GREEN
DATA CENTRE
Technology Roadmap

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The Green Data Centre Technology Roadmap sets out a framework to improve data centre sustainability. The roadmap aims to reduce energy consumption and improve the energy efficiency of the constituent systems of a data centre – facilities and IT. It assesses and makes recommendations on potential directions for Research, Development & Demonstration (RD&D) to improve the energy efficiency of data centres in Singapore's context.

Most data centres in Singapore, many in the middle of their lifespan, were designed and constructed without sustainability and energy conservation in mind. Looking ahead, the Singapore data centre industry is expected to experience strong and sustained growth. Consequently, the Roadmap covers the green initiatives that span existing data centres and new data centres.

Energy consumption in data centres can be broadly categorised into two categories – energy consumption by IT systems and energy consumption by facility systems (i.e. supporting infrastructure). Three main areas examined in this roadmap are facility, information technology (IT) systems, and an integrated approach to design and deployment of data centres.

Most of the global effort to improve the energy efficiency of data centres has centred on facility systems. Of facility systems, cooling has received the most attention as it is generally the single largest energy overhead. Singapore's climate, with its year-round high temperatures and humidity, makes cooling particularly energy-intensive compared to other data centre hubs globally. The following technologies to improve the energy efficiency of facility systems are assessed:

- Direct liquid cooling
- Close-coupled refrigerant cooling
- Air and cooling management
- Passive cooling
- Free cooling (hardening of IT equipment)
- Power supply efficiency

Notwithstanding the importance of improving the energy efficiency of powering and cooling data centres, the current focal point for innovation is improving the energy performance of physical IT devices and software. Deficiencies in these areas provide opportunities for innovation that would greatly improve the sustainability of data centres. The following technologies to improve the energy efficiency of IT systems are assessed:

- Software power management
- Energy-aware workload allocation
- Dynamic provisioning
- Energy-aware networking
- Wireless data centres
- Memory type optimisation

An integrated approach to design and deployment is absent in the data centre sector. The energy performance of a data centre depends on the complex interplay between its IT systems and facility (i.e. non-IT) systems. However, the two are often deployed with little regard for one another. This silo-ed approach to design and deployment often results in wasteful over-provisioning. Technology vendors have offered Data Centre Infrastructure Management (DCIM) tools as a way for end users to address the silo-ed nature of data centre deployments. However, very often the DCIM solutions themselves remain limited to the domains of their respective vendors, restricting their use to only piecemeal optimisations. The Roadmap explores future directions in advanced DCIM to enable the integration and automation of the disparate systems of the data centre.

Owing to the mission critical nature of data centre operations, end user risk aversion is the largest barrier to implementing more effective energy strategies in data centres. It is recommended that the most effective way to address risk aversion and promote the adoption of new energy efficient technologies is collaboration between the public bodies responsible for energy efficiency and the industry. To this end, proof-of-concept demonstrations are essential if the adoption of new technologies is to be fast-tracked in Singapore.

1.INTRODUCTION

Singapore leads Southeast Asia as a data centre hub, accounting for more than 60% of the region's data centre market. Buoyed by political and economic stability, well-developed telecommunications infrastructure, the presence of a large number of multinational companies, and government initiatives such as the planned Data Centre Park, the data centre industry is expected to continue to experience strong growth.

The data centre industry in Singapore and many developed countries are subject to intense scrutiny due to the high proportion of energy costs in the context of overall cost of ownership.

To identify new technologies relevant to Singapore, the Infocomm Development Authority (IDA) of Singapore appointed i_3 Solutions Group to prepare the Green Data Centre Technology Roadmap (the 'Roadmap') with input from the research community, industry, and government agencies. The Roadmap provides insights on methods and systems to reduce energy consumption in data centres without compromising system performance and security requirements. It also identifies key areas of research where academic institutions can take a lead role in developing emerging technologies.

1.1 PAST AND CURRENT INITIATIVES

In 2009, the Infocomm Development Authority of Singapore (IDA) initiated the Green ICT programme with the aim of improving the energy efficiency and competitiveness of the Singapore data centre industry.

Investment Allowance Scheme: The Investment Allowance was introduced to help data centre operators improve the energy efficiency of their existing data centres. Successful applicants are granted an investment allowance of 30% to 50% off the fixed capital expenditure incurred on retrofitting their data centres.

Green Data Centre Standard: Published in 2011, the Singapore Standard 564 (SS564) was developed by the Green DC Standards Working Group under the industry-led Information Technology Standards Committee. The SS564 aims to benchmark the energy efficiency of data centres, while providing a set of best practices for the industry.

BCA-IDA Green Mark for Data Centres: Jointly developed by the IDA and the Building and Construction Authority (BCA), the Green Mark is a rating system that encourages the adoption of energy efficient design, operation and management of data centres.

There is a clear necessity for data centre owners and users to account for their energy usage and adopt methods to reduce energy consumption. Although some efforts have already been made through the IDA-BCA Green Mark Standard and SS564, these have primarily focused on the facility (i.e. non-IT) systems of the data centre.

New technologies applicable to facility systems and IT systems will enable further improvements in energy performance beyond the best-in-class today. It is proposed that Singapore becomes an early adopter of beneficial new energy efficient technologies and apply medium to long term targets for energy efficiency.

1.2 OBJECTIVES AND APPROACH

The Green Data Centre Technology Roadmap considers how the carbon footprint of data centres in Singapore can be reduced. This is done in the context of ensuring the continued growth of the industry and enhancing Singapore's position as a data centre hub for the Asia Pacific region.

