear door heat exchangers (RDHX).

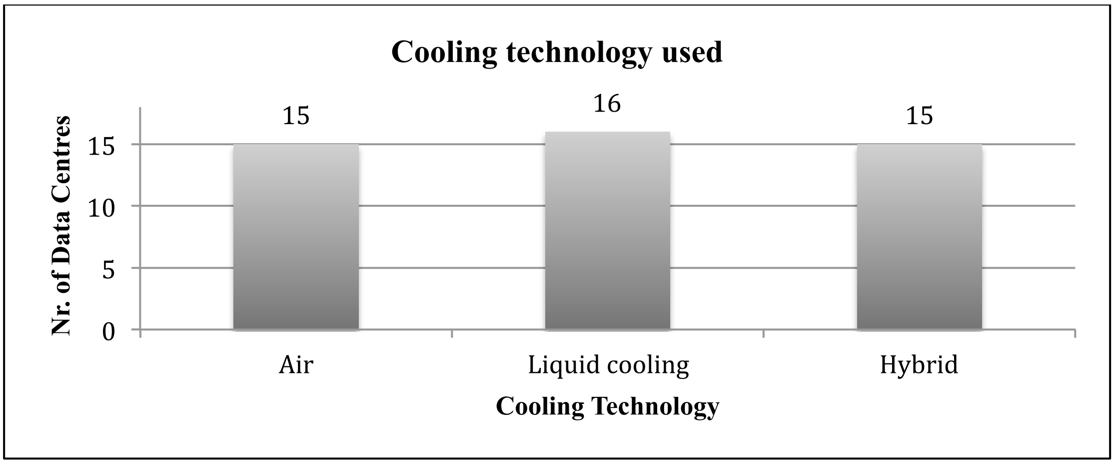

Figure 10: Cooling technologies used

Twelve of the sites use a combination of two cooling technologies, whilst the remaining six sites use all three cooling technologies. Figure 10 above shows the occurrences of the different cooling technologies within the sample.

Eleven of the sites have had retrofits for liquid cooling, whilst the five most recently built data centres were designed to include this from the outset. One of the main challenges sites face is deciding the temperature ranges for liquid cooling installations. Sites traditionally had cold water available at around 6°C as this is commonly used to supply the Computer Room Air Conditioning Units (CRACs) for air-cooled environments. This is, therefore, the supply that has typically been used for liquid cooling retrofits. Initial liquid cooled HPC systems were able to take advantage of this same temperature range, but newer systems have increasingly been designed to work with higher coolant temperatures. The ranges vary between vendors, making it difficult for sites to define and build a sustainable liquid cooling infrastructure. This causes increased costs and disruption, as parts of the liquid cooling infrastructure have to be modified or completely rebuilt with new system installations. Sites that were designed with liquid cooling tend to be less affected by this, as their designs comprise multiple temperature loops.

The TIA-942-A, the EN 50600 series and Uptime Institute Tier Classification do not address the topic of liquid cooling.

The 2014 edition of the ANSI/BICSI 002-2014 has been updated to include this cooling technology and explicitly refers to the following two ASHRAE Datacom Series handbooks already discussed for further guidance on the topic:

- Thermal Guidelines for Data Processing Environments, Third Edition
- Design Considerations for Datacom Equipment Centers, Second Edition

Regarding the extent to which HPC requirements for liquid cooling are covered by existing standards, we may conclude that although three of them do not cover the topic, two of them now do. In this sense, the most recent standards have caught up with HPC requirements in this area. The *Thermal Guidelines for Data Processing Environments* handbook includes a classification for liquid cooling temperature ranges[147] along with water quality requirements for the various cooling loops.[148] IT vendors frequently refer to the latter in their site-installation guides. The liquid cooling classification is increasingly being adopted too.

4.2.7 *Time to plan and build*

Although standards do not address this, time to solution is an area where HPC sites differ substantially from the wider industry. The average industry Tier IV data centre takes around thirty months to plan and build. If limited to planning and building alone, public research HPC centres are comparable. As the interview data shows, however, the overall process tends to be substantially longer for these sites.[149] The sites report times-to-solution from first idea to start of operation ranging from twenty-four months to 180 months. Only four sites were completed in less than thirty months. A further six completed in less than four years and seven took longer to complete.[150] It is worth noting that the U.S. has a different process for projects above or below the $10 million threshold. Below the threshold, projects are funded by programmatic funding and are

[147] ASHRAE 2015, p. 42
[148] ASHRAE 2015, p. 47
[149] This is mostly due to the time involved in getting political backing and funding for the project.
[150] Note that for the three oldest sites this data was not available.

subject to a less scrutinised process. Projects over and above this sum run as congressional line items and have a longer approval process and more extensive oversight.
A similar threshold applies for Switzerland at CHF10 million.[151] This means HPC sites need to either initiate the process for a new data centre very early – when it is difficult to justify the need to the funding agencies – or wait until the need is clearly manifest and hold out in a data centre they have outgrown until the process for a new one runs its course. Either way, the length of the process means technology could change between the start and completion of a project, and these changes may or may not be able to be integrated in the design stage.
Once built, funding agencies also expect these buildings to last for anything from twenty to fifty-five years and more. The sites interviewed agree that infrastructure is amortised sooner, and share the experience of computer lifecycles of three to five years. This means that a given building shell will see multiple infrastructure refreshes and many generations of supercomputers.

In answer to research question 2, we conclude that the ASHRAE Datacom Series best addresses the requirements of HPC. The newly updated ANSI/BICSI has integrated liquid cooling and refers back to ASHRAE for further details. The other standards do not cover the requirements of HPC environments. Table 5 (below) provides an overview of which requirements are covered by the standards.

	TIA	ANSI/BISCI	EN 50600	ASHRAE	UI
Power density	✗	✗	✗	✓	✗
Raised floor rating	✗	✗	✗	✗	✗
UPS	✗	✗	✗	N/A[152]	✗
Liquid cooling	✗	✓	✗	✓	✗

Table 5: Coverage of HPC requirements by standards

4.3 Research question 3 – defining design criteria not covered by standards

The following section discusses the main challenges mentioned by the sites in defining design criteria as well as approaches adopted in tackling them.

[151] Data was not collected for the other European countries.
[152] ASHRAE, given its focus, does not cover power distribution.

4.3.1 *Cooling technology, capacity & balance between technologies*

Mentioned by sixteen sites, the criteria interviewees found most challenging to define were related to the selection and capacity specification of cooling technologies. For new buildings, the challenge lay with deciding what cooling technologies to design for and how to define and balance their capacities. Older sites in the sample mostly faced the challenge of retrofitting for liquid cooling and having to define the temperature range and capacity for this.

Although sites agree that liquid cooling is a requirement for high-density environments and has the added benefit of being more energy efficient than air cooling, they are frustrated by the lack of uniformity in temperature ranges amongst different vendor solutions. Six sites also report increasing difficulty in fitting additional capacity into the same building envelope.

New sites have tackled this challenge by either specifying multiple cooling loops at different temperature ranges or supplying a single temperature range that forms a prerequisite for HPC-system tenders. They also allow for at least two different cooling technologies within their designs. In determining the cooling technology and capacity sites also make decisions about the environmental conditions they will offer as well as the potential for reutilisation of waste-heat for heating purposes. As far as overall cooling capacity goes, please refer to the description in the next chapter of the approach described by the interviewees.

4.3.2 *Power density & capacity*

The other equally challenging[153] design criteria are the power density and capacity. The increasing power density is pushing facilities to their limit. In some cases, the move from 208V to 480V distributions for most HPC systems has eased the pressure at least in terms of space required for the electrical infrastructure. All sites still run their electrical distribution in Alternate Current (AC). One site is looking at testing Direct Current (DC) that is known to transmit power more efficiently.

When designing new data centres six sites used power capacity as a starting point in defining their design. They all followed a fairly similar thought process that can be summed up as follows:

---

[153] Same number of mentions.

