

Iceotope Engineering Review

Desktop Study Report- Liquid and Air-Cooling
Compared

Iceotope Technologies Limited

Job No: 1027146
Doc Ref: ICTP-CDL-XX-XX-RP-Z-90200
Revision: —
Revision Date: 08 March 2021



Project title	Iceotope Engineering Review	Job Number
Report title	Desktop Study Report- Liquid and Air-Cooling Compared	1027146

Document Revision History

Revision Ref	Issue Date	Purpose of issue / description of revision
-	08 March 2021	First Issue

Document Validation (latest issue)

X _____ X _____ X _____
Principal author Checked by Verified by

Executive Summary

SECTION 1: Introduction

The purpose of this report is to introduce liquid cooling for data-processing equipment, with a focus on the chassis-level precision immersion and Direct-to-Chip solutions produced by Iceotope Ltd. Different forms of liquid cooling are described and compared, as are other hybrid forms of liquid to air cooling.

The performance and energy efficiency advantages that liquid cooling can offer over conventional air-cooling are assessed, based around four imaginary Data Centres, all of which are based upon real world scenarios. Calculated PUE and WUE performance figures and predicted OPEX savings are presented, along with an indication of the environmental benefits, to illustrate the advantages that an operator would achieve through the adoption of liquid cooling.

The report predicts a strong, growing need for data centre capacity and highlights the challenge of providing cooling to these facilities in ways that minimise their carbon footprint. Several factors driving the adoption of liquid cooling are highlighted, including:

- Increasing CPU power and associated increases in power densities, both at a rack level and within the overall white space.
- Emergence of Artificial Intelligence (AI) and other power-hungry applications.
- Legislation, at local, national, and international levels.
- Increased adoption of Waste Heat Recovery (WHR), which is greatly facilitated by the use of liquid cooling.
- The need to reduce carbon footprint, particularly through a reduction in power and water consumption.
- Limited availability of large plots of development land for low-density, single storey data centres – recognising that the higher power density achievable with liquid cooling can reduce the overall footprint of the facility, especially if constructed in multi-storey. Heat rejection may also simplify, by virtue of the operating temperatures of the FWS, which for Classes W4 and W5, would be higher than those of the air and chilled water system typically deployed within an air-cooled facility.

The report describes the 5 classes of liquid cooling defined by ASHRAE, ranging from W1, (with an FWS temperature range of 2°C - 17°C), to W5 (with temperatures >45°C).

Section 1.5.1 of the report introduces the following generic forms of liquid cooling:

- Direct-to-Chip (Single-phase and Two-phase)
- Full Immersion – otherwise known as “Tub” or “Open-Bath” (Single-phase and Two-phase)
- IT Chassis-Level Precision Immersion

Reference is also made to hybrid air-cooled solutions, such as In-Row and Rear-Door Units, noting that these can be useful to support clusters of high-density racks, but recognising that the cooling medium is still air, so the capacity of High Density (H-D) racks that can be supported is limited by the amount of airflow that can be delivered to them.

In section 1.6.1 the report assesses the diverse variety of mainstream applications into which liquid cooling is now being applied:

- Cloud & Hyperscale
- Enterprise & CoLocation
- Edge Computing & IoT

The report highlights the following generic advantages of liquid cooling compared to air-cooling:

- Ability to accommodate rising power densities – most air-cooled racks typically operate in the range of 7-10kW/rack. Higher power densities can be supported by adoption of supplementary cooling such as in-row and rear door coolers. However, Iceotope’s KU:L Box System can comfortably accommodate power densities of up to 60kW/rack.
- Significantly reduced building footprints – The report highlights that by increasing the Power Density from 8kW/rack to 60kW/Rack, it is possible to significantly reduce the number of racks required to support the load. This would

equate to a reduction in data hall footprint of 86%. Even for application where mixed high- and low-density loads are anticipated, where power is available, groups of high-density loads can be introduced into a low-density hall and greatly increase its capacity without the need for additional space.

- Simplified Heat Rejection - the higher FWS operating temperatures typically associated with liquid cooling provides for simplified and smaller heat rejection plant. The need for mechanical compressor-based cooling plant may be avoided and the need for evaporative cooling may only arise in specific localities, or under peak summertime conditions. The requirement to design the data hall to accommodate large air-handling cooling plant, or large numbers of CRAHs is avoided. This reduces the overall complexity of the building and would simplify multi-storey construction.
- Reduced CAPEX - due to the elimination of aisle containment systems, the reduction in cooling plant complexity and the potential to reduce the size and cost of both the building and the plot on which it sits.
- Energy Efficiency – water has a volumetric heat transfer capacity that is over 3,000 times greater than that of air, so cooling can be delivered far more efficiently through this medium. Furthermore, liquid cooling systems typically operate with a supply temperature in the range 40-48°C, compared to a typical range for air-cooling of 16-25°C, which will enable far greater use of free cooling.
- WHR is far more effective and easier to achieve, thanks to the higher operating temperatures cited above.
- Reduced water consumption - the higher operating temperatures noted above, means that the use of evaporative cooling is greatly reduced or even eliminated, resulting in greatly reduced water consumption.
- Reliability and life-expectancy of components is greatly enhanced with liquid cooling, due to the uniformity of cooling and the elimination of hot-spots.

Section 1.7.5 of the report provides definitions of ‘Power Utilisation Effectiveness’ (PUE) and ‘Water Utilization Effectiveness’ (WUE), before highlighting in the following section the shortcomings of these measures – most notably the failure of the PUE metric to recognise the proportion of server power that is consumed by parasitic loads, such as the dielectric fluid pumps within an liquid-cooled server, or the PSU cooling fans within an air-cooled server. Whilst the latter will typically consume around 10% of the total rack power, the former only consumes around 3%, thus providing a ‘hidden energy saving’ of around 7% of total rack ITE power.

To provide a more comprehensive, transparent, and accurate measure of efficiency, the report therefore recommends the use of ‘IT Power Usage Effectiveness’ (ITUE) as a useful measure of server specific efficiency. Furthermore, it suggests that ‘Total-Power Usage Effectiveness’ (TUE), calculated as $ITUE \times PUE$, provides a more precise indication of the data centre overall energy performance.

Section 1.7.9 considers local environmental considerations, such as the reduced noise propagation of liquid-cooled systems (both internal and external) and the ambivalence of liquid cooled systems to poor air quality and variations in humidity, that can both be so damaging for air-cooled systems.

The report then describes the 3 Classes of “Continuous Cooling” defined by the Uptime Institute (UTI), highlighting how the increased thermal storage in chassis and full immersion liquid-cooled systems compared to those using air, combined with the faster start-up capability of dry/hybrid coolers compared to chillers, makes liquid-cooled systems far less vulnerable to temporary disruptions in the power supply.

SECTION 2: Iceotope Liquid Cooling Technology

The report provides a summary of key facts regarding Iceotope, along with an overview of its various Liquid Enable Architecture and Technology Platforms, including:

- KU:L Heat Sink - part of the ‘Precision Delivery’ architecture, providing localized ‘bathtub’ containment of dielectric fluid over the centre of the chip.
- KU:L Cold Plate - maximum cooling for HPC applications.

- KU:L System - core reference platform, targeting users and applications where high volume reference form factors are used at scale.
- KU:L Hybrid System - using the core components of the KU:L system, but replacing the KU:L heat sink with a KU:L cold plate.
- Hyperscale Cloud KU:L Box - chassis-type server enclosures in an insulated, rigid plastic structure, partially filled with dielectric fluid, that is cycled through plate heat exchangers at the rear of the box.
- Rack - developed in collaboration with Schneider Electric, that houses up to eight KU:L Box 'Cubbies' plus an optional liquid-cooled power supply unit. Provides easy access, allowing an individual chassis to be partially withdrawn from the rack, whilst maintaining the power supply and without disrupting the flow of cooling water to the adjacent racks.
- KU:L 2 - a flexible and server adaptable solution, where current 1U and 2U air-cooled servers can be converted to liquid cooling by a simple swap of air cooling infrastructure with Iceotope's liquid cooling technology.

The report then considers the relative merits of Iceotope's chassis-level approach to liquid cooling compared to the more traditional 'Full Immersion' alternative, highlighting the following advantages:

- Far less dielectric fluid is required for a given cooling capacity. This has cost and environmental benefits.
- The reduction in fluid volume equates to a significant weight reduction, which for a large deployment would equate to an advantageous reduction in structural load.
- The KU:L Box cubbies each form a self-contained unit that is interchangeable with other cubbies within the rack, providing 'Plug and Play' functionality. In the event of a server failure, a replacement Cubby can be shipped to site, pre-built, tested and already filled with dielectric fluid, where it can be quickly exchanged for the failed unit. At all times, the replacement server would remain secure within its sealed Cubby, protected from any airborne contamination.
- Individual servers can be accessed and serviced without affecting adjacent equipment within the rack.
- With the depth of dielectric fluid being much shallower than with full immersion systems, it is far easier to access, inspect and if necessary replace individual components in a liver server tray, without creating the drips and spillages often associated with full immersion systems.
- The 'Plug and Play' functionality of KU:L Box is also advantageous for provisioning new installations. The racks can be delivered to site pre-populated with Cubbies, which would be tested at the factory and arrive fully charged with dielectric fluid.

Section 2.4 then describes the interface of these systems with the Facility Water System (FWS) and provides recommendations to optimise the FWS for efficient, reliable operation.

SECTION 3.0: The Liquid Cooled Data Centre

Here the report summarises the characteristics of the Liquid-Cooled Data Centre, including building form and layout, heat rejection plant and implications regarding footprint.

The ASHRAE Thermal Guidelines are explored in more detail, including a review of the heat rejection systems typically used and the potential utilisation of WHR. Diagrams are provided, showing the typical infrastructure arrangements that would be recommended in order to achieve different levels of supply temperature to the CDU.

Section 3.2 describes the issues to be considered when designing a Liquid-Cooled Data Centre, including the recommended temperatures and pressures in the Technology Cooling System (TCS) and Facility Cooling System (FCS). Consideration is then given to the relative merits of Low-Level versus High-Level Services Strategies.

The section concludes with a detailed review of the dielectric fluids typically used for ITE cooling applications, which fall into two broad categories:

Oils – both mineral and synthetic.

Fluorocarbons – for both single and two-phase applications.