The desired outcomes of the Roadmap are:

- Development of a framework for the research community, the industry and the Singapore Government to assess their technology options in reducing the carbon footprint of the data centre industry.
- Identification of directions in Research, Development & Demonstrations (RD&D) that will enable Singapore to achieve a leading position in Green Information and Communications Technology (ICT) technology development.

The assessment of the suitability of individual technologies for Singapore data centres is primarily based upon the extent of current product development and the relative impact each technology has on energy consumption and sustainability. It should be noted that some technologies complement each other to different extents, while others are mutually exclusive. Additionally, some technologies are best applied to new data centres, retrofitted to existing data centres or applied equally to both new and existing data centres.

The Roadmap presents recommendations on technological directions to improve data centre energy performance. They are quantified with indicative incremental adoption costs, deployment curves, abatement potential and marginal abatement cost.

Information presented has largely been gathered from the annotated reference sources. The unique nature of this report requires considerable input based on expert opinion. Where expert opinion is presented due to the lack of external references, explanations of the reasoning and assumptions are included.

2.ENERGY USE IN DATA CENTRES

Data centres comprise three main energy groups as shown in Figure 2.1. Each group can contribute to improving the sustainability and energy efficiency of the data centre.

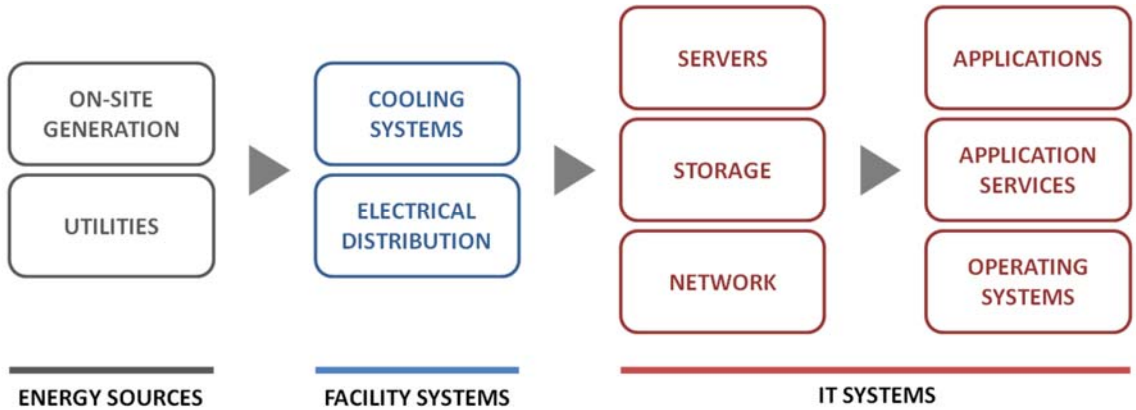

Figure 2.1: Data Centre Energy Groups.

2.1 ENERGY SOURCES

While a comprehensive discussion of energy sources is beyond the scope of this report, a short treatment is nonetheless included here for completeness.

Energy sources may reside outside the data centre (in the case of power from a utility company) or within the data centre (in the case of on-site generation options). In either case, data centres select their energy sources to achieve the best balance of cost-effectiveness and reliability.

Internationally, many large internet companies have made using renewable or green energy sources for their data centres a priority. However, the scale, climate and tidal range of Singapore restricts the use of geothermal, tidal, hydro and wind, either as viable sources of energy or for use as economisers for data centres. Approximately 96% of energy supplied in Singapore is fossil fuel based, derived from natural gas and other petroleum products [1].

2.2 FACILITY SYSTEMS

Facility systems refer to the infrastructure and processes needed to support the operation of IT equipment in a data centre. They comprise two principle components – cooling and electrical distribution. To date, most industry initiatives have been focused on improving the efficiency of data centre facility systems. This has been driven by the high proportion of facility system energy costs relative to the entire data centre. This is illustrated in Figure 2.2.

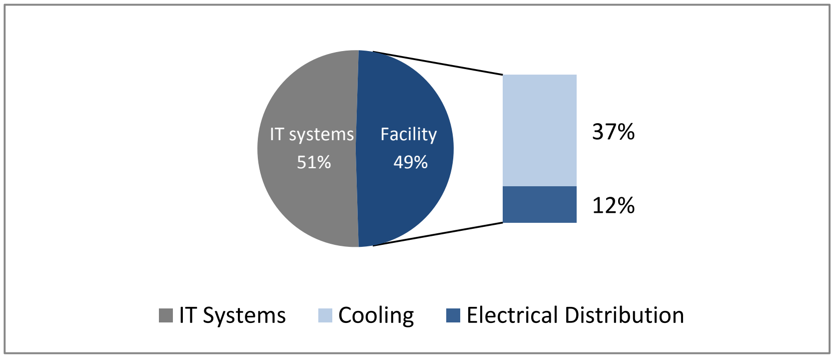

Figure 2.2: Energy profile of a typical data centre's Facility Systems in Singapore.

Power Usage Effectiveness (PUE) [2] is the primary metric used to measure the energy efficiency of a data centre's facility systems. PUE is calculated as follows:

$$\text{PUE} = \frac{\text{TOTAL DATA CENTRE POWER}}{\text{IT EQUIPMENT POWER}}$$

The ideal PUE is 1.0; in this case, all power supplied to the data centre reaches the IT equipment. In practice, power is required to provide cooling to the data centre. Power losses also occur in electrical conversion and distribution systems, and ancillary systems within the data centre. Numerous claims of low PUE have been made by companies around the world, with some stating PUE ratings as low as 1.05.