- Decide what compute capacity the site should be able to host as far out as imaginable.
- Extrapolate power and footprint for such a system based on today's technology, vendor R&D roadmaps and research.
- Assume the need to be able to host two of these systems in parallel for a limited amount of time when a new system comes online before the old one is retired. Add space, power and cooling capacity for miscellaneous systems and service area. This will allow you to extrapolate the amount of power and cooling capacity as well as the footprint you will need to design for.
- Design the envelope so it can accommodate the final capacity.
- Break infrastructure down into modules that can be added when needed.
- Pre-build connection points for addition of infrastructure modules.

Alternatively, one site has elected to build the optimal facility for the known requirements and keep it as generic as possible so as to be able to modify it to adapt to future changes in requirements.

4.3.3 *Foreseeing requirements of next generation of systems & disparate lifecycles*

The challenge of hosting numerous generations of machines in a building that has a much longer life cycle was discussed in chapters 2.3.3 and 4.2.7. Although only mentioned by eleven of the interviewees, this is a common challenge in the group. Approaches to mitigating this are further discussed in chapter 4.4.

4.3.4 *Environmental factors*

Six sites identified related environmental factors, such as the impact of their operation on their surroundings, meeting energy efficiency requirements and water usage limitations, as challenges. Energy efficiency has certainly become a key focus for all sites, although the sites vary as to whether this is driven by cost or by environmental considerations. Governments and legislative requirements are increasingly supporting the change in attitude at the societal level with regard to environmental consciousness.

4.3.5 *Raised floor ratings*

Five sites identified raised floor ratings as a challenge. As discussed in chapter 4.2.3, the rating is key to HPC sites. The main strategy that interviewees mentioned for defining these design criteria was to look at vendor requirements and their evolution and add some buffer.

4.3.6 *Commissioning of liquid cooling*

Two sites raised difficulty in commissioning liquid cooling infrastructures without the IT load being installed. One of the sites opted for point commissioning and delayed some of the testing and acceptance to when the first large machine was installed. The other site connected the district warm water to the facility to simulate the heat load during testing and commissioning. In both cases, the interviewees pointed out the importance of considering commissioning procedures early on in the design and construction process.

4.3.7 *Defining interfaces between infrastructure and IT equipment*

Five sites noted the difficulty of specifying the power and cooling distribution and connections for future systems with yet unknown requirements. The most frequently mentioned approach to mitigating this challenge is to prepare the main electrical and cooling distribution and finalise the distribution to the IT system once the requirements are known. However, two sites mention problems in applying this strategy, due to the limited amount of time available to execute the connections once the requirements are known. One site has mitigated this by tendering for long-term contracts with an architect and engineering partner (A&E) as well as contractors. This allows the site to commission for this work without having to go through another tender.

4.4 **Research question 4 - Future-proofing – challenges and strategies**

In part four of the interview, we asked sites about what future proofing measures they had put in place to mitigate the insecurities about the future evolution of technology and its requirements. The measures mentioned were grouped into topics discussed hereafter. It is worth mentioning that two sites indicated they had not invested in future-proofing as they would not have been able to get funding support for this. One site was criticized for investing in future-proofing measures.

4.4.1 *Design infrastructure for growth*

Fifteen sites mentioned future-proofing measures that fall within this topic. Sites having to upgrade existing infrastructures advocate strongly for planning for extra headroom and connection points, so as to be able to accommodate further growth and make the investment more durable. They also note the need to be careful to not end up with stranded[154] or trapped[155] capacity.

When designing a new building, ten sites in the sample support for designing for the final expected capacity. For infrastructure, the most frequently applied strategy is to build out those parts that are costly and disruptive in order to change, plan and execute the ones comparatively easy to modify in a modular fashion, so it can grow in step with capacity needs. In essence, flexibility is key.

4.4.2 *Building envelope – big vs. modular*

On this topic, there are two directions of thought. The first and most frequently represented one in the sample (9 sites) advocates for building as large a building envelope as can be justified in order to be able to react to changes within this. In this case the advice of one site to make sure you get your steel and concrete right is important, as is the advice of keeping some flexibility.

The second, pursued by one site and mentioned by another, is to build a modular envelope and optimise it based on the known requirements whilst keeping it generic enough in order to modify if or when requirements change.

4.4.3 *Plan for the future*

This subject was raised by four sites, which mentioned the importance of having a business plan. In the event that a site has multiple facilities the importance of also formulating a portfolio plan was raised by one site. The sites suggest that such plans should include growth scenarios and future options one may want to execute so as to avoid their obstruction. Fourteen sites also mentioned the value of continuously talking to vendors about their future technology roadmap. One site extended this advice to also comprise being well informed about how technologies evolved in the past in order to recognise patterns if they repeat and to be able to anticipate their requirements.

4.4.4 *Diversity in cooling technologies*

Four sites mentioned measures in this area, which can essentially be broken down into the following two strategies. The first consists in designing for the ability to cool with

---

[154] Due to overprovisioning.
[155] Due to underuse.

different technologies. The second includes cooling loops of different temperature ranges. Both of these measures apply primarily to new designs. They contain the challenge of selecting the balance of capacities for each of the cooling technologies (also mentioned in chapter 4.3.1) or that of selecting the capacity balance between temperature loops.

4.4.5 *Invest in a raised floor with high specifications*

This topic was only explicitly raised by two sites, but given the number of sites that have had to upgrade their raised floors and noted this as a challenge (see chapter 4.2.3) it is worth discussing. The advice regarding raised floors the interviewees gave was to invest in the best raised floor attainable and make sure that the entire raised floor surface is uninterrupted. The sites mention that although this increases the upfront cost, it is well worth the investment in order to ensure performance and avoid expensive and disruptive subsequent upgrades. Ratings for raised floor, support structure and cement slab need to be coherent in order to ensure the requisite performance of the structure.

4.5 Testing of hypothesis I

Based on the discussion of research questions one to four, we find data centre standards do not cover a number of the key requirements seen in HPC. This is largely due to the fact that the HPC sites included in this sample are frequently amongst the first to install new technologies and gain experience with them, which may later give rise to the formulation of a standard. The mission of these sites means that they will always be one step ahead of the standards when it comes to new technologies. This does not, however, make data centre standards entirely inapplicable or unhelpful for HPC sites. They can provide a good basis or framework upon which the necessary adaptations and solutions can be developed that will reflect the needs of HPC. Due to the fact that many IT vendors are now referencing the ASHRAE Datacom Series these should be compulsory reading.

4.6 Research question 5 - Scope for definition of a design standard for HPC

Research question five asks if there is sufficient similarity between the approaches of the various HPC sites to allow the definition of a standard for the design of public research HPC centres. Throughout the development of this thesis, it has become clear that due to the nature of the mission of these sites and that of design standards as articulated in the previous chapter, the formulation of a standard covering HPC requirements - including those that lie just around the corner in the future – is not what we should be aiming to achieve, as it would most likely be superseded by the next development step in technology. Instead, it would be preferable to invest in collecting and formulating best practices applied by the community to improve the sharing of this knowledge.

4.7 Testing of hypothesis II

Although, as mentioned above, the formulation of a standard may not be useful, this work has drawn on the data collected in the interviews to compile a first set of best practices applied within the peer community. The author's hope is that this will facilitate the sharing of available knowledge and provide a starting point for the community to elaborate a complete compilation of best practices. The usefulness of this is underlined by the fact that nine sites mentioned the importance of attending workshops on HPC infrastructure and seven sites mentioned having invested time in visiting other sites to learn their best practices prior to engaging in a design project for a new building.

Based on the interviews, the best practices mentioned have been collected, grouped and structured into the compilation presented in chapter 4.8. This first attempt is by no means complete or exhaustive and should be seen as a starting point for a more extensive compilation.

4.8 Compilation of best practices

The following chapter is based on the best practices mentioned by sites during the interviews. By best practices we mean approaches that proved successful for the interviewed sites. They have been grouped and structured into chapters for ease of reference.

4.8.1 *Management topics*

- Have a clearly formulated and agreed business plan and, if applicable, also have a portfolio plan before you start. This will ensure that you know your assets and their limits and can plan for costs throughout the lifecycle. For new projects this will allow you to clearly identify and defend your requirements for the design. Update these plans frequently.