When comparing the relative merits of the two basic forms of dielectric fluids, the report highlights the following key points for each generic type:

Oils:

- Lower CAPEX
- Lower OPEX
- Zero GWP & ODP (but are polluting in the environment if released)
- No evaporation (but can therefore be messy if spilt)
- Typically, 10-20 year service life
- Typically classified as 'Combustible Liquids', rather than 'Flammable'
- Flash point is well above normal system operating temperatures

Fluorocarbons:

- Higher cost
- Prone to evaporation if not fully enclosed, thus a higher GWP
- Zero ODP
- Non-polluting in the environment
- Typically, 30-year service life
- Non-flammable, non-combustible

SECTION 4.0: Performance Comparison

Here the report describes three scenarios that Data Centre Operators may face, each with differing objectives, priorities, and constraints. In each case, the scenarios reflect real-world situations, as far as possible, based upon the needs of three fictitious Clients.

For each scenario, a design that includes the complete or partial adoption of Liquid Cooling, is compared against a benchmark that uses the best available form of air-cooling technology, taking account of the location and specific application.

In each case, the operating envelope reflects those most often requested in real-world projects by the relevant type of data centre operator:

Hyperscale Operator (Air-Cooled) – Assumed supply air condition at the server inlet of 18°C to 29°C Dry Bulb and 25% to 80% relative humidity. Supply air volume generally equates to a 12K rise across the racks.

CoLo Operator (Air Cooled) – Assumed supply condition following TC9.9 Recommended Range. Supply air condition at the server inlet of 18°C to 27°C dry bulb and relative humidity range equating to -9°C dew point at the lower limit and 15°C dew point and 60% relative humidity at the upper limit.

Liquid Cooled or Mixed Air and Liquid Cooled Environments – Taken as being the same as those specified for corresponding air-cooled environments.

Detailed weather data for each scenario is taken from the 2017 ASHRAE bin file.

Wherever possible, annualised energy and water consumption figure for the main mechanical plant has been based upon actual plant selections obtained from vendors. Where this has not been possible, performance has been estimated based upon pro-rata the water and power consumption of different sized plant, performing under identical operating conditions to those described in each scenario.

In all cases, the same approach to calculating energy and water consumption has been adopted for both air and liquid-cooled options.

The report includes a detailed analysis of the results, which in terms of impact on the operating costs can be best summarised in the following tables:

Base Case	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Application	5MW Hyperscale		2.4MW Co-Lo	1MW Edge DC
Location	Richmond VA	San Jose CA	London UK	Singapore
Tech Comparison	Fan Wall with Hybrid Cooler		IEC	Chiller + CRAH
Cost of Power for Cooling	£1,106,742	£1,098,586	£550,785	£755,899
Cost of Water for Cooling	£23,049	£11,077	£ 3,119	£ 15,993
Total Cost of Cooling	£1,129,791	£1,109,663	£553,904	£771,893
Cost of Power for Computing	£4,194,288	£4,194,288	£1,342,172	£838,858
Total Cost of Power	£5,301,030	£5,292,874	£1,892,957	£1,594,757
Total OPEX Cost	£ 5,324,079	£ 5,303,951	£ 1,896,076	£ 1,610,750
Total ITE Rack Power [kW]	5,000	5,000	1,600	1,000
Total ITE Power [kW]	4,500	4,500	1,440	900
Total Cost per KW of ITE Rack Power	£ 1,065	£ 1,061	£ 1,185	£ 1,611
Total Cost per KW of ITE Power	£ 1,183	£ 1,179	£ 1,317	£ 1,790
Total Cost of Cooling per KW of ITE Power	£ 251	£ 247	£ 385	£ 858

Liquid Cooling	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Cost of Power for Cooling	£562,263	£562,263	£663,514	£136,368
Cost of Water for Cooling	£ 3,557	£ 1,256	£ 343	£ 676
Total Cost of Cooling	£565,820	£563,519	£663,858	£137,044
Cost of Power for Computing	£4,520,510	£4,520,510	£2,091,552	£904,102
Total Cost of Power	£ 5,082,773	£5,082,773	£ 2,755,066	£ 1,040,470
Total OPEX Cost	£ 5,086,330	£ 5,084,029	£ 2,755,409	£ 1,041,146
Total ITE Rack Power [kW]	5,000	5,000	2,400	1,000
Total ITE Power [kW]	4,850	4,850	2,244	970
Total Cost per KW of ITE Rack Power	£ 1,017	£ 1,017	£ 1,148	£ 1,041
Total Cost per KW of ITE Power	£ 1,049	£ 1,048	£ 1,228	£ 1,073
Total Cost of Cooling per KW of ITE Power	£ 117	£ 116	£ 296	£ 141

Impact	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Water usage per KW of ITE Power	-85.7%	-89.5%	-92.9%	-96.1%
Power usage per KW of ITE Power	-11.0%	-10.9%	-6.6%	-39.5%
Total Cost per KW of Rack Power	-4.5%	-4.1%	-3.1%	-35.4%
Total Cost per KW of ITE Power	-11.4%	-11.1%	-6.7%	-40.0%
Total Cost of Cooling per KW of ITE Power	-53.5%	-52.9%	-23.1%	-83.5%
Increased ITE Capacity per Hall	*7.8%	*7.8%	55.8%	*7.8%
Footprint Required for additional Hall	n/a	n/a	n/a	-74.1%

*Note: The 7.8% increase in ITE Capacity per Hall in Scenarios 1a, 1b & 3 derives entirely from the reduced power draw of the on-board pumps in the liquid cooled design, compared to that of the server fans in the air-cooled base case. Substantial further increases could have been achieved by increasing the rack density and number of racks in the liquid-cooled design.

The report finally assesses the environmental impact of the alternative designs, by first estimating the carbon-footprint, then by assessing the amount of water consumed. In all cases, the report shows significant improvements:

Base Case	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Application	5MW Hyperscale		2.4MW Co-Lo	1MW Edge DC
Location	Richmond VA	San Jose CA	London UK	Singapore
Tech Comparison	Fan Wall with Hybrid Cooler		IEC	Chiller + CRAH
Total Power Demand [kWh/year]	49,821,715	49,745,054	17,790,948	14,988,326
CO2 Emissions [kg/year]	11,615,435	11,597,562	4,147,782	3,494,378
Total ITE Power [kW]	4,500	4,500	1,440	900
kg of CO2 per kW of ITE Power [kg/kW per year]	2,581	2,577	2,880	3,883

Liquid Cooling	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Total Power Demand [kWh/year]	47,770,423	47,656,116	25,893,476	9,778,856
CO2 Emissions [kg/year]	11,137,196	11,110,547	6,036,805	2,279,842
Total ITE Power [kW]	4,850	4,850	2,244	970
kg of CO2 per kW of ITE Power [kg/kW per year]	2,296	2,291	2,690	2,350

Impact	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Percentage Reduction in CO2 Emissions per kW of ITE Power [kg of CO2/kW per year]	11.0%	11.1%	6.6%	39.5%

Note: Emissions estimate based upon data published by BEIS (Dept. for Business, Energy & Industrial Strategy). For 2020, UK Grid generation represents 0.23314 kg of CO2/kWh.

Litres of water consumed per year, per kW of ITE Power	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Air-Cooled	4,687	2,252	1,982	16,261
Liquid-Cooled	671	237	140	638
Reduction	86%	89%	93%	96%

Contents

1.0	Introduction	8
1.1	Objectives of the Report	8
1.2	Authority	8
1.3	Abbreviations and Acronyms	8
1.4	How did we get here?	8
1.5	A Liquid-Cooled Future?	8
1.6	Drivers Towards Adoption of Liquid Cooled Solutions	11
1.7	Liquid and Air-Cooling Solutions Compared	13
2.0	Iceotope Liquid Cooling Technology	18
2.1	Company Profile	18
2.2	Iceotope Technology	18
2.3	Comparisons with other Forms of Liquid Cooling	23
2.4	Interface with Facility Water System	23
2.5	Dielectric Fluid Selection	24
3.0	The Liquid-Cooled Data Centre	26
3.1	Characteristics of the Liquid-Cooled Data Centre	26
3.2	Design Considerations	27
3.3	Dielectric Fluids	27
4.0	Performance Comparison	31
4.1	Objectives of the Comparison	31
4.2	Basis of Design	31
4.3	Operating Thermal Envelope	32
4.4	Locations	33
4.5	Energy Modelling	34
4.6	Notional Data Centre Comparison Scenario # 1 – The Hyper-Scale Developer	35
4.7	Notional Data Centre Comparison Scenario #2 – The Co-Location Developer	38
4.8	Notional Data Centre Comparison #3 Tropical Region Performance	41
5.0	Conclusions	44
5.1	Summary of Comparison Results	44
5.2	Conclusions	46

1.0

Introduction

1.0 Introduction

1.1 Objectives of the Report

The purpose of this report is to introduce liquid cooling for data-processing equipment, with a focus on the chassis-level precision immersion and Direct-to-Chip solutions produced by Iceotope Ltd.

Different forms of liquid cooling are described and compared, as are other hybrid forms of liquid to air cooling that are often used to support high density deployments.

The performance and energy efficiency advantages that liquid cooling can offer over conventional air-cooling are also discussed. These are highlighted in Section 4.0 of the report, which describes four imaginary Data Centres that are based upon real world scenarios. Calculated PUE and WUE performance figures and predicted OPEX savings are presented, along with an indication of the environmental benefits, in order to illustrate the advantages that an operator would achieve through the adoption of liquid cooling.

1.2 Authority

Authority to Undertake this report was provided by David Craig, CEO of Iceotope Cooling Technologies Ltd.

1.3 Abbreviations and Acronyms

The following abbreviations and acronyms have been used in the text of this report:

ITE Power	Electrical Power supplied to data processing equipment within the rack
ITE Rack Power	Electrical Power supplied to IT equipment racks.
PUE	Power Usage Effectiveness. The ratio of total facility energy to IT Energy (dimensionless)
WUE	Water Usage Effectiveness. The ratio of Annual Water Usage to IT Equipment Energy (L/kWh)
FWS	Facility Water System
TCS	Technology Cooling System
CDU	Cooling Distribution Unit
Direct Air-Cooling	The cooling potential of air drawn directly from atmosphere, either at ambient temperature, or chilled by other means, and supplied directly to the IT equipment racks.
Indirect Air-Cooled	The cooling potential of air recirculated across an air-to-air heat exchanger, where the primary source of cooling is ambient air and the secondary side airflow is supplied to the technical space.
CWS	Condenser Water System
CRAH	Computer Room Air Handling Unit
OCP	Open Compute Project – Collaboration of Hardware developers that share designs and intellectual property, to enable mainstream delivery of the most efficient designs for scalable computing.