A 2012 site measurement study surveyed 23 data centres in Singapore and recorded an average PUE of 2.07 [3]. These findings are comparable to the average PUE ratings of 2.2 in the United States and 2.02 in Europe. However, a separate industry survey of 100 respondents in 2013 reported an average PUE of 2.61 [4]. A possible explanation for this discrepancy is self-selection by larger data centre operators in the 2012 Singapore study. On the other hand, the 2013 survey was able to capture a larger cross section of the data centre landscape.

Apart from PUE, other efficiency metrics include Water[1] and Carbon[2] Usage Effectiveness [5, 6]. Currently, only 24% of organisations globally collect data on carbon emissions and 34%

---

[1] WUE™ is the annual site water usage divided by IT equipment energy. This includes water used for humidification and water evaporated on-site for energy production or cooling of the data centre.

collect data on water usage [7].The Green Data Centre Primer further elaborates data centre performance metrics [8].

2.2.1 Cooling

On average, 37% of the total energy consumed by data centres in Singapore is used to cool IT equipment (Figure 2.2). There are many ways to reduce the energy involved in cooling. However, Singapore's climate poses a major disadvantage. The high temperatures and humidity of Singapore's tropical climate makes cooling data centres significantly more energy-intensive compared to other locations in the world.

The majority of data centres today operate the cooling systems of their data centres within the guidelines set by ASHRAE (formerly the American Society of Heating, Refrigerating and Air Conditioning Engineers). The guidelines set recommended operating envelopes with regard to temperature and humidity for various equipment classes (Figure 2.3) [9].

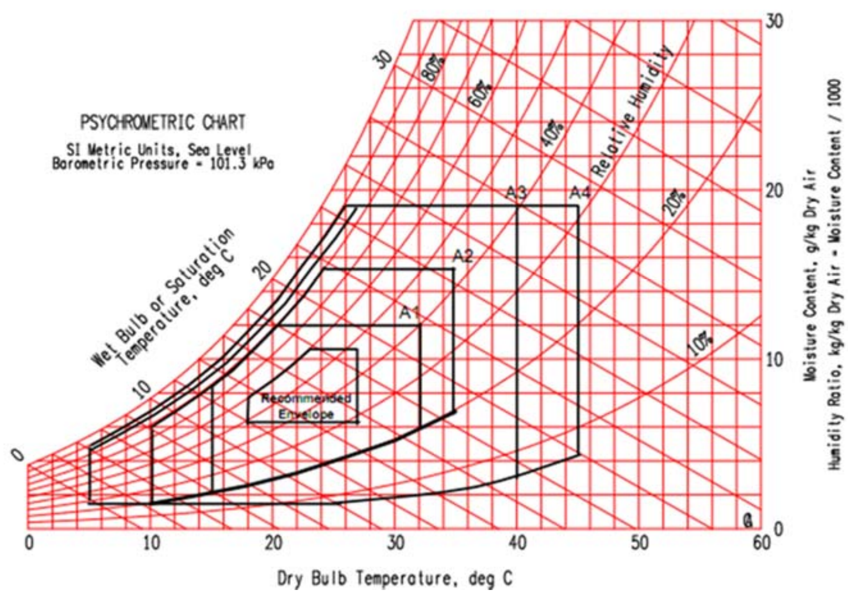

Figure 2.3: ASHRAE-recommended operating envelopes for data processing environments (Source: ASHRAE).

Current guidelines from ASHRAE for the upper limits measured at the air intake of IT equipment are detailed in Table 2.1. In practice, most data centres operate at temperatures well below the recommended upper limit set in the guidelines, resulting in high cooling costs.

Table 2.1: ASHRAE operating ranges for Class A1 equipment.

[2] CUE™ is the total CO_2 emissions caused by the total data centre energy divided by IT equipment energy. CUE is a source based metric and includes carbon generated in the energy distribution path.

	Recommended	Allowable
Dry Bulb Temperature	27°C	32°C
Dew Point	15°C	Not specified
Relative Humidity	60%	80%

The guidelines also provide an assessment of the increase in IT equipment failure rates associated with various operating temperatures. This is detailed in Table 2.2. The most common failure modes in IT equipment are power supply failure and storage device failure.

Table 1.2: ASHRAE IT equipment failure rates at various operating temperatures.

Inlet Dry Bulb (°C)	Increase in Failure Rate
20	0%
22	13%
25	24%
27.5	34%
30	42%
32.5	48%

The increase in failure rates at temperatures even well below the upper temperature limit is significant to mission critical data centres, where higher failure rates are undesirable. Failure rates published by ASHRAE are likely to deter organisations from adopting higher operating temperatures. Further RD&D is recommended to test the failure rates stated by ASHRAE. Vigorous quantification of failure rates will encourage data centres to operate at higher temperatures, thus reducing cooling needs.

2.2.2 Electrical Distribution

The most significant component affecting the energy efficiency of the electrical system is the UPS (uninterruptible power supply). Various UPS types are used in Singapore; the most common type being the on-line double conversion UPS. Assuming a typical 2N UPS configuration and an average UPS utilisation of 30%, typical of a mature data centre, the electrical system efficiency is likely to be approximately 89%; excluding transformer and distribution losses.

An alternative to the on-line mode of operation is the off-line mode where the UPS' normal condition is off and power flows through the static bypass. Operating off-line can improve efficiency up to 98%. Distribution losses, which account for around 37% of the average losses in the electrical system, can be reduced, though not to the same extent as the UPS system.

Another method to reduce losses in the electrical distribution system is the use of DC power over the current AC power distribution systems. The inherent benefit that DC data centres have over AC data centres is a reduced component count. This inevitably results in a higher reliability configuration since there are fewer conversion steps compared to an AC system. Typical AC data centres in Europe and Asia (excluding Japan) have five conversion steps associated with normal power delivery to the server while DC systems require only three steps. As a result DC systems have less equipment usage, less heat generation and increased reliability because there are fewer complex components and potential points of failure.