- Build support for your project at each level so the funding request can be defended and supported by each in turn.
- Get full backing from all stakeholders to ensure everyone is on the same page and focused on the same goal.
- Know and understand government or other policy makers' requirements that will affect your project. This may include energy efficiency, funding thresholds, carbon emissions, etc.
- Beyond the expertise in IT systems, invest in hiring (or at least contracting) the necessary skills in terms of electrical and mechanical engineering as well as building automation. This will be invaluable during the design and construction phase, but will also allow you to fine tune and optimise the facility when in operation as well as ensuring that you have key knowledge about your facility in house.
- Ensure facility and IT system managers are both involved throughout the full design and construction process to ensure optimal fit with requirements, shared responsibility for decisions and a smoother and faster handover to operations on completion of the project.
- Ask a lot of questions, exchange on best practices with peers and get regular updates from vendors about their technology roadmaps.
- Take the long view: understanding how technology changed in the past is valuable as a background when guessing at the future. It will allow you to recognise repeat patterns and pre-empt them.
- Be aware of ASHRAE Datacom Series. ANSI/BISCI can provide a helpful basis if starting from scratch.
- Prior to talking to an A&E firm, put together a catalogue of your design and performance requirements as well as any standards you want the design to adhere to.
- Define no more than five design features considered make or break to the success of the project. This will ensure you do not make trade-offs in the wrong places and that you know where you can compromise.
- Review the different possible tendering processes and choose the one that best suits your needs. (See next section for further details)

4.8.2 *Tendering processes*

- Data centre experience is imperative for your A&E partner. Experience of HPC is desirable.
- Possible tendering options:
    - Each trade is tendered singularly. High administrative overhead, potentially longer process, coordination onerous, responsibility and warranty grey zones.
    - A&E partner and general contractor: Separate tenders are put out for an A&E partner and a general contractor. Only two contractual partners. Potentially shorter process, responsibilities and warranty clearer. (Also referred to as design-bid-build in the U.S.)
    - Design-build (also referred to as Integrated Design): Single tender results in a contract with a combined team of A&E and contractors that go through the design and construction process together. Collaborative and iterative process. Allows for on-going optimisation and shorter turn-around of the project. The DOE High Performance and Sustainability Guidance[156] requires DOE sites to employ integrated design principles. This form of tendering is fairly recent and in Europe it is not yet widely applied because tendering processes and standard contracts have not yet been adapted to allow for them. Progress is being made on this front though, so it is worth looking into in detail.
- Consider buying the rights to bids that do not win so as to be able to integrate their best ideas in your final design.
- Consider tendering for framework contracts (also referred to as master contract) with A&E and key trades once you are in operation. This allows you to work with the same partners over a number of years, thereby accumulating knowledge and experience. It also enables you to turn around infrastructure upgrade projects faster because you do not need to tender for each one of these separately but can go directly to your framework partners.

---

[156] http://www.wbdg.org/pdfs/hpsb_guidance.pdf

4.8.3 *Building envelope*

There are two trains of thought on this subject. In order to determine which best fits your requirements consider the following:

- Does your facility have a PR role that will require it to be architecturally pleasing and designed to accommodate frequent visitors?
- Do your funding agencies prefer large and infrequent or smaller and incremental funding requests for infrastructure measures?
- How straightforward is it to get planning permission for a new building or building extension at your site?

Once you have your answers to the above, you should be able to choose which of the following two approaches best fits your situation:

- Big envelope: this approach entails building the biggest building envelope you can justify in order to be able to accommodate for changes within this framework. In this approach, the envelope is set but the infrastructure can still be designed and executed in a modular fashion. Within this, it is important to invest in good quality for long lifecycle items (e.g. the raised floor) and include options for future developments in the design stage. Check that your site will be able to accommodate your future requirements in terms of electricity and water and check options for waste-heat re-use or free cooling.
- Modular envelope: this approach aims to provide the optimal structure for the known requirements that is sufficiently generic to be able to adapt to changes. Think of this as a box to which you attach the requisite services. Both the envelope and the infrastructure are modular.

4.8.4 *Raised floor*

Again, there are two lines of thought, which are: to have a raised floor or to build slab on grade (i.e. not have a raised floor). The choice depends on the situation and preferences. In both cases, the full surface of the machine room should be uninterrupted by structural elements in order to have full freedom when siting systems.

- Raised floor: Allows you to run connections to IT equipment under the floor and if desirable to separate certain types of connections (e.g. water and electricity). The crux lies in the choice of the depth of the raised floor and the load ratings for it. If you elect to have a raised floor, consider the following:

- Specify both a static and a point load rating.
- Make sure the ratings are coherent and aligned with those for the raised floor structure and supporting building structure.
- Invest in the highest rating you can afford, as this is more cost effective and less obstructive than a later upgrade.
- Raised floor depth:
  - Check for limitations regarding this in building codes and other regulations as they may impact your operational processes.
  - Within these limitations maximise the depth of your raised floor in order to be able to accommodate all the connections as well as airflow if needed.
  - Consider making the raised floor a full storey below the machine room. This will allow you to locate electrical and mechanical distribution infrastructure directly below the systems, thus shortening the connections, improving efficiency and reducing the surface needed for these installations in other parts of the building. It also improves the ability to deploy infrastructure in a modular way more comfortably and without impacting operations in the machine room. Be sure to check fire codes if you pursue this as this will result in a fire zone of considerable volume.
  - Ensure the raised floor structure can support the weight of cable trays and piping.
- Slab on grade: Once you specify the load ratings for your cement slab, you do not need to worry about them again. All connections are run overhead together with lighting, fire detection and possibly extinguishing systems. The crux lies in the coordination of these.

4.8.5 *Electrical infrastructure*

- Define a standard electrical distribution and connection type for non-HPC hardware that will be hosted in your machine room. This makes installation faster and allows you to move hardware around in the room more easily.
- HPC systems will mostly require 480V distribution.

- For HPC systems, build the main distribution up front. For the final distribution to the system you can either:
  - Reserve the build-out until you have the full requirements for the system and do this as part of the site preparation. Having framework contracts in place makes this considerably easier to turn around in the limited time available.
  - Build out the full distribution, in which case you will need to specify this in future systems procurements and require vendors to connect to them. This approach is unlikely to work for sites procuring first-of-a-kind systems.
- Keep distribution paths as short as possible.
- Integrate the ability to meter and monitor as much as possible. (See *Measuring and monitoring* for more detail)
- AC or DC: currently none of the sites use DC distribution but one site is preparing to test it. DC has efficiency advantages but does also come with an increased arc flash capability and you will require specially trained facility staff to work in this environment. Because DC distribution is not mainstream at this point the technology is more expensive than AC.
- Look at possibilities for using alternative / renewable energy sources.
- UPS: Evaluate the quality of your utility power supply in order to ascertain the frequency at which you may experience brownouts or blackouts. This evaluation, combined with your business case, should inform your decision to put in place a power protection strategy or integrate a future option for this in your design.
- Make sure the utility company has sufficient power available for your planned capacity and be aware that load swings may adversely impact the utility.
- Earthing and electromagnetic compatibility (EMC) are of paramount importance in any data centre.

4.8.6 *Cooling*

Most sites have now integrated liquid cooling in their design. However, in most cases it remains necessary to maintain a mixture of cooling technologies for an HPC site. When selecting them, the following should be considered:

- What are the environmental conditions you want to provide in the machine room that will allow you to accommodate a wide range of vendor technology?
- Can you re-use your waste heat or make it available to a third party for use now or in future? Is waste-heat re-use a design driver and does it require you to use specific temperature ranges for your cooling? Note that this will impact your design choices as well as your future hardware choices.
- Evaluate and be aware of the trade-off between the savings of higher coolant temperatures vs. the impact this may have on the compute performance of your systems.