1.4 How did we get here?

The design of data centre cooling systems has been steadily evolving over the last 25 - 30 years. During this time, air-cooling solutions – where heat is carried away from hardware components by a flow of cool air passing across the server or motherboard, have predominated. Whilst there has long been a role for the use of liquid – either water, or some other dielectric fluid – as the primary means of conveying heat away from the equipment, this is was typically only for highly specialised main-frame computer systems, produced by the likes of IBM and Honeywell-Bull. Most computer racks operated at comparatively low power density and there was no need for the increased capacity that liquid cooling could provide. Environmental conditions within the technical space were typically maintained within a relatively tight control band, by today's standards, and energy-intensive compressor-based cooling systems were the norm. There were no universally recognised design parameters for environmental conditions within the technical space and average PUEs were correspondingly high.

In the '90s and early '00's, things began to change. Average server power began to rise. The emergence of CMOS (Complementary Metal Oxide Semiconductor) transistor technology and later of dual and multi-core processors, allowed more efficient and powerful CPUs to be accommodated on server motherboards. The corresponding increase in heat flux density was addressed by the introduction of heat sinks and CPU cooling fans. 2004 saw the publication of the first edition of ASHRAE's Thermal Guidelines for Data Processing Environments by Technical Committee 9.9. This recommended higher server inlet air temperatures and a broadening of the thermal envelope than had hitherto been the norm, with the aim of promoting the exploitation of air-side economisation. The trend in increasing rack power has continued from an average of 4-5kW/rack 8-10 years ago, to around 7-10kW/rack or more, today.

Over the same period, there has been unprecedented growth in the number and size of data centres world-wide, driven by the expansion of the internet and the evolution of a digital society. In the early '60s, the Canadian philosopher Herbert Marshall McLuhan, predicted the creation of the World-Wide Web before it had been invented and first coined the phrase the 'Global Village' to describe an interconnected world in which electronic media would replace the visual printed word. In 1990, very few computers were connected to the internet and only 0.5% of the world's population had access to it. Today, it is estimated that up to 62% of the world's population is online (<https://www.internetworldstats.com/stats.htm>).

According to the International Energy Agency (IEA), global energy consumption by data centres now accounts for up to 200TWh of electrical power per annum, or around 1% of all global power consumption. The Uptime Institute cites a higher range of 400-500TWh, representing around 2% of global consumption. Many data centre operators do not disclose the associated water consumption of their cooling systems, although it has been estimated that approximately 1.8 litres of water are used by the average data centre for every 1kWh it consumes. Based upon the IEA's more conservative estimate of annual data centre power consumption, this equates to $3.6 \times 10^{11} \text{ m}^3$ of water consumed globally, each year.

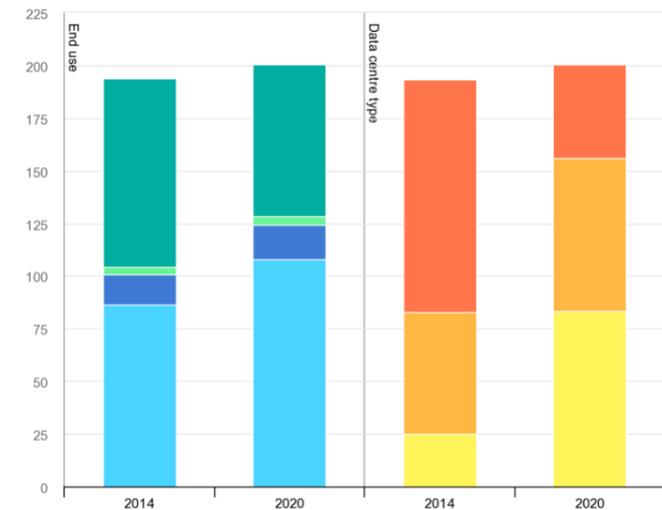


Figure 1.1: Global Data Centre Energy Consumption TWh (2014-2020) (<https://www.internetworldstats.com/stats.htm>)

1.5 A Liquid-Cooled Future?

The insatiable demand for ever more data and more powerful processors to stream, handle and manipulate it, is showing no sign of slowing down. This, and a combination of other technical and environmental factors, are combining to bring about conditions that will promote an inevitable shift towards the widespread adoption of liquid cooling:

- With increasing CPU power, the rise in processor heat flux density is approaching the limit of what can be satisfactorily addressed with air-cooling. Processors cannot physically get much larger than they already are, due to spatial constraints on the motherboard. So, unless a fundamentally new type of chip is developed that improves efficiency, any further performance gains can only be achieved by further increases in power. However, this may result in heat flux density being pushed to levels that risk over-heating. CPU vendors are therefore faced with no alternative but to seek more effective forms of cooling.
- Coupled to these practical constraints, emerging applications, including Artificial Intelligence, Edge and graphical application requiring high performance GPUs, are putting ever-greater demands on server performance. Trying to accommodate these high-density requirements in an air-cooled environment would be both difficult and inefficient from a spatial, cooling and power perspective. This is especially so in compact Edge data centres, where the overall footprint of the facility may be constrained in a tight urban setting.
- Legislators are aware of the environmental impact of data centres and the huge amounts of power and water that they consume. Sustainability targets imposed by Governments and Municipality Planning Authorities are becoming ever more stringent and there is also growing public pressure for Data Centre Owners to clean up their act. The UK was the first major economy to pass legislation to bring all greenhouse gas emissions to net zero by 2050 and the EU has also set a target of becoming climate neutral by 2050. Given the significant improvements in average PUE that have

taken place over the last few years, it is difficult to see how further step-change improvements in the energy efficiency of air-cooled systems are going to be achieved. The cooling and electrical infrastructure serving new-build data centres are typically already energy efficient. However, in this regard, liquid cooling offers some tangible advantages:

- One area that is likely to see development is the exploitation of data centre waste heat recovery (WHR), especially in localities that are close to major conurbations, where district heat infrastructure may already exist. The adoption of liquid cooling would greatly simplify that process and improve its efficiency, as there would be no need to harvest heat from the cooling system airflow and the facility water-system (FWS) can operate at temperatures far greater than any air-side cooling system.
- Even where WHR is not feasible, liquid cooling is demonstrably more energy efficient and cost-effective than many other forms of cooling. ASHRAE TC 9.9 has proposed Liquid Cooling Classes W1 to W5, which are indicated in Table 1.1 below. for the design of FWS infrastructure. Under North American and Northern European climatic conditions, all Classes can exploit free cooling to some degree. Even under peak summertime ambient conditions, Classes W4 and W5 would allow heat rejection entirely by dry-air coolers, or more efficient evaporative cooling units, or cooling towers. Energy-intensive compressor-based mechanical cooling would not be required.

ASHRAE Liquid Cooling Class	Facility Water System Temperature Range
W1	2°C - 17°C
W2	2°C - 27°C
W3	2°C - 32°C
W4	2°C - 45°C
W5	>45°C

Table 1.1 ASHRAE Liquid Cooling Classes

(Source: TC9.9 Liquid Cooling Guidelines for Datacom Equipment Centers 2nd Edition. ASHRAE 2014)

Server liquid cooling systems may be categorized into two broad groups: Immersion Cooling and Direct to Chip/Cold Plate Cooling. Some vendors of in-rack or in-row cooling units have also sought to describe their products as high-density liquid cooling solutions, but the key difference with this type of equipment is that the cooling fluid is used to cool air, which in turn cools the server. So, whilst in-rack and in-row may certainly boost local cooling capacity to address higher rack cooling loads, these are not liquid cooling solutions.

The image below summarises the current categories and sub-categories for liquid cooling:

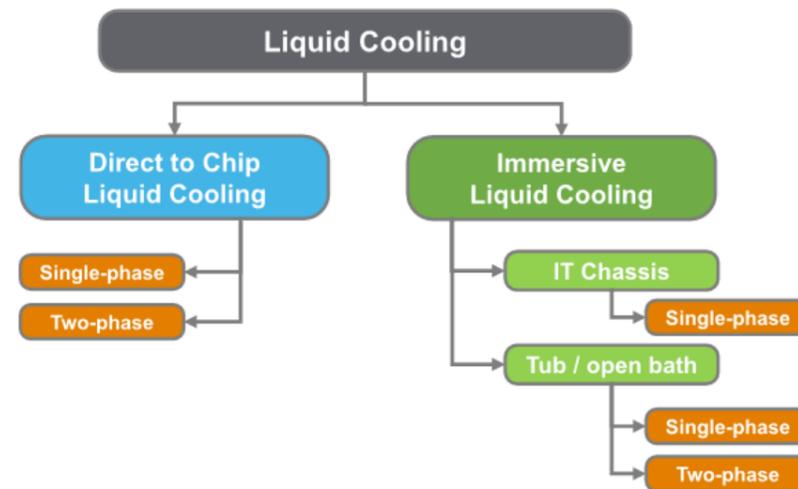


Figure 1.2: Liquid Cooling Categories (Source: Schneider Electric)

1.5.1 Immersion Cooling

Immersion Liquid Cooling refers to tub or open bath systems. With this method, the IT equipment is completely submerged in the fluid. Servers are typically installed vertically within a liquid filled tub. The system can be realized using either single phase or two-phase coolant.

In the single-phase system, the heat within the dielectric fluid is transferred to a water loop via heat exchanger(s), either by means of an internal circulation pump or by natural convection. The heat exchanger can be located either within the tub or externally and is connected to the FWS or to a fluid cooler. The typical cooling medium is oil based dielectric fluid and is always maintained in the liquid phase. Figure 1.3 and Figure 1.4, below, illustrate the operating principles.