2.3 INFORMATION TECHNOLOGY (IT) SYSTEMS

2.3.1 Hardware

Information technology (IT) infrastructure plays a major role in the energy efficiency of a data centre. Of the various physical IT components, servers account for the majority of physical IT energy use (Figure 2.4). This is followed by storage devices and network devices.

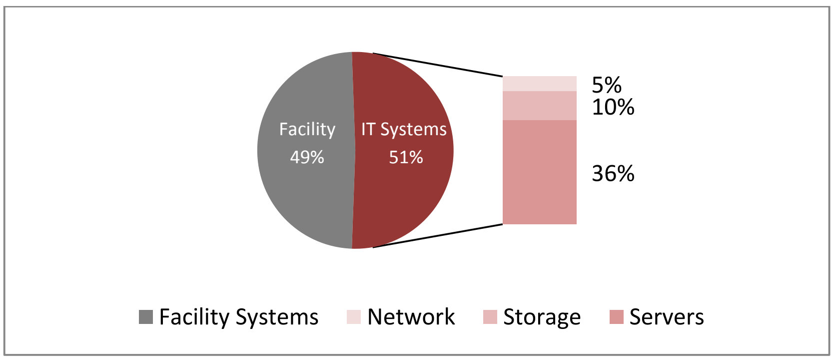

Figure 2.4: Energy Profile of a Typical Data Centre's IT Systems in Singapore.

Efforts have been made to reduce the energy consumption of servers and increase their energy efficiency. Techniques such as energy-aware workload management tools, voltage and frequency scaling and power capping have been introduced by hardware manufacturers. Currently the most significant issue affecting server energy performance is the power used at zero processor utilisation. Intuitively one would expect this to be zero or close to zero; instead the average is 50% of the maximum power consumption [10]. Techniques such as device power-down and invoking processor sleep states can address this problem. The drawback of this is a latency penalty as additional time is required to wake from a sleep state in response to demand spikes.

Storage capacity is the fastest growing segment of physical IT equipment. The emergence of cloud computing, in particular the public cloud is driving an extraordinarily high growth rate in storage. Globally storage capacity is expected to grow to approximately 40YB[3] by 2020 [11]. Some measures have been taken by manufacturers to address energy usage, such as

[3] 1YB= 10^{24} bytes

using different drive technologies and new techniques for deploying data storage. There are, however, many aspects of storage research that can be addressed in the context of sustainability and energy efficiency.

Contemporary data centres use both fibre-optic and copper cables for networking. As network speeds increased in conjunction with the commoditisation of fibre products, a larger proportion of fibre is used in the data centre. Similarly, advances in fibre and copper technology have enabled higher throughputs across the data centre network. In parallel, network switches have become more energy efficient through a combination of improved power supply efficiency and component design. Energy saving techniques involving the use of intelligent routing methods, node power-down and sleep states could potentially realise significant network energy savings.

2.3.2 Software

Software that is optimally configured for energy performance directly reduces the energy consumed by IT equipment and has a cascading effect on the energy consumption of downstream devices. The Software-to-Facility Energy Cascade Effect makes software one of the most important areas to address in terms of reducing energy consumption. This is illustrated in Figure 2.5.

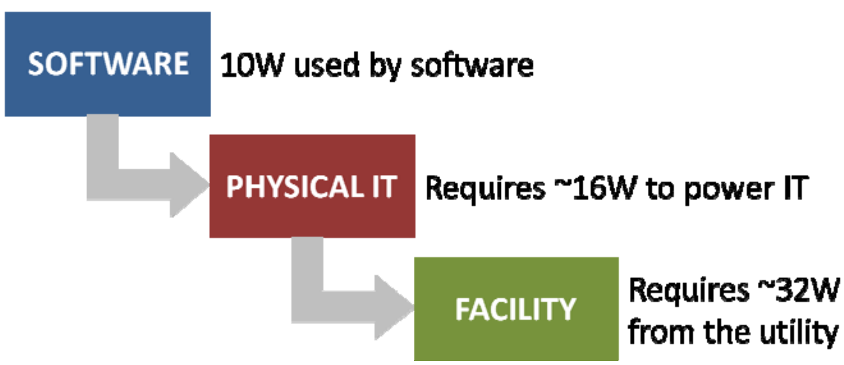

Figure 2.5: The Software-to-Facility Energy Cascade Effect.

With the exception of mobile platform designers, for whom extending the battery life is a major factor in platform software and application design, most applications and operating system developers do not consider how energy efficiency and usage is impacted by software. In the case of fixed platforms, the impact of energy-aware language structures, coding and compilers have largely been overlooked by software developers.

The industry trend towards cloud based services (IaaS, PaaS and SaaS) particularly in public and hybrid cloud environments will inevitably become more energy-focused. Cloud service providers operate a usage-based charging model for software and infrastructure. Since energy costs usually account for a large proportion of a data centre's operating expenditure, cloud providers will have to prioritise energy efficiency to maximise their return on investment.

2.4 DESIGN AND DEPLOYMENT

Technologies should be complemented with design and deployment best practices to ensure that they are deployed to their fullest effect. In most cases, the focus has been on "right-sizing" data centre deployments – meeting IT demands without over-provisioning resources.

An essential and often overlooked factor in the design of data centres is part-load performance. Usually, energy efficiency is considered at full design load, but this is rarely relevant since it does not convey the progressive increase in data centre occupancy and hence energy use of the data centre.