No matter which cooling technology(ies) are chosen, some basics will apply:

- Separate hot and cold.
- Use pumps with Variable Frequency Drives (VFD's) so you can accommodate different levels of operation efficiently.
- Size pumps to match operational range in order to avoid issues with cavitation.
- Define a standard cooling distribution for your non–HPC equipment (e.g. hot and cold aisles, in-row cooling, rear door heat exchangers). This makes installation faster and allows you to move hardware around in the room more easily.
- Keep distribution paths as short as possible.
- For HPC systems, build the main distribution up front. For the final distribution to the system you can either:
    - Reserve the build-out until the full requirements for the system are received and do this as part of site preparation. Having framework contracts in place makes this considerably easier to turn around in the limited time.
    - Build out the full distribution, in which case you will need to define this in future systems procurements and require vendors to connect to them. This approach is unlikely to work for sites procuring first of a kind systems.
- Commissioning for a large capacity cooling plant in the absence of the heat load from the IT is non-trivial and should be considered from the outset. Two possible approaches are:
    - Point commissioning of flow rates and temperatures with final testing and acceptance delayed to when the load is installed.

  - Use of district hot water distribution to simulate heat load during commissioning phase.
- For liquid and hybrid cooling you will also want to consider the following:
  - Define a temperature differential (ΔT) and a pressure differential (Δp) that you expect to see across your distribution as they will impact your overall cooling capacity.
  - Consider having cooling loops at different temperature ranges to be able to accommodate a wider range of hardware. Run in cascade formation for improved energy efficiency.
  - Monitor and control water chemistry – even when not required by a system vendor.
  - Divide different cooling loops (external, internal, IT) with heat exchangers so as to be able to separate different water chemistry requirements as well as avoid undesirable reactions between different materials used. This also reduces the amount of water in case of leakage or need for chemical treatment. This setup will facilitate running different temperature loops in a cascade formation.
  - Be aware of the weight that liquid cooling will add to your raised floor when you are defining the criteria for the same.
  - Due to the higher densities of liquid-cooled systems there is less temperature buffer. For this reason, where systems are on protected power the cooling *distribution* should have N+1 pumps and be on protected power so as to ensure continued circulation and avoid hardware damage.

4.8.7 *Fire protection*

Fire protection is an important topic that will be subject to extensive regulations due to the amount of power present in these facilities.

- Invest in a Very Early Smoke Detection Apparatus (VESDA) system.
- Fire suppression: there are different directions of thought on this topic. The best choice for your site will depend on your preferences, budget and priorities.
  - In the U.S., the predominant choice is for either pre-action dry piping or wet-pipe systems.
  - In Europe, numerous sites use gas or water mist.

4.8.8 *Measuring and monitoring*

- Invest in the ability to monitor and measure as much as possible as it will not only help with preventative maintenance (see next section for more detail) but also allow you to do such things as accurately measure TCO and energy efficiency. You cannot improve things that you cannot measure.
- Consider investing in a Data Centre Information Management (DCIM) system.

4.8.9 *Once in operation*

- Measure and monitor as much as possible as this will facilitate understanding what the "normal" situation of your infrastructure is, thereby enabling you to detect failures before they arise and execute preventative maintenance.
- Following commissioning spend time and effort optimising the infrastructure for maximum efficiency. This can noticeably reduce operating costs.
- Ensure that the requisite operational procedures are in place in order to avoid outages due to human error.
- When tendering for an HPC system, integrate the requirements of applications, systems and facilities to ensure procurement of a system best fitted to all requirements or that at least fits within the boundary conditions.

5 Conclusion

5.1 Summary of findings

The work has shown that data centre standards, with the exception of the ASHRAE Datacom Series, are not well known within the HPC community. Although the design requirements of HPC are not entirely covered by existing standards, most recent publications have closed the gap and the standards can be a useful base framework to build and extend upon when designing an HPC data centre.

Since HPC sites work at the forefront of technology, they frequently gain the experience on which standards will later be based. For this reason, HPC will by definition always be one step ahead of any standard one may formulate. Nonetheless, there is great value in sharing experiences gained within the peer community and the information collected in the interviews run for this study has allowed this first compilation of best practices to be put together.

5.2 Limitations and Implications

The best practices collected are by no means complete and exhaustive, but should be seen as a basis upon which the community can build. They represent the opinions of the sites in the sample, based on their experiences. At times approaches between sites differ fundamentally and in such cases both directions of thought are represented in the compilation. Any errors are the responsibility of the author alone and it should be noted that the next disruptive change in technology might well require a substantial review of any of the best practices in the compilation.

5.3 Conclusion

Defining design criteria for a data centre with a life expectancy of several decades, that must be able to accommodate multiple generations of HPC systems, is a complex task.

It requires a combination of engineering expertise, great curiosity and educated guesses about future IT technology developments to allow the derivation of clear design requirement for the A&E partner whilst maintaining a maximum amount of flexibility to adapt to future changes.

The goal of this study was to provide an overview of relevant existing data centre standards and investigate to what extent they are known and applied by public research HPC data centres. The differences in requirements between HPC and the wider data centre industry were to be investigated. Through interviews, the challenges and approaches of HPC sites in defining design criteria and future proofing their centres were to be studied.

The desired outcome of this study was to provide managers of future HPC data centre construction projects with both an overview of existing standards and a collection of best practices applied by the community. Having set out on a quest for a design standard for HPC data centres, the elaboration of this study has shown the formulation of best practices to be a more appropriate and useful goal.

Best practices for HPC data centres must continue to evolve with the technology they host. The current compilation can therefore only be a starting point for the community to build on. Further work in this context would be desirable and most useful to the community.

Bibliography

Alger, D. (2005): Build the best data center facility for your business. Indianapolis, Ind. 2005

American Society of Heating Refrigerating and Air-Conditioning Engineers (2007): Structural and vibration guidelines for datacom equipment centers. Atlanta, GA, 2007

American Society of Heating Refrigerating and Air-Conditioning Engineers (2008a): Best practices for datacom facility energy efficiency. Atlanta, GA, 2008

American Society of Heating Refrigerating and Air-Conditioning Engineers (2009a): Design considerations for datacom equipment centers. Atlanta, GA, 2009

American Society of Heating Refrigerating and Air-Conditioning Engineers (2009b): Real-time energy consumption measurements in data centers. Atlanta, GA, 2009

American Society of Heating Refrigerating and Air-Conditioning Engineers (2011): Green tips for data centers. Atlanta, GA, 2011

American Society of Heating Refrigerating and Air-Conditioning Engineers (2012): Datacom equipment power trends and cooling applications. Atlanta, GA, 2012

American Society of Heating Refrigerating and Air-Conditioning Engineers (2013): PUE : a comprehensive examination of the metric. Atlanta, GA, 2013

American Society of Heating Refrigerating and Air-Conditioning Engineers (2014a): Liquid cooling guidelines for datacom equipment centers. Atlanta, GA, 2014

American Society of Heating Refrigerating and Air-Conditioning Engineers (2014b): Particulate and gaseous contamination in datacom environments. Atlanta, GA, 2014

American Society of Heating Refrigerating and Air-Conditioning Engineers (2015a): Thermal guidelines for data processing environments. Atlanta, 2015

American Society of Heating Refrigerating and Air-Conditioning Engineers (2015b): Server Efficiency - Metrics for Computer Servers and Storage. Atlanta, GA, 2015

American Society of Heating Refrigerating and Air-Conditioning Engineers. TC 9.9 Mission Critical Facilities Technology Spaces and Electronic Equipment (2008b): High Density Data Centers. Atlanta, GA, 2008

Anixter (2007): *Data Center Intrastructure Resource Guide*. Anixter. Retrieved from http://www.anixter.com/datacenterguide [Accessed on 07.05.2016]

ASHRAE (Firm) (2016): IT equipment design impact on data center solutions. Atlanta, GA, 2016

Bartels, A. (2011, August): [INFOGRAPHIC] Data Center Evolution: 1960 to 2000. *Rackspace*. Retrieved from http://blog.rackspace.com/datacenter-evolution-1960-to-2000/ [Accessed on 09.05.2016]

Beaty, D. L. (2013): *Why Datacom?* In ASHRAE Journal March 2013 (p. 74 – 78).