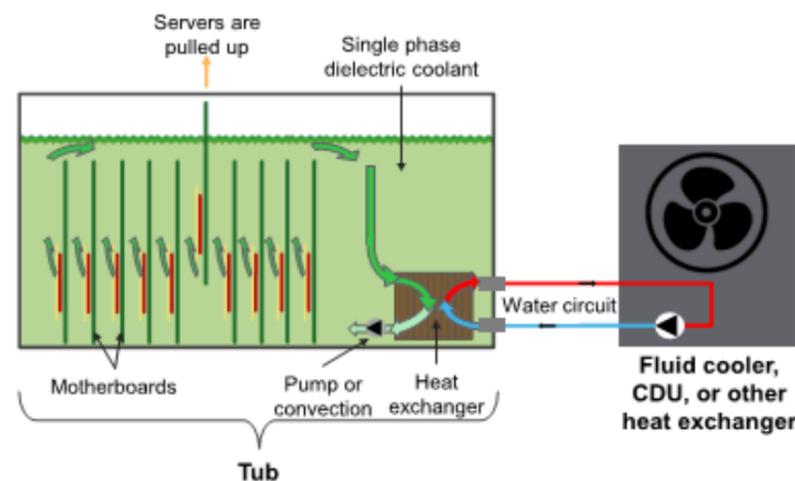


Figure 1.3- Immersion Cooling Tank with Single Phase Dielectric Fluid (Source: Schneider Electric)

In the two-phase system, the electronic components are completely submerged in a two-phase dielectric coolant. The fluid effectively acts as a refrigerant. Heat produced by the CPU is absorbed by the fluid, causing it to change from the fluid

to vapor phase. The vapor is then passed through a heat exchanger to be condensed back into its fluid phase before being returned to the tank.

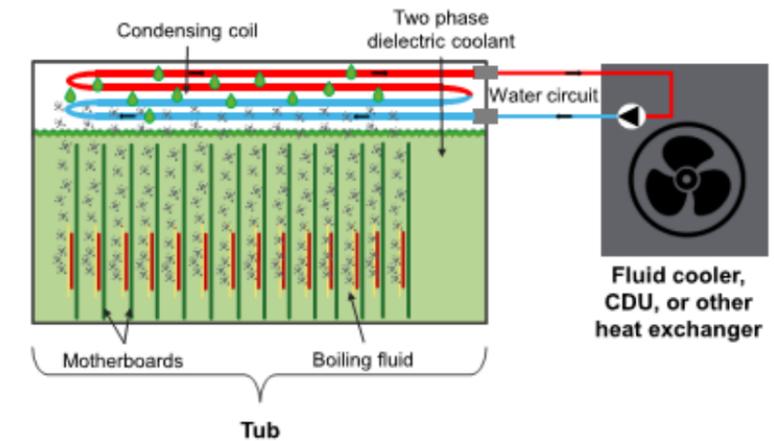


Figure 1.4: Immersion Cooling Tank with Two-Phase Dielectric Fluid

(Source: Schneider Electric)

- Both these solutions offer high capacity cooling of the electronic components and maintain the equipment at a steady temperature, by virtue of full immersion within the tank.
- Both utilise a relatively high volume of dielectric fluid, to fully immerse the electronic components.
- The tank configurations are typically of a proprietary, vendor specific form factor that cannot easily be integrated into the rack aisle layout of a conventional air-cooled data centre, although it is not uncommon for high density liquid-cooled racks to be installed alongside low density air-cooled equipment.
- The two-phase fluid solution offers very efficient cooling, by virtue of the latent heat absorbed during the change from liquid to vapor phase. However, the low temperature and pressure at which the fluid will evaporate means that the enclosure must be tightly sealed to prevent loss of vapor to surroundings. If the enclosure needs to be opened to carry out maintenance on the equipment, then some fluid loss is inevitable.
- Servers requiring maintenance must be physically removed from the bath of dielectric fluid. With single phase fluids, the equipment is usually placed onto a special service tray that collects any fluid run-off and drains this back into the tank. Care needs to be taken to avoid spillage onto the surrounding floor, which could potentially be hazardous and would certainly be a nuisance within the data hall environment.

1.5.2 Chassis-Level Precision Immersion Cooling

Chassis-Level Precision Immersion Cooling refers to a cooling strategy in which the liquid coolant and IT equipment are fully contained within a chassis-type enclosure. The CPU and other electronics components are partially or fully submerged in the dielectric fluid, which is circulated over the equipment by pumps. Capillary tubes and conduits are arranged to ensure a steady flow of coolant across heat producing components, thereby ensuring effective cooling.

The warm fluid accumulates in the bottom of the chassis, from where it is drawn by the pumps and passed across a heat exchanger, the primary side of which is cooled by the FWS. The cooled fluid is then re-circulated back to the equipment. With this typology of liquid cooling, all fans within the server can be removed and all the electronics are placed in an environment which is inherently slow to react to any external changes in temperatures. The system is also immune to the influence of humidity and pollutants. The IT chassis coolant circuit is hydraulically separated from the facility water system by an embedded heat exchanger and an internal micropump provides coolant circulation. The absence of any fans allows the system to operate in near silence.

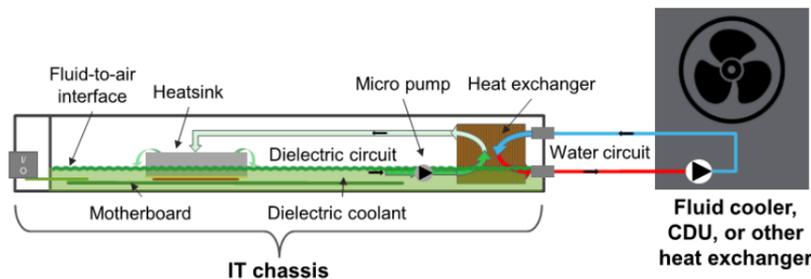


Figure 1.5: Chassis-Level Precision Immersion Cooling.(Source: Schneider Electric)

1.5.3 On-Chip and Cold-Plate Technologies

On-Chip cooling is realized by taking the coolant directly to the hottest components (CPUs, GPUs, etc.) via a cold plate heat exchanger that is in direct contact with the chip sets on the server. The IT components are therefore physically separated from the coolant liquid which can be either water or dielectric fluid. This method requires air cooling fans to dissipate residual heat from remaining components within the IT server.

Single phase on-chip systems use water as coolant; the fluid circuit can be connected to either to the Facility Water System or a dedicated Technology Cooling System to a fluid cooler, as shown in figure 1.6 below, or to a water-cooled CDU, which is served by the FWS.

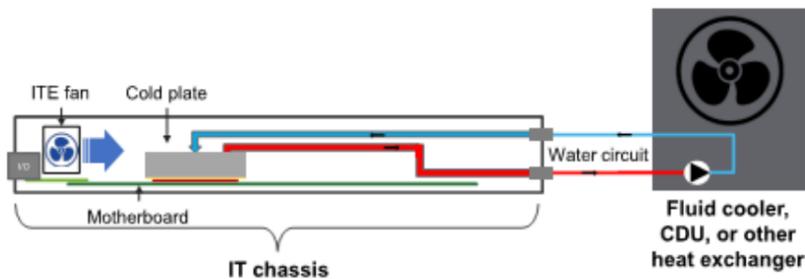


Figure 1.6: Cold Plate Single Phase Cooling (Source: Schneider Electric)

Two-phase on-chip cold plate systems use engineered dielectric fluid or refrigerant. The fluid in liquid phase captures the heat from the cold plate and in doing so, evaporates into the vapor phase. The hot gas is passed to an external condenser / CDU to release its heat and be condensed back into its liquid state, before being returned to the cold plate. The refrigerant circuit may either be pumped or operate on the thermo-siphon principal. Figure 1.7 below illustrates the operating principles of the system.

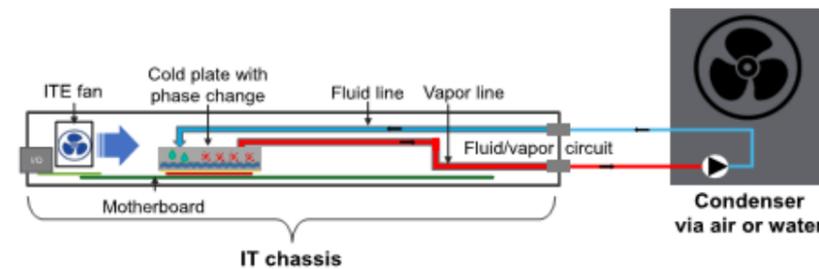


Figure 1.7: Cold Plate Two Phase Cooling (Source: Schneider Electric)

- Both Cold Plate solutions offer high capacity cooling directly to the heat producing components, whilst avoiding the complexity and inconvenience for maintenance arising from full immersion of the server.
- Whilst the cold plates can deal with comparatively large power loads, they typically do not have the cooling capacity of full or chassis-level precision immersion systems.
- The need for a conventional air-cooling system to dissipate residual heat from other components remains. The data hall would therefore have two cooling systems, rather than one, leading to increased installation costs. However, this could be an effective solution, when a limited number of high-density racks need to be deployed within a conventional air-cooled data hall that would otherwise be unable to support them.
- The effectiveness of the system is highly dependent upon maintaining good surface contact between the cold plate and the CPU. This is generally achieved by means of a heat conducting paste that maintains good conductivity and can tolerate the flexing and movement that may occur due to temperature changes in the CPU.
- The cold plates would offer very little in the way of thermal capacitance in the event of an interruption in the flow of cooling fluid and unless means of maintaining the continuity of cooling were provided, such as UPS-backed circulating pumps, then a power failure affecting the mechanical plant could result in rapid over-heating of the IT equipment.
- With water-cooled cold plates, any water leakage from the cold plate, or its connecting pipework within the rack would likely result in a catastrophic failure of the server.

1.5.4 Hybrid Solutions

In-Door Cooling

An In-Door or Rear-Door Cooling system comprises a heat exchanger coil and EC fan array that is mounted on the rear door of the server rack. The fan array provides additional air handling capacity to pull a large volume of air through the server, to satisfy the high-density cooling load within the rack. The rear door mounted cooling coil cools the discharge air, so that the rack does not reject its heat into data hall. The cooling coil may be connected to the Facility Water System or be filled with refrigerant and connected to a remote DX condenser. This could either be linked to the FWS for water cooling or be externally mounted and air-cooled.

Figure 1.8 below illustrates a typical arrangement.



Figure 1.8: In Door Cooling (Image source: Stulz)

- With this technology, no liquid is introduced in the IT rack, although in the case of water-cooled rear-door units, water is present in the data hall. Some DC Operators who are sensitive to bringing water into their data halls, opt for DX versions for this reason.
- Cooling to the IT equipment is provided by airflow, so although rear-door cooling can support high density applications, the equipment is not liquid-cooled, and the system cannot match the densities of immersion systems.
- As the air discharged from the rack has been cooled, the technology is less suitable for integration into hot or cold aisle architecture. In an open architecture hall, the cool discharge air can be recirculated back into the rack.
- To ensure resilience, some Rear Door units are provided with hot swappable redundant fans.
- Because the cooling medium is air, the capacity of In-Door cooling is limited by the amount of air that can be physically circulated through the rack. Cooling capacities are typically up to approx. 30-35kW/rack.