Four factors that limit data centre capacity are physical space, cooling, power and network connectivity. Poor design and operational practices often fail to consider that these factors can grow at significantly different rates. This results in one of these design factors creating a bottleneck that prevents the data centre from realising its full design capacity and return on investment. The problem is further compounded by the use of fault tolerant UPS systems. At 2N topology[4], the UPS rarely operates above 50% utilisation. In reality, it would spend most of its tenure operating at utilisation levels between 10% - 30%. It is widely acknowledged by industry experts that IT and facility departments are not aligned to common objectives when it comes to the data centre. IT departments cannot predict their infrastructure requirements beyond 3-5 years, yet data centre construction capital expenditure is usually amortised over 15 to 20 years. This problem is referred to as the *Technology-Facility Paradox*.

In a typical scenario, during the initial stage of data centre planning, IT departments approximating future requirements tend to over-provision. The problem is exacerbated by the MEP (mechanical, electrical and plumbing) engineering teams who, due to their lack of understanding of IT infrastructure, further over-provision with the use of conservative designs. The outcome is unnecessarily high capital and operating expenditures for data centre deployments.

[4] 2N topology is a UPS architecture providing multiple-redundancy. In this case, each power supply is connected to its own UPS, providing protection from both power supply failure and failure of the UPS.

2.4.1 Multi-Tiering

Current data centre designs are generally wasteful in terms of both embedded energy and operational energy use. A typical data centre comprises a single tier of MEP design that is applied uniformly across the computer floor. Nonetheless, from the end-user's perspective, not all service-lines or applications are of the same criticality. Single-tier data centres offer unnecessarily high redundancy for applications that do not require it.

Multi-tiering is the design of data centres to provide multiple levels of reliability that commensurate with service-line or application criticality, the underlying application and technology infrastructure. A simple two-tier arrangement is shown in Figure 2.6. Multi-tiered data centres avoid redundancy, which reduces capital costs and fixed losses, and when used in conjunction with a progressive modular design, avoids the over-provisioning of MEP systems. The resulting data centre is also more adaptive to IT changes, and is able to provide variable and scalable availability.

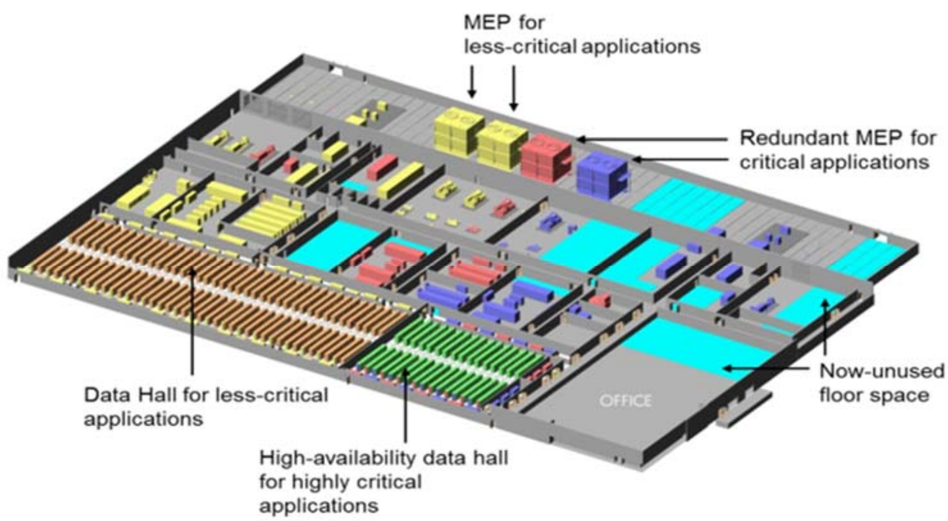

Figure 2.2: A Simple Two-Tier Data Centre Deployment (Source: HP)

2.4.2 Modular Provisioning

Over-provisioning occurs when the capacity of MEP systems deployed far exceed the needs of the data centre's IT load. This results in energy wastage due to fixed losses. Fixed losses are inherent to all electrically powered equipment. They occur when equipment is powered on and remain constant irrespective of whether the equipment is utilised.

Modular provisioning is the design of data centres to allow for IT and MEP resources to be deployed in progressive stages. In turn, MEP systems are able to better match the actual IT load. This approach to designing and deploying data centres has significant benefits in terms of deferred capital expenditure and operating expenditure. Energy costs can be reduced by the elimination of fixed losses. Figure 2.7 illustrates the improved efficiency profile of modular deployments compared to conventional monolithic deployments.

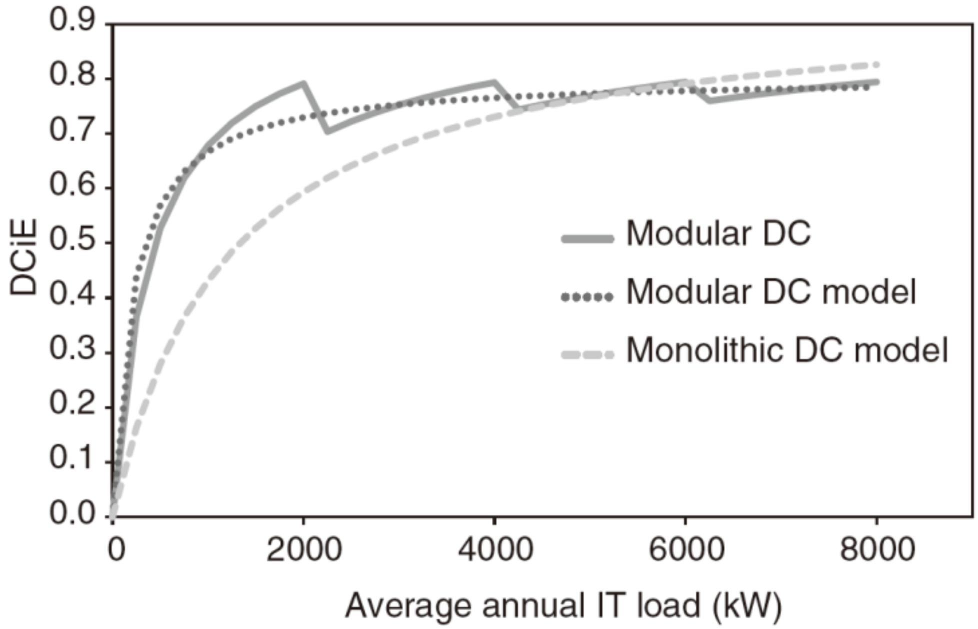

Figure 2.3: Efficiency of Modular compared to Non-Modular Deployments [12].