Bernhardt, M. (2010): *Exascale: The Beginning of the Great HPC Disruption*. The Exascale Report. Retrieved from http://insidehpc.com/exascale-the-beginning-of-the-great-hpc-disruption/ [Accessed on 27.06.2016]

BICSI (2014): ANSI/BICS 002-2014, Data Center Design and Implementation Best Practices. Tampa, FL, 3rd edition 2014

Bilderbeek, P./ Coors, L. (2016, April): The need for a new data center design standard. *DCD Intelligence.* Retrieved from http://www.datacenterdynamics.com/research/open-research/the-need-for-a-new-data-center-design-standard/95999.article [Accessed on 11.05.2016]

British Standards Institution (2012): EN 50600-1 Information Technology - Data centre facilities and infrastructures - Part 1: General concepts. London, 2012

British Standards Institution (2014a): Information Technology. Data centre facilities and infrastructures. Building construction. London, 2014

British Standards Institution (2014b): Information technology. Data centre facilities and infrastructures. Power distribution. London, 2014

British Standards Institution (2014c): Information technology. Data centre facilities and infrastructures. Environmental control. London, 2014

British Standards Institution (2014d): Information technology. Data centre facilities and infrastructures. Management and operational information. London, 2014

BSI (2015): Information technology. Data centre facilities and infrastructures. Telecommunications cabling infrastructure. London, 2015

BSI (2016): Information technology. Data centre facilities and infrastructures. Security Systems. London, 2016

Ceruzzi, P. E. (2012): Computing : a concise history. Cambridge, 2012

Coors, L. (2016, July): Data center standards need an update. *DatacenterDynamics*, Retrieved from http://www.datacenterdynamics.com/content-tracks/power-cooling/data-center-standards-need-an-update/96557.fullarticle [Accessed on 12.07.2016]

Demetriou, D. W./Kamath V./ Mahaney. H. (2015): Understanding the True Total Cost of Ownership of Water Cooling for Data Centers. *ASME 2015 International*

Technical conference and Ehibition on Packaging and Integration of Electronic and Photonic Microsystems InterPACK2015, San Francisco, CA, pp. V001T09A013; 9 pages.

Donnelly, C. (2016, April): CW@50: The changing face of the datacentre over the past 50 years. *ComputerWeekly*, Retrieved from http://www.computerweekly.com/feature/CW50-The-changing-face-of-the-datacentre [Accessed on 15.07.2016]

European Commission, D.-G. f. C. N., Content and Technology (2015): *High performance computing in the EU: progress on the implementation of the European HPC strategy: final report*. EU Bookshop. Retrieved from http://bookshop.europa.eu/en/take-two-pbKK0415496/

Ezell, J./ Atkinson, R. D. (2016): *The Vital Importance of High-Performance Computing to U.S. Competitiveness*. I. T. I. F. Retrieved from http://www.itif.com [Accessed on 27.06.2016]

Fosdick, L. D. (1996): An Introduction to high-performance scientific computing. Cambridge, Mass., 1996

Glanz, J. (2012): Power, Pollution and the Internet. *New York Times* 22.09.2012. Retrieved from http://www.nytimes.com/2012/09/23/technology/data-centers-waste-vast-amounts-of-energy-belying-industry-image.html?_r=0 [Accessed on 12.06.2016]

ITRS (2015): *The International Technology Roadmap for Semiconductors 2.0: 2015*. www.semiconductors.org. Retrieved from http://www.semiconductors.org/main/2015_international_technology_roadmap_for_semiconductors_itrs/ [Accessed on 15.07.2016]

Joseph, E. C./ Conway, S./ Sorensen, R. (2014): *IDC HPC ROI Research Update: Economic Models For Financial ROI and Innvoation from HPC Investments* [presentation]. hpcuserforum.com. Retrieved from https://hpcuserforum.com/presentations/coloradosept2015/PublicIDCDOEROIResearchUpdate8.19.2015.pdf. [Accessed on 21.05.2016]

Joseph, E. C./ Dekate, Ch./ Conway, S. (2014): *Real-World Examples of Supercomputers Used for Economic and Societal Benefits: A Prelude to What the Exascale Era Can Provide*. CASC. Retrieved from
http://casc.org/wp-content/uploads/2014/07/IDCReportRealWorldExamplesOfBenefitsOfSupercomputers.pdf. [Accessed on 14.06.2016]

Kaplan, J. M./ Forest, W./ Kindler, N. (2008): *Revolutionizing Data Center Energy Efficiency*. The Sallan Foundation. Retrieved from http://www.sallan.org/pdf-docs/McKinsey_Data_Center_Efficiency.pdf [Accessed on 23.05.2016]

Kaufmann, W. J./ L. L. Smarr (1993): Supercomputing and the transformation of science. New York 1993

Koomey, J. (2011): *Growth in data centre electricity use 2005 to 2010*. Oakland: Analytics Press. Retrieved from http://www.analyticspress.com/datacenters.html [Accessed on 06.06.2016]

Jones Lang LaSalle (2016): *Data Center Perspective North America*. (Winter 2016). Retrieved from http://www.us.jll.com/united-states/en-us/research/6651/US-North-America-Data-Center-Perspective-2016-JLL [Accessed on 15.05.2016]

Jones Lang LaSalle (2011): *Perspectives on Data Centre Solutions*. Retrieved from http://www.us.jll.com/united-states/en-us/research/ [Accessed on 15.05.2016]

McCarthy, K. J. S. (2011): High-performance Computing (HPC) Is Likely to Become Mainstream. *Facilitiesnet.* Retrived from http://www.facilitiesnet.com/datacenters/article/High-performance-Computing-HPC-Is-Likely-to-Become-Mainstream-Facilities-Management-Data-Centers-Feature--12242 [Accessed on 13.06.2016]

Miller, R. (2012, October): Koomey: Data Center Energy Use is Moderating. *Data Center Knowledge*. Retrieved from http://www.datacenterknowledge.com/archives/2012/10/15/koomey-data-center-energy-use-moderating/ [Accessed on 21.06.2016]

Mitchell, R. L. (2010, January): Data center density hits the wall. *Computerworld,* Retrieved from http://www.computerworld.com/article/2522601/it-management/data-center-density-hits-the-wall.html?page=2 2016. [Accessed on 14.05.2016]

Murray, C. J. (1997): The supermen : the story of Seymour Cray and the technical wizards behind the supercomputer, New York, (1997)

Nispel, M. (2013, April): What Makes The Enterprise Data Center Different? *Enterasys Networks,* Retrieved from http://www.extremenetworks.com/what-makes-the-enterprise-data-center-different/ [Accessed on 28.07.2016]

Northwetern University, (2012): Return of the Large Data Center. Retrieved from http://www.it.northwestern.edu/bin/docs/dc_architecture.ppt; [Accessed on 29.07.2016]

Nutt, A. (2008, September): History of the Data Centre. *Articlesbase.com*. Retrieved from http://www.articlesbase.com/technology-articles/history-of-the-data-centre-556616.html [Accessed on 25.07.2016]

PRACE (2016): Preparing Europe for the exascale generation. In *PRACE Annual Report 2015 (*p. 27 – 29). Retrieved from http://www.prace-ri.eu/IMG/pdf/Prace-annual-report2015_lowres.pdf [Accessed on 17.05.2016]

Pricket Morgan, T. (2015, December): Looking ahead in the datacentre with Intel. *TheNextPlatform*, Retrieved from www.nextplatform.com/2015/12/10/looking-ahead-in-the-datacenter-with-intel [Accessed on 10.07.2016]

Restivo, K. (2015): In: *Data Center Evolution: From Cabinets to Colocation, Cloud Computing and Beyond* [presentation]. Retrieved from http://www.slideshare.net/Panduit/data-center-evolution-from-cabinets-to-colocation-cloud-computing-and-beyond [accessed on 17.05.2016]

RIKEN (2016): Flagship 2020 Project [presentation] Retrieved from http://www.aics.riken.jp/fs2020p/en/ [Accessed on 27.07.2016]