In-Row Cooling

An In-Row Cooling unit consists of a heat exchanger coil and EC fan array mounted within a cabinet that is sized to be installed within a row of IT racks. The coil is connected to the facility water system or DX circuit. The hot air is drawn into the cabinet from the rear (hot aisle) side of the unit by the fan array, passes through the cooling coil and is then supplied back into the cold aisle through a discharge louvre on the front of the unit. Typically, In-Row units are deployed in pairs within each aisle, to ensure redundancy.

Figure 1.9: In Row Cooling (Source: Stulz) below illustrates a typical unit.



Figure 1.9: In Row Cooling (Source: Stulz)

- In-Row units are intended for deployment within cold/hot aisle arrangements. The introduction of hot or cold aisle containment can improve efficiency, by eliminating short-circuiting, although this would also be dependent upon maintenance of good airflow management within the racks, especially those that are only partially populated.
- As with In-Door Cooling technologies, no liquid is introduced into the IT rack. Cooling air is supplied to the IT equipment via the cold aisle. Water is present in the data hall, although this can be avoided by using In-Row units fitted with DX coils.
- In Row units can be deployed within selected aisles of a conventional, air-cooled data centre, to provide a localised boost to the hall's cooling capacity, to support clusters of high-density racks. In this sense, In-Row can offer a good retro-fit option, for technical space that might otherwise be incapable of supporting high-density deployments.
- A draw-back of In-Row cooling is that it occupies rack space within the aisle, thereby reducing the amount of IT equipment that can be accommodated. For bespoke hall layouts, designed around clusters of high-density racks, this may not be an issue. But for a CoLo Operator, or Hosting Company concerned about maximising the amount of lettable rack space within their facility, this could represent a loss of income.
- In-Row units are available for several vendors and typically available in two frame widths: 300mm and 600mm, and range in capacity from 10kW up to 60kW, for the full width units.
- As with In-Door cooling units, the cooling medium is air, so the capacity of H-D racks that can be supported is limited by the amount of airflow that can be delivered to them.

HAC Cooling Solutions

Hot Aisle Containment Cooling consists of a hot aisle containment system with the addition of integral cooling heat exchangers and fans. The operating principle is like that of In-Row cooling; warm air is drawn from the contained hot aisle and is cooled by passing it across a cooling coil. It is then delivered back into the cold aisle for recirculation back through the IT equipment racks. The

difference is that the fan array and cooling coils are contained within modules mounted above the hot aisle (item 108 in the illustration below).

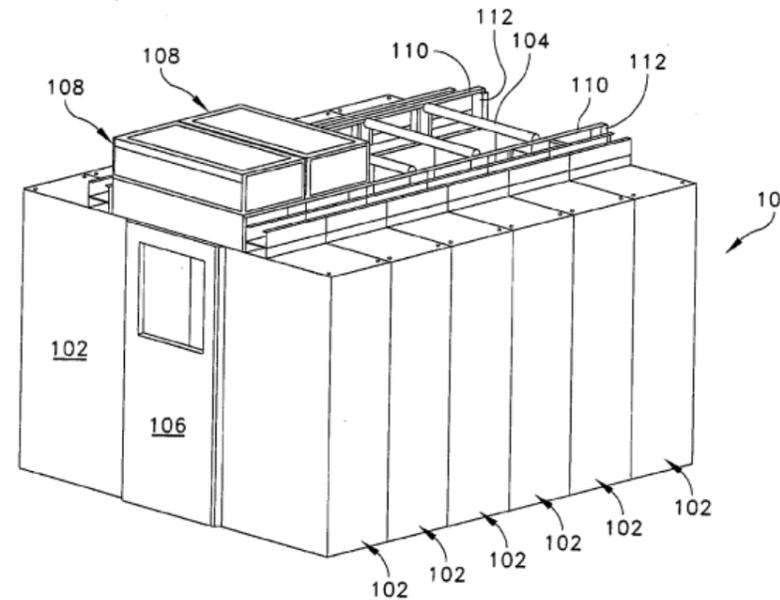


Figure 1.10: Hot Aisle Containment (Source: Schneider Electric)

- An advantage of this arrangement over in-row cooling is that there is no loss of rack space within the aisle. As with In-Row cooling, it is intended to be a supplementary means of meeting increased cooling demands for high density equipment.
- As with In-Door cooling units, the cooling medium is air, so the capacity of High Density (H-D) racks that can be supported is limited by the amount of airflow that can be delivered to them.

1.6 Drivers Towards Adoption of Liquid Cooled Solutions

1.6.1 Applications

Although historically associated with specialist and academic users, Liquid Cooling is now being considered for a diverse variety of mainstream applications:

- **Cloud and Hyperscale** – Given the size and scale of hyperscale deployments, there is increasing need to optimise energy and water efficiency. While still predominantly reliant upon direct air + evaporative cooling, the extent of their operations is now pushing Hyperscale owners to consider alternative ways of reducing their operating costs and carbon footprint. Exacerbating this challenge is the increase in applications requiring high-density, high-power racks, for data analytics and graphical streaming. If placed in conventional air-cooled data halls, these racks would consume a disproportionate amount of cooling airflow, which would be problematic for surrounding racks. It would also be wasteful, in terms of space utilisation within the data hall. Deploying a mix of air and liquid-cooled racks within the same environment instantly resolves these issues.
- **Enterprise and CoLocation** – Enterprise and CoLocation Operators have become adept at offering their customers a reliable, competitively-priced

environment for their IT Infrastructure, based upon a traditional, air-cooled approach. However, the growth in customers that are requesting support for Artificial Intelligence and High-Performance Computing racks is having a disruptive effect upon existing data centre facilities, which were not designed to support such high levels of power and cooling. As a short-term solution, some CoLo Operators have investigated the feasibility of scattering high density racks around the data hall, to mitigate against the impact they have on surrounding equipment. However, this is not ideal solution as H-D racks often need to be deployed in clusters, for data sharing, or other purposes. The impact this is having is driving some CoLo Operators to consider retrofitting existing data halls to be able to support a percentage of high-density racks. In new facilities, they are looking at creating dedicated, liquid-cooled suites within their data centres. This is also seen as a way of soaking up any 'trapped capacity' e.g. spare electrical power that is available to the facility, but which has not been utilised.

- **Edge Computing and IoT** – This is a very rapidly expanding sector. In recent years, there has been a proliferation in the numbers and types of connected devices being used in both commercial and domestic context, from handheld 'smart' phones, intelligent CCTV cameras and terminals, to intelligent devices intended for home use. Trains, cars, and commercial vehicles are also increasingly being provided with on-board capabilities. To optimise the performance of these devices and reduce dependency upon the speed and reliability of the network connection, it makes sense for the necessary data storage and processing capacity to be provided as close to the users as possible, in the form of Edge-based systems. These are typically small, remotely managed installations, mostly in urban environments, but also in more remote locations, such as along arterial transportation routes. What they have in common is the need to operate autonomously and provide a high degree of reliability - sometimes in quite harsh environments – whilst only requiring very low levels of on-site maintenance. Liquid cooling is ideally suited to this application, in that it can easily handle the enhanced cooling loads that may arise, especially with Edge-AI. Chassis level precision immersion cooling also encloses sensitive electronic components in a secure containment, protected from atmospheric pollution and humidity. The higher operating temperatures of liquid-cooled systems allow heat rejection via simple dry-air systems, avoiding the need for more complex compressor or evaporative-based cooling systems. The steady operating temperature of the liquid cooling fluid also promotes reliability, by avoiding temperature swings that can occur with direct air cooling.

1.6.2 Rising Power Density

All CPUs need to perform their tasks as efficiently as possible, to consume the minimum amount of power and avoid overheating. Historically, to increase the processing capacity of their products, chip manufacturers have added transistors and increased clock rate, both of which resulted in increased heat output. This is problematic, as it reduces cooling efficiency and induces thermal stresses in the processor.

Figure 1.11 below shows the expected future trend in heat load growth for standard air-cooled 42U racks, for different applications. Apart from

communications, all sectors are expected to increase, with scientific and HPC cooling loads increasing significantly.

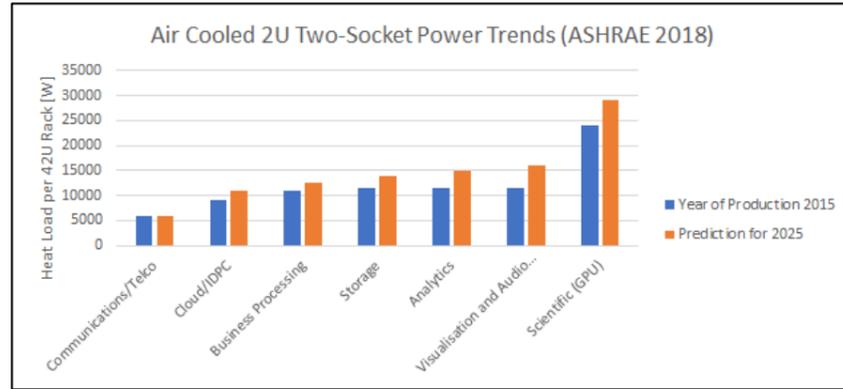


Figure 1.11: Air-Cooled 2U two-socket server power trends (Source: TC 9.9 Liquid Cooling Guidelines for Datacom Equipment Centers. ASHRAE 2014)

In recent years, the trend in rising chip power has been mitigated to an extent by the introduction of multiple cores. This feature allowed the execution of multi-threaded code, which increased chip capacity without incurring a corresponding increase in power demand and heat output. However, there are limits to how far this strategy can be pursued and to achieve further step-changes in performance, chip manufacturers have had to increase power, leading to a further rise in heat output. Furthermore, miniaturisation and other developments in processor design have given rise to unequal heat production across the chip's surface, leading to 'hot spots' and asymmetrical cooling requirements.

1.6.3 Energy and Water Efficiency

With an air-cooled data centre, increased rack power may lead to a corresponding rise in infrastructure power consumption, due to greater fan power demand and/or the presence of supplementary cooling systems, such as in-row cooling introduced to locally boost cooling in the vicinity of high density racks.