3.TECHNOLOGY DEVELOPMENT

3.1 DIRECTIONS IN RESEARCH AND INNOVATION

In this section, we review potential research directions in **facility, information technology (IT) systems, and integrated approach to design and deployment of data centres**. A total of **13 technology areas** are presented and analysed. The areas were selected in consultation with the research community, industry, as well as with input from government agencies. Together, they represent the major technological developments in data centre energy efficiency pursued today.

3.1.1 Facility Systems

The high cost of cooling is Singapore's main competitive disadvantage when it comes to hosting a data centre. Even in locations with favourable climates, cooling is relatively expensive. Apart from operating temperatures closer to the upper limit of the ASHRAE recommendations, energy efficiencies of data centres can be improved by reducing cooling needs. Cost of cooling can be reduced by improved cooling technologies and reducing the need for active cooling. Technologies to enable these two approaches are elaborated below.

Direct Liquid Cooling or immersion cooling is the use of dielectric liquids as alternatives to air as cooling mediums. Some electrical transformers and supercomputers have used dielectric fluid submersion for many years. Until recently, this method involved higher costs compared to air cooling. Direct liquid cooling is similar to direct-to-chip water cooling, except that unlike direct-to-chip water cooling, the dielectric fluid passes over all surfaces inside the IT equipment. The fluid is then transferred to an external heat exchanger where it is cooled for recirculation back to the IT equipment. Trials involving immersion cooling have demonstrated significant energy savings. However, perceptions of increased complexity and restrictions on the type of hardware that may be used have hindered widespread deployment.

Close-Coupled Refrigerant Cooling is the use of liquid refrigerants with micro-channel heat exchangers to bring cooling as close as possible to the actual heat sources. Beyond being more efficient than air as a cooling medium, the close proximity of the cooling and heat source allows for greater precision in the application of cooling. Thus data centres do not need to be overcooled to address hotspots. Trials have reported energy efficiency improvements of 80% over traditional water cooled chillers.

The following are measures to reduce the need for energy-intensive cooling.

Air and Cooling Management refers to the broad suite of technologies and accompanying practices that allow for the optimisation of data centre cooling resources. It remains the single largest near and medium term opportunity to improve data centre energy efficiency due to the continued popularity of air as a cooling medium. Today, best practices with regard to air and cooling management are well-known, if inconsistently followed. Examples of recent research efforts include the use of computational fluid dynamics (CFD) simulations to aid in the design of racks and servers, to improve airflow and heat management.

Passive Cooling aims to eliminate or reduce the need for energy-intensive active cooling systems by improving passive heat transfer from chips, servers or racks. An example of this is the use of heat pipes to reduce hotspots on chips or servers. Passive cooling can reduce data centre cooling demand by up to 40%, significantly reducing data centre operating costs. This potentially allows higher computational densities to be achieved without the complexity and cost of liquid cooling architectures or immersion tanks.

Free Cooling (Hardening) refers to improving the protection of IT equipment against adverse environmental conditions, namely particulate matter, temperature and humidity. The need to cool the data centre has traditionally been a large energy overhead. Globally, the data centre industry has responded by exploiting free-cooling techniques in locations where practicable and/or moving operations to those locations. Singapore's hot and humid tropical climate restricts the use of free-cooling. Improving the environmental operating envelop of IT equipment would allow data centres to be operated with lower levels of cooling than is possible today.

With regard to power distribution, efforts are focused on reducing conversion losses or avoiding conversions altogether.

Power Supply Efficiency refers to improving the efficiency of power transfer between the energy source and IT equipment. This is largely a function of losses in electrical conversion. Incentivised by industry programmes like 80 Plus, power supply efficiencies have improved significantly since 2006, improving from 67% to over 90% at ideal loads.

3.1.2 IT Systems

Due to the Software-to-Facility Energy Cascade Effect (Figure 2.5), developments in software energy efficiency can effect a magnified improvement throughout the data centre technology stack. However, software has until recently, received little to no attention from the data centre industry. Conversely, software energy efficiency has been a priority for the mobile platform industry for many years.

To some extent, new research can take advantage of the gains made by mobile platforms. However, there are profound differences in data types used by mobile and fixed platforms;

more work is needed to understand the techniques that can be translated and where new techniques need to be developed.

Significant progress has been made in improving energy efficiency using dynamic voltage and frequency scaling. However, there are numerous areas that require further research and development to progress IT energy performance.

Suggested areas for research include:

Software Power Management, which refers to development and translation of techniques to reduce computational demand by software when not performing useful work. Device idle power consumption has improved considerably. However, it remains a major barrier to energy proportional computing. The best levels of idle energy are approximately 25% of the maximum power draw for the most recent generation of servers; the average is approximately 50%. The use of sleep states during CPU idle-time is fundamental to increasing battery life in mobile platforms. Although the use of sleep states is widespread in mobile devices these techniques are rarely used in fixed platform IT. This omission presents a significant opportunity to reduce data centre energy consumption.