Saarinen, J. (2016): EU Consortium Gathers Core Capabilities to Build Exascale HPC Prototype. Retrieved from http://www.itnews.com.au/news/europe-to-build-exascale-supercomputer-prototype-414471 [Accessed on 01.08.2016]

Telecommunications Infrastructure Association (2014): Telecommunications Infrastructure Standard for Data Centers Telecommunications Industry Association. www.tia.org. 2014

The Council on Competitiveness (2014): *Solve: The Exascale Effect: the Benefits of Supercomputing Investment for U.S. Industry*. C. o. Competitiveness. http://www.compete.org/, Coucil of Competitiveness. Retrieved from http://www.compete.org/reports/all/2695 [Accessed on 19.06.2016]

The Council on Competitiveness (2015): *Modeling, Simulation and Analysis, and High Performance Computing*. NDEMC. T. C. o. Competitiveness. Washington, D.C. Retrieved from http://www.compete.org/reports/all/2938 [Accessed on 03.06.2016]

The Data Center Journal (2016, February): *Indutry Outlook: Preparing for HPC*. Retrieved from ww.datacenterjournal.com/industry-outlook-preparing-for-hpc/ [Accessed on 23.05.2016]

The White House (2015): *Executive Order - Creating a National Strategic Computing Initiative*. Retrieved from https://www.thewhitehouse.gov/the-press-office/2015/07/29/executive-order-creating-national-strategic-computing-intitiative [Accessed on 23.05.2016]

Trader, T. (2016, June): China Debuts 93-Petaflops "Sunway" with Homegrown Processors. *HPCWire*. Retrieved from https://www.hpcwire.com/2016/06/19/china-125-petaflops-sunway/ [Accessed on 20.06.2016]

Turner, W. P./ Seader J. H./ Renaud V./ K. G. Brill (2008): White Paper: Tier Classifications Define Site Infrastructure Performance, retrieved from: www.uptimeinstitute.org [Accessed on 13.05.2016]

United States Congress, (1991): *High-Performance Computing Act of 1991. Public Law No: 102-194.* Washington, DC. Retrieved from https://www.nitrd.gov/congressional/laws/102-194.pdf [Accessed on 09.06.2016]

United States Congress House Committee on Science (2003): Supercomputing: is the U.S. on the right path?: hearing before the Committee on Science, House of Representatives, One Hundred Eighth Congress, first session, July 16, 2003, Washington, 2003

United States Environmental Protection Agency (2007): Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431, Washington (2007) [Accessed on 28.06.2016]

Uptime Institute (2012): *Data Centre Site Infrastructure Tier Standard: Topology*. Uptime Institute. Retrieved from http://www.uptimeinstitute.com [Accessed on 13.05.2016]

Uptime Institute (2013): *Data Center Site Infrastructure Tier Standard: Operational Sustainability*. Uptime Institute. Retrieved from http://www.uptimeinstitute.com [Accessed on 13.05.2016]

Uptime Institute (2014): *2014 Data Center Industry Survey*. Uptime Institute. Retrieved from https://journal.uptimeinstitute.com/2014-data-center-industry-survey/ [Accessed on 13.05.2016]

Vetter, J. S. (2013): Contemporary high performance computing : from petascale toward exascale. Boca Raton, 2013

Internet resources

https://www.ashrae.org/

https://www.bicsi.org/

https://www.nitrd.gov/

http://www.thegreengrid.org/

https://uptimeinstitute.com/

http://www.delloro.com

http://www.opencompute.org/about/

http://www.thegreengrid.org

https://www.top500.org/

http://www.aics.riken.jp/

http://www.exanest.eu

http://beyondcmos.ornl.gov

http://www.nersc.gov/

http://science.energy.gov/

https://www.lrz.de/

https://eehpcwg.llnl.gov/

www.bsigroup.com/

http://www.iso.org/

http://energy.gov/about-national-labs

http://www.prace-ri.eu

http://www.nrel.gov

https://www.pawsey.org.au

http://nci.org.au

http://www.riken.jp/en/

http://www.wbdg.org/

http://iet.jrc.ec.europa.eu/

Appendix 1 – Computer performance by orders of magnitude

	Number of FLOPS	Description
Deca-scale	10	Ten
Hecto-scale	10^2	Hundred
Kilo-sacle	10^3	Thousand
Mega-scale	10^6	Million
Giga-scale	10^9	Milliard / Thousand Million
Tera-scale[1]	10^{12}	Billion
Petascale[2]	10^{15}	Billiard / Thousand Billion
Exascale[3]	10^{18}	Trillion
Zetta-scale	10^{21}	Trilliard / Thousand Trillion
Yotta-scale	10^{24}	Quadrillion

[1] First achieved in 1997 on Intel ASCI Red supercomputer

[2] First achieved in 2009 on IBM Roadrunner supercomputer

[3] Current projections expect this performance to be reached around 2020.

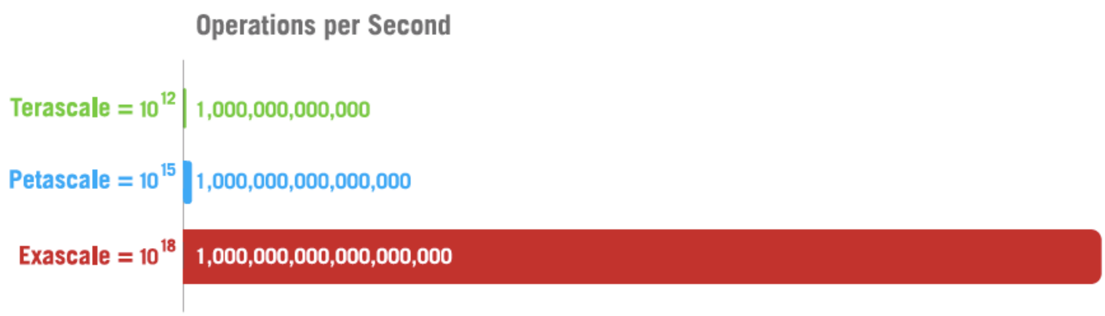

Figure 11: Computer performance – Ezell/ Atkinson 2016 [157]

Note: the step from one scale to the next represents a 1000-fold increase

[157] Ezell/ Atkinson 2016, p. 7

Appendix 2 – Data Centre metrics

Power Usage Efficiency:

As defined by The Green Grid

$$PUE = \frac{Total\ Facility\ Energy}{IT\ Equipment\ Energy}$$

Data Centre Infrastructure Efficiency:

As defined by The Green Grid

$$DCiE = \frac{IT\ Equipment\ Energy}{Total\ Facility\ Energy} * 100\%$$

Corporate Average Datacentre Efficiency:

As defined by McKinsey and the Uptime Institute

$$CADE = Facility\ Efficiency * Asset\ Efficiency$$

Appendix 3 – overview of national HPC strategies by country

Country	**HPC Strategy / Program and Description**	**Investment Level**
United States	National Strategic Computing Initiative (NSCI)	$320 million/year
China	13th Five-Year Development Plan (Develop Multiple Exascale Systems)	$200 million/year (for next five years)
Japan	Flagship2020 Program	$200 million/year (for next five years)
European Union	ExaNeSt; PRACE; ETP4HPC	$1.1 billion total allocated through 2020 (annual allocations N/A)
India	National Supercomputing Mission	$140 million/ year (for five years from 2016 – 2020)
South Korea	National Supercomputing Act	$20 million/year (for five years from 2016 – 2020)
Russia	HPC Focus of Medvedev Modernisation Program	N/A

Figure 12: Summary of National HPC Strategies by Country based on Seager (2010) [158]