However, the adoption of liquid cooling may help to offset, or reverse this trend:

- Chassis or full-immersion liquid cooling eliminates the need for large scale air-cooling systems and delivers cooling to the IT equipment in a much more energy-efficient manner. Compared with air at sea level, water is over 784 times denser and has a specific heat capacity that is over 4 times greater. Hence, water has a volumetric heat transfer capacity that is over 3,000 times greater than that of air.
- In a conventional air-cooled data centre, where chilled water is used to cool the air that is to be supplied to the racks, a heat exchange process takes place at the CRAH cooling coil. Power is then consumed by the CRAH fans that supply air to the racks and further power is consumed by the cooling fans inside the servers. With liquid-cooled, where the water is supplied directly to the racks, the energy consumption associated with the CRAH and server fans is eliminated. Some parasitic power consumption is consumed by internal fluid pumps that circulate the dielectric fluid over the chip, but this is minimal in comparison to the server fan power consumption.

- The higher operating temperatures at which liquid cooling systems typically operate will almost certainly reduce the need for evaporative cooling and heat rejection. This will have a beneficial impact upon water consumption and WUE.

1.6.4 Facility Footprint as a Function of Power Density

- Increased Power Density.** Liquid Cooling Technology allows superior heat rejection rates from the IT equipment. A consequent advantage, compared to air cooled solutions, is the prospect of accommodating a much greater ITE capacity within the same data hall area than would normally be possible with air cooled servers, due to their lower power density. This results in fewer racks being required to achieve the same overall ITE, or if additional power is available to the site, of a greater ITE capacity being accommodated within the same data hall footprint. Hence, there is an opportunity to either achieve a smaller building envelope, or to maximise the potential of an otherwise comparatively small site.

As an example of this, Table 1.2: Relationship between Power Density, Rack Count and Data Hall Size, and Figure 1.12: Curve showing the Relationship between Rack Power Density and Data Hall Area, illustrate the relationship between Rack Power, Rack Count and Data Hall Area, for a notional 1.0MW data hall:

Data Hall ITE: [kW]		1000
Rack Footprint [m ²]		0.7308
Power Density [kW/rack]	Racks [Quantity]	Data Hall Area Occupied by Racks [m ²]
5	200	146
8	125	91
15	67	49
20	50	37
30	33	24
45	22	16
60	17	12

Table 1.2: Relationship between Power Density, Rack Count and Data Hall Size

The number of Racks required, and the space needed to accommodate them within the data hall decrease with inverse proportionality when the rack power density is increased.

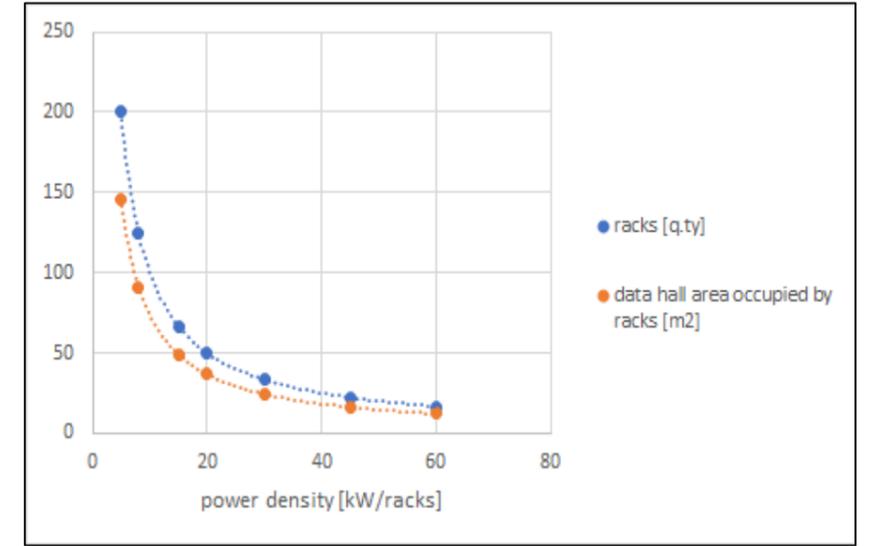


Figure 1.12: Curve showing the Relationship between Rack Power Density and Data Hall Area

- Plot Area vs. Power Density.** As introduced in the preceding paragraphs, Data Hall size can potentially be reduced by increasing the rack ITE capacity. Many of the spaces within the Data Centre, e.g. plant, ancillary, admin and support areas – would be unaffected by this. However, if increased density is carried through an entire multi-hall facility, then it is reasonable to state that this could reduce the overall footprint of the building. This unlocks the possibility of either using a smaller campus plot or boosting the capacity of a site that would otherwise be limited in the ITE that it could accommodate. This could lead to very substantial savings, in terms of construction and land cost.

1.6.5 Opportunities for Waste Heat Recovery

- Heat Capture –** The elevated temperature at which many liquid-cooled installations operate, particularly in the W4 and W5 classes, offers great opportunities for the capture of waste heat that would otherwise be released to atmosphere. This could be delivered into a local district heating network via an interposing heat exchanger. Subject to the operating temperature of the network, the injected heat could either be boosted up to the required supply temperature, via heat pumps, and delivered directly to customers, or used to pre-heat cool return water before it is sent back to the district heating energy centre, and thereby reduce the load on centralised heating plant. This has the potential to be a 'win-win' arrangement for both parties; not only would local users of the district heating network benefit from the supply of heat from the facility, but the facility itself would benefit from being able to shut down heat rejection plant that would otherwise be needed to cool the FWS. Given the pressures to reduce carbon emissions and the challenging targets some national governments have set to achieve Net Zero Carbon, Municipality Planning Authorities are increasingly looking at large-scale data centres as a potential source of low-carbon, waste heat and are encouraging this approach by introducing Planning Conditions to require data centre operators to facilitate waste heat recovery. This is particularly the case in some Nordic Countries, where many of the larger

conurbations have established district heating infrastructure and grid supplied power contains a high proportion of renewable energy.

- **Heat to Facility** – Recovered waste heat can also be used within the facility itself. Admin, ancillary and support areas may well require heating during the winter season. This could be provided by waste heat recovery from liquid-cooled racks with the minimum amount of additional infrastructure. Such an arrangement would assist in occupied area of the facility meeting Building Energy Codes.

1.7 Liquid and Air-Cooling Solutions Compared

1.7.1 Flexibility and Deployment

- **Server Configuration and Form Factor.** Liquid-cooled servers can be built from the same electronic components originally produced for air cooled applications. It is also possible to convert air-cooled servers to liquid cooling, although this is not common practise. With air-cooled servers, the layout and orientation of components must be optimised for the passage of cooling airflow across the motherboard. This imposes restrictions on the layout that do not apply to liquid-cooled servers, allowing them to be more compact. The Iceotope KU:L Liquid-cooled servers can be installed in a standard 1U rackmount, although the larger KU:L Box Cubbies take up 4U.

Unit	Rackmount Dimensions (W x H x D)
1U	19" x 1.75" x 17.7" 19" x 1.75" x 19.7" 19" x 1.75" x 21.5"
2U	19" x 3.5" x 17.7" 19" x 3.5" x 20.9" 19" x 3.5" x 24"
3U	17.1" x 5.1" x 25.5"
4U	19" x 7" x 17.8" 19" x 7" x 26.4"
5U	19" x 8.34" x 19.67" 19.1" x 8.75" x 26.4"
6U	19" x 10.5" x 19.5"
7U	17" x 12.2" x 19.8"

Table 1.3 - Standard Rackmount Dimensions

- **Rack Level Deployment.** Standard Open Compute Project (OCP) V2 Racks, commonly adopted for air-cooled solutions, are designed to also accommodate liquid-cooled technology. By keeping the front access zone free to operate the servers, flow and return TCS manifolds and connections can be installed to the rear, in what would otherwise be the 'hot' aisle. Designated 'hot' and 'cold' aisles become redundant, as there is no airflow through the racks and no need for any kind of containment. However, it is recommended to adopt a conventional data hall rack layout, with the front and rear sides of each row of racks facing on to one another, to form 'Access' and 'Service' aisles, the latter being for TCS pipework headers and connections to each rack. This is not only the most spatially efficient way of arranging rack, but as with air-cooled data halls, it facilitates cooling

distribution, and provides for the most efficient layout of electrical power busbars and fibre trays. The only differences are the omission of airflow containment systems (providing valuable CAPEX savings) and the provision of TCS pipework headers running down the length of the Serviced aisle.

1.7.2 Power Density and Operating Temperature

- **Power Density Limitations of Air-Cooling vs. Liquid Cooling.** Air cooled racks are commonly designed to support average power densities in the order of 7 to 10 kW per rack, resulting in data hall power densities of roughly 2.5-3.5 kW/m². Higher power densities are achievable, but difficulties may arise when this is combined with conventional hot and cold aisle architecture, as the air velocity and pressure within the cold aisle may become excessive, leading to problems with air distribution and leakage. To overcome these issues, it may be necessary to increase cold aisle widths, or provide supplementary cooling in the form of rear-door or in-row cooling units. These measures tend to increase the spatial requirements of the racks, resulting in a less efficient layout and fewer racks in the hall. Although individual rack power levels may have increased, the overall power density of the data hall has not appreciably changed. Liquid-cooled racks are not subject to the airflow constraints and can support much higher power densities without detriment to cooling distribution.
- **Operating Temperature of Components and Cooling Medium.** The average idle temperature for microprocessor chips is 45-50°C, rising to 80-85°C when operating under full load for a prolonged period. As a rule, processors perform better at lower temperatures and incorporate features to protect themselves against damage, due to overheating. Modern, multi-core Intel CPUs operating at full load in 'Turbo Boost' mode (i.e. over-clocking at faster than the marked frequency) can reach operating temperatures of up to 100°C, but beyond this level performance starts to be affected. Sensitive electronics components have a finite lifespan and running them at higher temperatures tends to shorten it. For this reason, the 'Tjunction Max' setting (the maximum thermal junction temperature that a processor will allow prior to using internal thermal control mechanisms to reduce power and limit temperature) for Intel processors, is generally set at 100°C.
- **Cooling to Component.** With air-cooled systems, the cooling air entering the server is controlled at the temperature of the cold aisle (up to 27°C for TC9.9 Recommended Range), whereas with liquid-cooled systems, the temperature of the dielectric fluid in contact with the electronic components is much higher. For the Iceotope KU:L Box solution, this is >50°C. However, as the dielectric fluids used for liquid cooling have much greater thermal capacity than air, they are more effective in dissipating heat. Processors run cooler and therefore more efficiently.
- **Delta T and its impact upon component reliability.** With the KU:L Box semi-immersive approach, the Delta T of the coolant being circulated around the microprocessor chips is in the order of 5-6K. This provides for a very uniform distribution of cooling and avoids the risk of steep temperature gradients across the surface of components. Whereas with air-cooled servers, there is a tendency for cooling airflow to become hotter as it passes across the server, leading to uneven cooling and potential hot

spots. Even with cold-plate systems, where the electronic components are not fully immersed, heat is conveyed away from the largest heat source – the processor – more effectively than with air, allowing it to run cooler.