Energy-Aware Workload Allocation is the automated migration of workloads in virtualised environments to make better use of the most energy efficient servers or data centres across a network. Where and when possible, workloads could potentially be consolidated onto the most energy efficient servers or distributed in an energy efficient manner. Unused servers can then either be put into sleep mode or (time permitting) be powered off. All this needs to be achieved while maintaining quality-of-service and security standards.

Dynamic Provisioning is the automated scaling of a data centre's compute, storage and network resources, and by extension its energy expenditure, to match demand. This is a key feature of energy proportionality in the data centre. The main challenge is dealing with the inherent trade-off between energy proportionality and latency in ramping up resource availability in response to demand spikes.

Energy-Aware Networking is the design and use of the network in a manner that minimises its energy footprint. The need to ensure network resilience and meet demand spikes has led to the practice of deploying a large number of devices to achieve this. The aim of energy-aware networking is the dynamic minimisation of network energy footprint while maintaining performance and resilience.

Wireless Data Centres is the use of wireless networking in data centres to augment or replace cabled communications. While interest in wireless data centres is mainly in cable management (or the lack thereof) and the associated physical flexibility afforded, some energy efficiency potential has been noted.

Memory Type Optimisation in an energy efficiency context determines how data is stored and accessed using energy efficient technology whilst meeting execution quality-of-service standards. Among the most significant developments are those in non-volatile memory devices, offering a significant impact on the energy efficiency of IT at silicon level due to its inherent ability to retain its state under zero power. Depending on the workload, this could significantly reduce energy consumption.

3.1.3 Integration

Sustainability in data centres must be approached in an integrated manner, with opportunities pursued, at each level of the data centre technology stack. So far optimisation efforts have restricted themselves to either the facility systems or IT systems. Further opportunities are afforded by integrating facility systems, IT systems and overall alignment with actual service-line requirements.

Data Centre Infrastructure management (DCIM) market offerings are currently disparate and not always comparable with each other. This is because the scope of DCIM is not subject to any widely accepted definition, with many domains sitting under the broad DCIM heading. The majority of DCIM toolsets do not offer a solution within every domain. Domains typically include, but are not limited to:

- Utility power
- Mechanical and electrical plant
- Asset management
- Capacity management
- Space visualisation
- Simulation

Advanced DCIM should enable centralised monitoring, management, and intelligent capacity planning of critical systems. Subsequent generations of data centres will increasingly benefit from operational advantages that DCIM tools offer provided implementations include the integration of IT, Facilities and Application Management

3.2 COMPARATIVE IMPACT ASSESSMENT

In this section, the environmental and economic potential of each technology described above are assessed. The methodology used is described in the Appendix.

We consider the impact of each technology applied in isolation until 2030. This has the following important implications for the calculations that follow: i) they are useful for quantitatively assessing the relative impact of each technology, but ii) they are not useful for determining the absolute quantitative impact of each technology.

The technologies are presented here solely in terms of their environmental and economic potential. However, we must remain cognisant of the fact that technology development and deployment is determined not just by technical potential, but also by business and market considerations. However, the analysis presented here forms a useful starting point for the research community and business leaders to assess their technology options for increased sustainability of data centres in Singapore.

3.3 RESULTS AND ANALYSIS

The results of the calculations are illustrated in Figure 3.1. Most of the technologies are NPV-positive (i.e. they appear in the upper positive area of the plot).

The largest opportunities for energy efficiency in Singapore data centres are split evenly between technologies for the facility (in blue) and for IT systems (in red). Integration (in green) also presents a significant opportunity.

In reality, improvements to the facility are likely to be the first to yield results in terms of actual energy savings. This is because facility improvements are often less disruptive to data centre operations compared to similarly ambitious technologies for IT systems. This is reflected by the faster deployment curves for facility technologies.

Nonetheless, the even spread of opportunities across the data centre technology stack emphasises the need to pursue energy efficiency in a concerted manner at all layers of the data centre technology stack.

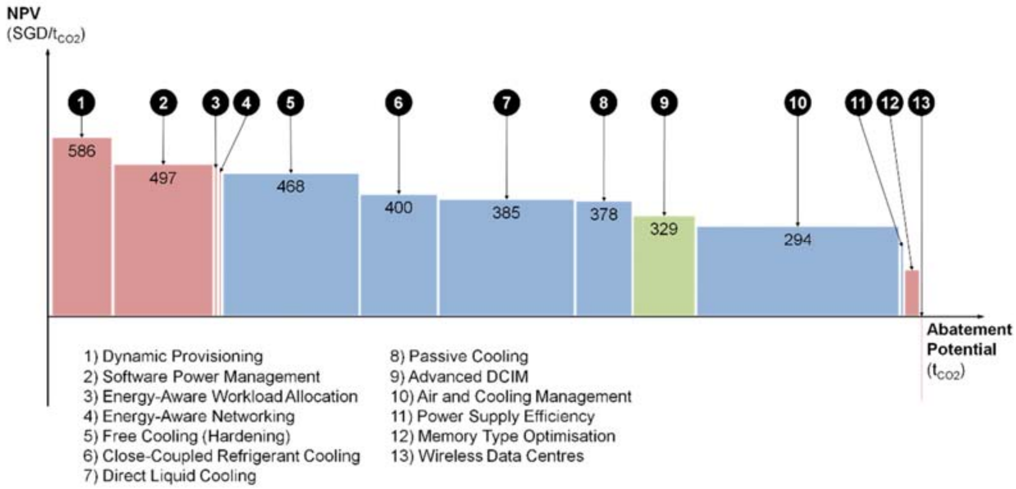

Figure 3.1: Abatement Potential vs Net Present Value.