[158] as quoted in Ezell/ Atkinson 2016, p. 41

Appendix 4 – Detailed overview of the ASHRAE Datacom Series

Thermal Guidelines for Data Processing Environments first appeared in 2004 and the current fourth edition was published in 2015. It was the first document to provide complete information regarding temperature and humidity requirements for IT equipment. It is universally applied today and endorsed by the major IT vendors. The original range of recommended temperatures and humidity ranges has been expanded since the 1st edition and new environmental classes for data centres defined, that allow the use of free-cooling techniques in most of the world's climates. The document provides data centre operators with the information needed to optimize the energy efficiency of their data centre by allowing larger ranges of temperature and humidity without adversely impacting the IT equipment and performance. ASHRAE defines 5 envelopes for air-cooled environments. The *recommended* envelope reflects the environmental conditions that facilities should be designed to achieve under normal circumstances. The four allowable envelopes define extended environmental conditions for which vendors test the functionality of their equipment. This standard has been expanded to comprise definitions of environmental classes for liquid-cooling. [159]

Datacom Equipment Power Trends and Cooling Applications was first published in 2005 and updated in 2012. It discusses the implications of trends in IT power density and cooling technology on the design of a data centre. It provides an overview of the power density trends of the various individual components as well as outlining approaches for determining floor space, power and cooling requirements for a data centre project. The guide provides designers and owners the information needed to more adequately plan for future technology developments. It also provides an overview of air- and liquid-cooled systems aimed at supporting future loads.[160]

Design Considerations for Datacom Equipment Centres first appeared in 2005. The current and second edition was published in 2009. The first part of this publication covers the basic design criteria for data centre facilities including Heating Ventilation and Air Conditioning (HVAC) loads, cooling systems, air distribution and liquid cooling. The second part of the publication provides supplementary information on ancil-

---

[159] ASHRAE 2015a, p. xi – xii, 1 - 8
[160] ASHRAE 2012, p. ix, 1 – 14

lary spaces, contamination, acoustics, structural and seismic design, fire suppression, commissioning, availability and redundancy, and energy efficiency.[161]

Liquid Cooling Guidelines for Datacom Equipment Centres was first published in 2006 and the current second edition appeared in 2014. This document addresses difficulties found in providing adequate cooling to IT equipment via the different existing air-cooled solutions due to increasing power densities. IT provides guidelines for liquid cooling strategies and distribution topologies. Requirements for liquid cooling infrastructure including chilled-water piping, electrical power sources and connections, monitoring, reliability and availability as well as commissioning are discussed. ITE vendors frequently refer to the water quality requirements described in this publication.[162]

Structural and Vibration Guidelines for Datacom Equipment Centers addresses the increasing requirements of datacom facilities in terms of structure and vibration performance. This guideline provides requirements to keep a datacom facility up and running even during and after more extreme natural or man-made events. These requirements go beyond those set out by building codes that are focused on life safety issues.[163]

Best Practices for Datacom Facility Energy Efficiency was first published in 2008. The second and current edition appeared in 2009. This guideline provides detailed information on how to minimise life-cycle cost of a data centre and maximise energy efficiency by applying sustainable design approaches. This is a topic of great importance for data centres given that their energy usage is significantly higher than that of office buildings and their round the clock operation means their also have three times the annual operating hours of other commercial properties and the environmental conditions within a facility have a strong impact on energy consumption. The increasing power density of IT equipment can have a significant impact on the Total Cost of Ownership (TCO). The book covers the topics of environmental criteria, mechanical equipment and systems, economisers, airflow distribution, controls and ener-

---

[161] ASHRAE 2009a, p. 1 - 4
[162] ASHRAE 2014a, p. 1 - 8
[163] ASHRAE 2007, p. 1 - 8

gy management, electrical distribution equipment, datacom equipment efficiency, liquid cooling, total cost of ownership, emerging technologies and future research.[164]

High Density Data Centres – Case Studies and Best Practices was published in 2008 motivated by the difficulties existing data centres were encountering in providing adequate cooling within a high density environment. The book discusses seven ventilation schemes that are frequently applied in the industry. It also provides a summary on best practices for new data centre building designs, accommodating future growth, raised-access and non-raised access floor designs, localised rack cooling and energy management and efficiency.[165]

Particulate and Gaseous Contamination in Datacom Environments was first published in 2009. The current and most recent edition appeared in 2013. The original publication was motivated by an increase in hardware failures in the mid-2000s following changes in product design to accommodate the EU Restriction of Hazardous Substances (RoHS). The first edition concentrated on monitoring, preventing and controlling particulate and gaseous contaminations in data centres. The second edition provides an update to all chapters based on knowledge gained since the original publication.[166]

Real-Time Energy Consumption Measurements in Data Centres was published in 2010. The motivation for this guideline was rooted in the increasingly intense discussion around energy consumption in data centres and their impact on electricity supply. The EPA report to the U.S. Congress on server and data center energy efficiency in 2007 showed that data centres in the United States were consuming 1.5% of the total electricity consumption and that this figure was would rise to 2.9% by 2011. The U.S. Department of Energy and The Green Grid set out to reduce annual energy consumption by implementing real-time energy efficiency measures that would allow data centres to measure and improve their energy consumption. The energy efficiency metrics PUE and Data Center infrastructure Efficiency (DCiE) formulated by The Green Grid in 2007 are discussed.[167]

---

[164] ASHRAE 2008a, p. 1 - 11
[165] ASHRAE 2008b, p. 1 - 6
[166] ASHRAE 2014b, p. v – vi, 1 - 11
[167] ASHRAE 2009b, p. 3 – 13

Green Tips for Data Centers, published in 2011, discusses techniques for optimizing energy efficiency and carbon footprint of existing data centres that can be easily implemented at relatively low cost. The guide covers optimizations divided into the following chapters: energy management, environmental conditions, air management, the cooling plant, IT power distribution, lighting and IT equipment.[168]
PUE™: A Comprehensive Examination of the Metric is written in collaboration with The Green Grid and was published in 2013. Following the publication of the PUE metric in 2007 it was rapidly adopted by the industry. The PUE metric expresses the ratio of total facilities energy to IT equipment energy. The more total power is actually used by the IT equipment the lower the PUE and the more efficient the data centre. The book aims to provide executives with a high level understanding of the concepts surrounding the PUE metric at the same time as providing those implementing and reporting data centre metrics with in-depth application knowledge and resources.[169]

Server Efficiency – Metrics for Computer Servers and Storage was published in 2015. It provides an in depth description of the tools available to quantify the energy consumption of IT equipment with the aim of providing managers with the information needed to relate product requirements to their specific environments. It introduces the metric of performance per watt for IT equipment.[170]

IT Equipment Design Impact on Data Center Solutions was published in 2016. It provides detailed information about cooling design and thermal management and trends for IT equipment and their impact data centre operations. This publication is motivated by the realisation that the needs of the digital economy and the related costs are growing faster than the hardware it depends on is developing. Short and disruptive hardware and software changes are proving increasingly challenging for data centres that are traditionally planned with a life expectancy of ten to twenty years and show that the data centre of the future will need to be flexible, scalable and adaptable in order to cope with this reality.[171]

---

[168] ASHRAE 2011, p. ix-x, 1 - 9
[169] ASHRAE 2013, p. 1 - 8
[170] ASHRAE 2015b, p. xv – xviii, 1 - 13
[171] ASHRAE 2016, p. ix – xi, 1

Appendix 5 – Interview introduction and questions

Part I: Design standards

Q 1: Are you familiar with any of the following design standards?

- TIA-942 (Telecommunications Infrastructure Standard for Data Centers)
- ANSI/BICSI 002-2014 Data Center Design and Implementation Best Practices (American National Standards Institute)
- EN 50600 Series: Information technology – Data centre facilities and infrastructures (European Committee for Electrotechnical Standardization)
- ASHRAE Datacom Series (American Society for Heating, Refrigeration and Air-Conditioning Engineers)
- Uptime Institute Tier Classification System

Q 2: Have you used these standards to help inform design decisions for your data centre?

Q 3: Did your design team use any of the above as a reference point for your data centre project?

Q 4: Did you use any other design standard that I have not listed here?

Part II: Commonly observed attributes found in non-HPC data centres

Q 5: The information in the following table is inspired by the Uptime Institute White Paper: Tier Classifications Define Site Infrastructure Performance 5th edition, (2008) with some additions. How do they compare to your reality / requirements?