- **Extreme and harsh environments.** Liquid cooling is well suited for use in extreme and harsh environments, as sensitive electronic components are protected within a secure enclosure. There is no risk from airborne contamination or humidity fluctuations, even when the equipment is deployed into heavily polluted environments.

1.7.3 Free Cooling

Free Cooling is one of the most commonly used strategies to save energy and is driven by the difference between the prevailing ambient dry bulb air temperature, and the cooling medium; be this either air (be that direct, or indirect), or water.

Free cooling is used extensively in data centres and can be very advantageous, but is nonetheless subject to constraints:

- Direct Air-Cooling, where the equipment is cooled directly with outside air, is a mainstay for many Hyperscale Operators, as they can accept supply air temperatures in the Allowable Range and a wide range in internal humidity (e.g. 20-80% rh). This strategy can be very effective when deployed in regions with a temperate, sub-tropical climate, where the ambient air temperature falls typically within the supply air envelope. Under peak summertime conditions, when the ambient temperature exceeds the desired supply air temperature, the incoming air can be cooled with wetted media. However, this strategy is less effective in hot, humid tropical regions, where the ambient temperature may be much higher than the supply condition and the use of direct evaporative cooling is ruled out, due to high prevailing humidity. In hot, arid areas, the amount of water needed to maintain evaporative cooling can also be excessive and prohibitive. For this reason, most data centres built in these regions rely upon energy intensive mechanical cooling systems (e.g. chillers).
- For CoLo and Enterprise Operators, Direct Air Cooling is problematic, as most of their customers will require that the environment in the Data Hall conforms to the TC 9.9 Recommended Range, or something close to it. They may also be unwilling to accept the risk of contamination from polluted outside air. Even in temperate Northern European locations, this factor precludes the use of Direct Air Optimisation, as summertime humidity levels are too high. Hence, many of these Data Centre operators will use either Free-Cooling Chillers, or Hybrid Indirect Air Cooling (IAC Units) that incorporate both evaporative and mechanical trim cooling. These are both very effective cooling strategies but are more energy intensive than Direct Air Optimisation or Liquid Cooling.

As can be seen from Table 3.1: ASHRAE Liquid Cooling Classes, the operating temperature of Facility Water Systems associated with liquid-cooled racks can be up to 25°C warmer than the cooling medium used in an air-cooled data centre. For example, the operating range for a chilled water CRAH may be ~16-25°C, whereas the FWS serving an Iceotope KU:L Box operates at 40-46.5°C. The warmer operating temperatures of liquid cooled systems compared to air-cooled systems permits:

- Better external cooling system heat exchange efficiency and potentially smaller heat rejection plant.
- Free Cooling availability for a much larger proportion of the year, possibly all year round in some localities.
- Possibility of selecting a larger delta T between flow and return cooling water, resulting in reduced cooling water flow rates and lower pumping power consumption.
- Possibility of drastically reducing dependency upon mechanical and evaporative cooling, allowing omission or reduction in chiller plant size (chillers may still be required to support other systems) and a significant reduction in evaporative cooling water storage infrastructure, allowing energy and cost savings. Note that a reduction in compressor power requirement would also allow more of the available site power supply to be diverted to increase the available ITE capacity of the facility and/or permit smaller electrical infrastructure.

It is therefore possible to expect a significantly better heat rejection power economy by adopting liquid cooled solutions compared to air cooled solutions on equal cooling demand terms.

1.7.4 Evaporative Cooling, Water Usage, Storage, Treatment and Blowdown.

Evaporative Cooling passively exploits the water liquid to vapor transition to adsorb the energy from air and consequently lower its dry bulb temperature. It is generally convenient from a power consumption point of view because it reduces the need for energy intensive mechanical cooling (compressor power). However, water pumping sets and water treatments plant will still be needed, all of which will consume electricity.

Evaporative Cooling strategies finds a convenient application in arid climates where the cooling effect produced by water vaporisation can be maximised. Nevertheless, the process is entirely dependent upon a wholesome water supply being available on site. Constraints relating to security and capacity of water supply must be considered and usually lead to the adoption of some precautions:

- Water needs to be stored on site to allow evaporative cooling systems to function in the event of an interruption in the supply, or to act as a buffer, to make up for any shortfall in supply capacity, during periods of peak consumption.
- Water should be treated before being stored and used in evaporative cooling systems. During winter months, when the demand for evaporative cooling may drop off dramatically, untreated stored water will rapidly become a liability and a potential hazard, giving rise to bacterial and algae growth within the storage and distribution system. Special measures may be required to avoid stagnation within storage tanks and pipework, and to protect these from frost damage. The safety of operatives is also a paramount consideration. Considering the risks, some operators simply drain down their entire systems, accepting the attendant water loss and reduction in WUE.
- Evaporation of process water will cause mineral content concentration to cycle up within the system. To prevent scaling and other damage to the

equipment, periodic down-down is normally required, to allow the mineral content to be removed and diluted with fresh make-up water. To obtain better utilisation of the available supply and to stabilize system water ph, it is not uncommon to provide softening and/or acid dosing to the supply, or to remove mineral content via reverse osmosis. All these processes incur capital and operational costs and consume water in themselves. Some of the treatment chemicals used are highly corrosive and must be handled with care, by trained personnel.

- For a large scheme, the volume of wastewater discharged to the public sewer under peak conditions can be considerable. If the Process Water system has been treated with chemicals, then discharge to surface water drainage soakaways or watercourses is not an option. This may trigger the requirement for a costly upgrade of the municipality's foul drainage infrastructure. Mitigation measures include wastewater recycling, which uses RO to remove the mineral content from the discharge water. However, this generates further costs, especially if there are limitations on the ph or conductivity of water that may be discharged to the sewers, as this may generate the need for further treatment.
- Cooling Towers associated with chillers or large-scale fan coil wall systems also make use of evaporative cooling, but are optimised to cool water, rather than air. They are nevertheless subject to the same issues and constraints.

The higher operating temperatures associated with a FWS serving liquid cooled racks mean that in many climates, dependency upon evaporative cooling may either be reduced to a narrow window of operation during the peak summertime period or eliminated altogether.

1.7.5 Comparison of PUE and WUE

Infrastructure power demands are generally driven by several factors; primarily the choice of cooling system and the climatic conditions under which it has been designed to operate. As a rule, for a conventional Chiller + CRAH cooling system, the largest single consumer of power in the facility, next to the IT load, will be the compressors. Progressive evolution of data centre design over the last few years has sought to lower the cooling energy, which has invariably meant reducing the dependency upon mechanical cooling, or in the case of Direct Air Optimisation + Evaporative Cooling, eliminating compressor power altogether.

The Green Grid has developed metrics to assess the relative effectiveness of the data centre infrastructure. PUE – 'Power Utilisation Effectiveness', is a dimensionless parameter, that shows energy use efficiency, where:

$$PUE = \frac{\text{Total Energy used by the Facility}}{\text{Total Energy used by the IT Equipment}}$$

and the partial PUE associated with the cooling system energy can be expressed as:

$$pPUE_{cooling} = \frac{\text{Total IT Equipment Energy, plus Cooling Infrastructure Energy}}{\text{Total IT Equipment Energy}}$$

The ideal value of PUE is 1.0, implying that all the energy used on the site goes to the IT Equipment. There is no theoretical upper limit for PUE.

The other commonly used metric is WUE – 'Water Utilization Effectiveness'. This is the expression of the site's water consumption relative to the IT Load. As such, it has the dimensions of litres/kWh, where:

$$WUE = \frac{\text{Annual Site Water Usage of the Facility}}{\text{Energy used by the IT Equipment}} \text{ l/kWh}$$

Setting aside debate focused on 'WUE source', which deals with the off-site consumption of water to generate the power that is supplied to the facility, WUE has taken on greater importance in recent years, as the growth in data centre capacity, coupled with greater usage of evaporative-cooling based systems. The theoretical ideal value for WUE would be zero, i.e. no water usage at all. However, in practice, this may not be a good outcome for the total operating cost of the facility, as there is often a trade-off between using water and consuming more electricity. As can be seen from the preceding sections, evaporative cooling plays an important role in improving the overall energy efficiency of the cooling system. In most localities, water is a cheaper and more plentiful commodity the power, so using water to reduce power consumption and thereby lower operating costs, makes perfect sense. But this also comes at a cost, as explained in 1.7.4. The balance between water cost and availability and the cost of electricity is also less clear-cut in some localities. Ironically, hot, arid regions where evaporative cooling would be at its most effective are also likely to be the regions where water is a scarce commodity and either not available in the quantity required, or quite expensive and carrying high embodied carbon in its production and delivery to the site (e.g. energy intensive desalination plants in Middle Eastern locations).

In this context, Liquid cooling offers an attractive alternative to air-cooling, as the higher water temperatures associated with the FWS Heat Rejection plant mean that it can still can operate effectively at relatively high ambient temperatures, only requiring evaporative or mechanical cooling trim under peak conditions. In more temperate zones, heat rejection can be achieved by a single stage heat rejection system. A dry or hybrid adiabatic cooler can provide the right temperature for the FWS without the need for interposed trim chillers and all the power usage. In this way, liquid cooling can achieve a significant improvement in PUE, without incurring any significant penalty in WUE, when compared with conventional air-cooled systems.

1.7.6 ITUE and TUE

It has long been recognised that PUE is a relatively crude metric, as it considers only the power delivered to the rack, and not the use to which this power is put. The Green Grid therefore recognises IT Power Usage Effectiveness (ITUE) as an alternative power usage indicator. The purpose of ITUE is to give a power usage efficiency indication at server level. It is defined as follows:

$$ITUE = \frac{\text{Total Energy into the IT Equipment}}{\text{Total Energy into the Electronic Components}}$$

ITUE is a server only calculated value. No other power demands, external to the server, are considered. Therefore, it calculates a server specific efficiency.

A second indicator, Total-Power Usage Effectiveness (TUE), defines the relationship between PUE and ITUE and provides a more precise indication of the data centre overall energy performance.

$$TUE = ITUE * PUE$$

As can be seen from the equation, TUE is the product of the multiplication of ITUE, a server specific value, and the PUE, a data centre overall value.