4. CONCLUSION

There is a necessity for data centre owners and users to account for their energy usage and to introduce methods to reduce energy consumption. New technologies applicable to facilities, IT infrastructure and software will enable further improvements in energy performance beyond what is available from the best-in-class today.

It is recommended that Singapore focus on the following 3 areas to improve the long-term energy sustainability of data centre industry:

- Lowering the high cost of cooling in Singapore's tropical climate via i) energy efficient cooling methods, and ii) the hardening of IT equipment to withstand higher temperature and humidity levels, thereby reducing the need for cooling.
- Improving the energy proportionality of data centre IT systems via i) software power management, ii) the development of energy efficient hardware and architectures, and iii) the automated optimisation of resource provisioning and workload allocation in software-defined data centre environments.
- Integration and optimisation across the traditionally silo-ed IT systems and facility systems of the data centre.

5. REFERENCES

[1] EMA, Singapore Energy Statistics 2013, 2013.

[2] The Green Grid, PUE: A Comprehensive Examination of the Metric, 2012.

[3] NEA-HP, Data Centre Energy Efficiency Benchmarking, 2012.

[4] Digital Realty, Singapore Survey Results, 2013.

[5] Green Grid, Water Usage Effectiveness: A Green Grid Data Centre Sustainability Metric, 2011.

[6] Carbon Usage Effectiveness: A Green Grid Data Centre Sustainability Metric, 2010.

[7] Uptime Institute, 2012 Data Centre Industry Survey, 2012.

[8] NCCS-NRF, Green Data Centre Primer, 2014.

[9] ASHRAE, Thermal Guidelines for Data Processing Environments, 2011.

[10] Google, The Case for Energy-Proportional Computing, 2007.

[11] IDC, The Digital Universe in 2020, 2012.

[12] Fujitsu, Characterisation of Data Centre Energy Performance, 2012.

6. APPENDIX

APPENDIX A

Methodology

The merit of each technology is assessed using two metrics:

1. **Abatement Potential, α_{2030}**, the cumulative mass of CO_2 emissions avoided from the present to 2030; and
2. **Net Present Value, β_{2030}**, the economic benefit per unit mass of CO_2 emissions avoided from present to 2030.

For each of the identified technologies, assumptions are made about the:

- **Efficiency Potential, μ** , the efficiency gain as a ratio of the total energy consumption of a data centre;
- **Cost, C** , the incremental cost as a percentage of the capital outlay of a data centre;
- **Market Penetration, P_{20XX}**, measured as a percentage of the Singapore data centre market by capital outlay in 20XX.

For each technology,

$$\alpha_{2030} = E_{2030} \times P_{2030} \times \varepsilon_{2030} \times \mu,$$

$$\beta_{2030} = \frac{(E_{2030} \times P_{2030} \times t_{2030} \times \mu) - (N_{2030} \times P_{2030} \times C)}{\alpha_{2030}}$$

Where,

E_{2030}, is the projected electricity consumption by data centres in Singapore in 2030,

t_{2030}, is the projected price per unit energy in Singapore in 2030,

N_{2030}, is the projected capital outlay of data centres in Singapore in 2030, and

ε_{2030}, is the projected Grid Emission Factor of Singapore in 2030.

Table A.1 details the assumptions for the assessed technologies.

Technologies are grouped according to the Energy Group on which they bring about energy efficiencies, regardless of where the technology is actually deployed. The Facility Group of technologies reduces the energy consumption of facility systems; the IT Group (encompassing both physical IT and software) reduces the energy consumption of IT systems; and the Integration Group reduces energy consumption across the entire data centre.

Table A.1: Model input (in orange) of the assessed technologies.

		μ	C	P_{2020}	P_{2025}	P_{2030}
Facility	Air and Cooling Management	0.1	0.50	80	100	100
	Passive Cooling	0.1	1.00	0	30	50
	Free Cooling (Hardening)	0.2	0.01	0	40	80
	Close-coupled Refrigerant Cooling	0.2	1.00	5	25	50
	Direct Liquid Cooling	0.4	2.00	0	25	50
	SMPS Efficiency	0.0	0.10	0	25	50
IT	Software Power Management	0.4	1.00	0	10	50
	Dynamic Workload Allocation	0.0	0.01	0	10	30
	Dynamic Provisioning	0.4	0.50	0	0	40
	Energy-Aware Networking	0.0	0.01	0	10	30
	All Wireless Data Centres	0.0	0.30	0	10	30
	Memory Type Optimisation	0.0	0.50	10	50	75
Integration	Advanced DCIM	0.0	0.50	20	50	75

Disclaimer, Limitation of Liability

This report is the result of a collaborative effort between the i_3 Solutions Group Pte Ltd, the Infocomm Development Authority and other Singapore Government Agencies. Users of this report shall make their own independent business decisions at their own risk and, in particular, without undue reliance on this report. Nothing in this report shall constitute professional advice, and no representation or warranty, expressed or implied, is made in respect of the completeness or accuracy of the contents of this report. i_3 Solutions Group accepts no liability whatsoever for any direct or indirect damages resulting from any use of this report or its contents. A wide range of experts reviewed the draft reports; however, the views expressed do not necessarily represent the views or policy of i_3 Solutions Group or other contributors to this report.

Lead Agencies

iDA SINGAPORE

Commissioning Agencies

NCCS
NATIONAL CLIMATE CHANGE SECRETARIAT
PRIME MINISTER'S OFFICE SINGAPORE

NATIONAL
RESEARCH
FOUNDATION
PRIME MINISTER'S OFFICE
SINGAPORE