	Data centre industry	**Your site**
Building type (tenant / stand alone)	Tier I + II: tenant Tier III + IV: stand-alone	
Staffing shifts (none/ 1/ 1-2/ 24hrs)	Depending on Tier level	
Load per cabinet	1 - 15kW	
Total power available	< 1MW - ~40MW	
Surface of machine room	< 50 m^2 - > 25’000m^2	
Ratio of support space to raised floor	Tier I + II: 20 – 30% Tier III + IV: 80 – 100+%	
Raised floor height	30 – 110 cm	

	(12 – 42 in)	
Raised floor rating for uniform / static loads	730 - 1220 kg/m^2 (150 - 250 psf.)	
Raised floor rating for concentrated/ point loads	567 kg – 680 kg (1250 lbs.– 1500 lbs.)	
Amperage at rack	16 A, 32 A	
Utility voltage	Tier I + II: 208, 480 Tier II + IV: 12 – 15kV	
Equipment on UPS	All	
Cooling medium (air/ liquid/ hybrid)*	air	
Time to plan and build (in months)	Tier I + II: 3 – 6 Tier III + IV: 15 – 30	
Expected life of data centre	10 – 20 years	

* Based on the medium that removes the heat at its source.

Q 6: For those attributes where you note a significant difference between your reality and non-HPC data centres, can you comment on what was the driver behind your design choice and why this may be so different?

Part III: Design criteria that were challenging to define

Q 7: During your design project, which criteria did you find most challenging to define and why? (Business case, structural, power, cooling, future-proofing, energy efficiency, other)

Q 8: What approach did you use to define these criteria? Note: I am particularly interested in understanding the thought process that led to the decisions as these could help advise future project managers as to possible approaches

Q 9: Do you have SLA's that drive your design requirements?

Part IV: Future-proofing your data centre

Q 10: Did you specifically integrate options or possibilities in your design to allow the building to adapt to future changes in requirements? (Future-proofing)

Q 11: Can you give me concrete examples of such options?

Q 12: What challenges do you see coming down the pipeline that you expect will need addressing in the next 2 -5 years?

Part V: Other

Q 13: Is there any other important advice you would give someone setting out on a design project for a public research HPC data centre?

Q 14: Are there any aspects that you feel should be addressed and have not been covered by the questions?

Appendix 6 – List of public research HPC centres comprised in sample

Site Name	*Country*	*Rank Top500**	*Affilia-tion*
Argonne National Laboratory	USA	5	DOE
Barcelona Supercomputing Center (BSC)	ESP	93	PRACE
Commissariat a l'Energie Atomique, Très Grand Centre de calcul du (CEA - TGCC)	FRA	53, 74	PRACE
Centre Informatique National de l'Enseignement Supérieur (CINES)	FRA	44	PRACE
Centre national de la recherche scientifique - Institut du développement et des ressources en informatique scientifique (CNRS/ IDRIS)	FRA	70	PRACE
Center for Scientific Computing (CSC)	FIN	59	PRACE
Centro di supercalcolo, Consorzio di università (CINECA)	ITA	37	PRACE
Forschungszentrum Jülich (FZJ)	DEU	11, 50	PRACE
Höchstleistungsrechenzentrum Stuttgart (HLRS)	DEU	8	PRACE
IT4Innovations National Supercomputing Center	CZE	48	PRACE
KTH – Royal Institute of Technology	SWE	52	PRACE
Lawrence Berkeley National Laboratory – National Energy Research Science Center (LBL – NERSC)	USA	40, 72	DOE
Lawrence Livermore National Laboratory (LLNL)	USA	3, 12	DOE
Leibniz Rechenzentrum (LRZ)	DEU	23, 24	PRACE
Los Alamos National Laboratory (LANL)	USA	6, 65	DOE
National Computational Infrastructure, Australian National University (NCI)	AUS	86	Other
National Renewable Energy Laboratory (NREL)	USA	N/A	DOE
Oak Ridge National Laboratory (ORNL)	USA	2	DOE
Pacific Northwest National Laboratory (PNNL)	USA	27	DOE
Pawsey Supercomputing Centre	AUS	68	Other
Poznan Supercomputing and Networking Center	POL	80	PRACE
RIKEN Advanced Institute for Computational Science (AICS)	JPN	4, 84	Other
Sandia National Laboratory (SNL)	USA	6, 65	DOE
Science and Technology Facilities Council – Daresbury laboratory (STFC)	GBR	49	PRACE
SURFsara	NLD	69	PRACE
Swiss National Supercomputing Centre (CSCS)	CHE	7, 92	PRACE
University of Edinburgh (EPCC)	GBR	41, 81	PRACE

Table 6: List of public research HPC centres comprised in sample

* November 2015 list

Note: The National Renewable Energy Laboratory (NREL) is not listed in the Top500 list despite its compute capacity of 1.19 petaflops.

Appendix 7 – Common attributes found in data centres

	Tier I	Tier II	Tier III	Tier IV
Building Type	Tenant	Tenant	Stand-alone	Stand-alone
Staffing shifts Staff/shift	None None	1 Shift 1/Shift	1+Shifts 1-2/Shift	"24 by Forever" 2+/Shift
Useable for Critical Load	100% N	100% N	90% N	90% N
Initial Build-out kW per Cabinet (typical)	<1kW	1-2 kW	1-2 kW	1-3 kW
Ultimate kW per Cabinet (typical)	<1 kW	1-2 kW	>3 kW[†,‡]	>4 kW[*,†]
Support Space to Raised-Floor Ratio	20%	30%	80-90+%	100+%
Raised-Floor Height (typical)	12 inches	18 inches	30-36 inches	30-42 inches
Floor Loading lbs/ft (typical)	85	100	150	150+
Utility Voltage (typical)	208, 480	208, 480	12-15 kV	12-15 kV
Single Points-of-Failure	Many + Human Error	Many + Human Error	Some + Human Error	Fire, EPO + Some Human Error
Representative Planned Maintenance Shut Downs	2 Annual Events at 12 Hours Each	3 Events Over 2 Years at 12 Hours Each	None Required	None Required
Representative Site Failures	6 Failures Over 5 Years	1 Failure Every Year	1 Failure Every 2.5 Years	1 Failure Every 5 Years
Annual Site-Caused, End-User Downtime (based on field data)	28.8 hours	22.0 hours	1.6 hours	0.8 hours
Resulting End-User Availability Based on Site-Caused Downtime	99.67%	99.75%	99.98%	99.99%
Typical Months to Plan and Construct	3	3-6	15-20	15-30
First Deployed	1965	1970	1985	1995

[*] 3.5 kW per cabinet over large areas is acceptable for traditional air-cooling designs.

[†] Higher kW/cabinet densities require a greater ratio of support space to computer floor (at least 1:1 at 3 kW/cabinet, 2:1 at 6 kW/cabinet, 3:1 at 9 kW/cabinet, etc.) Generally, deeper raised floors are required for higher densities.

[‡] Most sites have a difficult time maintaining stable and predictable cooling for racks in the 1-2 kW range. Major process improvements are required before entertaining rack densities above 2 kW. See Institute white paper *How to Meet "24 by Forever" Cooling Demands of your Data Center*.

Figure 13: Common attributes found in data centres as published by Uptime Institute 2008[172]

[172] Uptime Institue 2008, p. 14

Declaration of honour

I hereby confirm on my honour that I personally prepared the present academic work on the topic of public research HPC data centres and personally carried out the activities directly involved with it. I also confirm that I have used no resources other than the ones declared. All formulations and concepts adopted literally or in their essential content from printed, unprinted or Internet sources have been cited according to the rules for academic work and identified by means of footnotes or other precise indications of source.

The support provided during the work, including significant assistance from my supervisor has been indicated in full.

The academic work has not been submitted to any other examination authority. The work is submitted in printed and electronic form. I confirm that the content of the digital version is completely identical to that of the printed version.

I am aware that a false declaration will have legal consequences.

Mezzovico, *22. August, 2016*]

[*Ladina Gilly*]