(Source: *PUE in the past, Towards TUE in the future*; Henry Coles Lawrence Berkeley National Laboratory Environmental Energy Technologies Division (EETD) High Technology and Industrial Systems Group; BrightTalk: March 19, 2014).

TUE highlights an important attribute of liquid cooling, **which is the comparatively small proportion of server power that is consumed by parasitic loads**. The dielectric fluid pumps within the liquid-cooled server, and PSU cooling fans within an air-cooled server, are both powered at rack level and are included within the Total Energy of the IT Equipment, for the purposes of calculating PUE. However, the dielectric fluid pumps in an Iceotope KU:L Box only account for approximately 3% of the total rack power, whereas the power consumption of the PSU cooling fans is generally around 10%. The difference in TUE between typical air and liquid-cooled installations is approximately 7% of the Total ITE capacity of the facility.

1.7.7 Data Hall Architecture

- **Open Architecture** in Data Halls generally consists of channelling cooling air via the raised floor void and delivering it into the cold aisle via floor grilles located in front of the racks. The air passes through the servers into the hot aisle, absorbing heat from the equipment on the way. Hot air discharged from the rear of the racks can rise and circulate freely via natural buoyancy, and eventually finds its way back to the cooling units for re-circulation. Open Architecture is the most basic form of rack layout, and the most flexible. Racks of different heights and format can be grouped in the same aisle. Open Architecture does have the distinct draw-back of allowing re-circulation of warm air back into the cold aisle. This can lead to mixing and elevated server inlet temperatures, especially at the top of the rack. In extreme cases, this can cause overheating.
- **Cold Aisle Containment** is an evolution of Open Architecture. The supply air is channelled from the cooling units through raised floor void and delivered into the cold aisle by floor grilles. The cold aisle is enclosed by doors at each end and a roof that spans between the racks. Hot air discharged from the rear of the racks can rise and circulate freely back to the cooling units but is prevented from re-circulating back into the cold aisle by the containment system.
- **Hot Aisle Containment** reverses the cold aisle containment concept. Instead of containing the cooling air supply within the cold aisle, the exhausted hot air is contained within the hot aisle and conveyed back to the cooling units, or exhaust fans, via the ceiling void. Supply air can be delivered via a raised floor void, but this is more commonly omitted and a 'flooded room' approach adopted, where the entire data hall effectively becomes part of the cold aisle.
- **Airflow Management**. The drive to improve energy efficiency through the widespread adoption of containment system has highlighted the need for good airflow management at rack level. This generally involves fitting blanking plates across vacant U trays within the rack and adequate sealing of the containment systems, to prevent wasteful airflow leakage and to facilitate correct differential pressure control between the hot and cold aisles.

- **Liquid Cooled Architecture**. Immersion and Semi-immersion type liquid-cooled racks do not require any form of containment partitions and at first glance appear like Open Architecture. Racks can be disposed in rows, accessible at both front and rear ends. It is desirable to arrange pipework, fibre and power connections to access the servers from the front, while hydronic connections are accessible from the rear. There is no need for cooling air to pass through the racks. Nonetheless, some air circulation and cooling are required in the data hall to dissipate radiant and conductive heat gains from the racks, pipework, and cabling – also to provide a comfortable environment for the operatives. Facility Cooling Water connections can be connected to the rack manifolds from mains at high level or routed below the raised floor.

1.7.8 Building Format and Plant Space

- **Data Hall footprint as a function of Power Density**. Increasing power density from the typical air-cooled values by the adoption of liquid cooling results in the need for fewer racks and potentially reduced data hall footprint. Alternatively, the same data hall size can be used to accommodate rack of greater ITE capacity, providing an increase in overall capacity.
- **Plant Location**. An important consideration for many air-cooled data centres is the location of the cooling plant. This is particularly the case with large Direct or Indirect Air-Cooled hyperscale facilities. Air cooling plant, such as fan walls, externally mounted AHUs and IAC units must be considered in the overall design of the building and often occupy a significant proportion of the data centre footprint. Consideration must be given to intake and exhaust air paths, to avoid re-circulation of warm exhaust air and proximity to the technical space. The physical size of the plant must be considered, along with the need to minimise the length air distribution paths. Locating AHU and IAC plant at roof level is a good solution for single-storey builds but is less well suited to multi-storey applications.
- Air Cooled strategies involving chilled water-cooling terminal units and equipment (such as CRAHs, Fan Coil, etc.) installed in the data hall, or dedicated CRAH corridors, must be sized to support the entire cooling load, plus redundancy. The space dedicated to the cooling equipment within the data hall can be large and occupy a relevant portion of the data centre footprint. By comparison, liquid cooled strategies require minimal presence of room cooling emitters to cover the cabinet and cable radiant and convective heat gains, but the bulk of the heat load will be conveyed away from the racks by the FWS. The location of CDUs interposing between the FWS and TCS must also be considered.
- **Cooling Plant Footprint**. Typical liquid cooled medium temperatures are warmer than the temperatures encountered in air-cooled strategies. As such, adequate heat rejection can be achieved via the adoption of single stage heat rejection plant. Dry coolers or adiabatic coolers may handle the cooling load without the need of interposing chiller plant, leading to space savings for plant and ancillaries. The need for evaporative trim cooling is limited to peak summertime conditions only or eliminated altogether, meaning that water storage requirements are reduced or removed.

- **High Rise Data Centres**. Difficulties may arise with the adoption of Direct and Indirect Air-Cooling Strategies for high rise developments, due to the risk of exhaust air re-circulation. The extensive chiller plant often required with chilled water-based systems can also be difficult to accommodate, as the area of the roof is comparatively small, relative to the total area of technical space on multiple floors. The advantage that liquid cooling brings, of a simpler cooling system, coupled with more compact heat rejection plant described in the preceding paragraphs, is amplified on high rise and multi-storey buildings, where the available space is often restricted to the building roof. The ability of liquid cooling to operate far more effectively than air-cooling in a multi-storey configuration, is likely to become increasingly significant as large plots for hyperscale data centres become scarcer.

1.7.9 Environmental Considerations

- **Noise propagation – Internal**. Liquid cooled servers do not require any of the cooling fans normally associated with the PSUs and other components found within air-cooled servers and are therefore effectively silent in operation. Some plant noise may be present within the data hall, from CRAHs or AHUs providing background cooling to the room, and circulation pumps within CDUs. However, this results in much lower noise levels within the data hall than would typically be encountered within an air-cooled facility.
- **Noise propagation – External**. Noise levels produced by external heat rejection plant will be broadly similar for air and liquid-cooled data centres, but the omission, or reduction in the scale of chiller plant is likely to be beneficial. The absence of fan noise from extensive AHU and IAC unit installations normally be present for a Direct or Indirect air-cooled data centre would also be beneficial, in terms of external noise propagation.
- **Air Quality and Ambient Humidity**. With an air-cooled server, heat producing components are cooled by a steady flow of cooling air passing across the surface of the motherboard. Inherent in this arrangement is the potential exposure of sensitive electronic components to airborne contaminants. These falls into two categories; gaseous and particulate. The ASHRAE TC 9.9 Committee has produced a White Paper on this topic: *2011 Gaseous and Particulate Contamination for Data Centers*. This highlights that the two most common modes of IT equipment failure due to environmental contamination are copper creep corrosion on printed circuit boards and corrosion of silver terminations in miniature surface-mounted components. Sulphur-laden gases in the cooling air supply, especially where elevated humidity levels are also present, have been shown to cause corrosion to surface-mounted copper and silver pathways and connections. Particulate contamination, e.g. dust, can impair the cooling of electronic components and interfere with moving parts, such as cooling fans. Dust contaminated with ionic sulphur or chlorine-based salts that settles on electronic hardware can also cause corrosion.

Contamination can appear in the data hall from a variety of sources. In the case of gaseous contamination, this is invariably introduced via ventilation or direct air-cooling systems. In urban areas where a lot of airborne pollution may be present, gaseous filters can be provided on fresh air ventilation plant, but this is not a practical solution for large, hyperscale direct air cooling systems, due to

the cost of the filters and the energy penalty incurred due to additional fan power consumption.

Sources of particulate contamination range from dust and dirt brought in on the clothing of personnel, dust particles from cardboard packaging and other items which are occasionally (and incorrectly) allowed in, and residual dust and dirt from construction. No filter system is completely effective and none will be effective in dealing with dust brought in by other means, or dust that was already present within the space.

On the other hand, servers that are contained within sealed immersion, or semi-immersion liquid cooling enclosures are very well protected from both gaseous and particulate contamination and are largely immune to the problems that these may cause.

1.7.10 Continuous Cooling and Rapid Re-Start

As with air-cooled system, a utility power failure that interrupts the operation of the cooling infrastructure within a liquid-cooled data centre will place the IT equipment in jeopardy of overheating. For certified Tier IV 'Fault Tolerant' installations, the Uptime Institute requires the inclusion of 'Continuous Cooling' e.g. the uninterrupted supply of cooling medium to the rack, during a power failure. UTI also recommend Continuous Cooling where power densities exceed 4kW/rack, regardless of Tier. For air-cooled facilities housing high power density racks, this is an essential consideration, as the racks have no residual thermal capacity of their own and any interruption in the supply of cooling air will immediately lead to a dramatic increase in temperature, followed by equipment shut-down via internal over-heat protection devices. To highlight this risk, UTI has conducted experiments on a 6kW/rack average data hall, which demonstrated that intake air temperatures at the server typically reached the upper limit of the ASHRAE TC 9.9 'Recommended Range' thermal envelope within 60 seconds following the loss of cooling air movement.

For air-cooled installations, in their White Paper No. 809: '*Continuous Cooling is Required for Continuous Availability*', UTI has proposed three classes of Continuous Cooling:

- **Class A Uninterruptible Cooling:** Cooling fans and chilled water pumps supported by UPS power. Chilled water supply maintained via thermal storage tanks. Room conditions can be maintained within normal controlled range until main cooling plant can be re-started via generator-backed power supplies.
- **Class B Continuous Cooling:** Cooling fans are supported by UPS, but chilled water pumps are not, and thermal storage may not be present. Air circulation is maintained within the data hall, to mitigate temperature rise within the equipment.
- **Class C Interruptible Cooling:** No cooling equipment is supported by UPS. In the event of a power failure all cooling to the equipment ceases until the plant can be re-started via generator-backed power supplies.

With all classes of continuous cooling, consideration must be given to temperature rise within the room and cooling system following restoration of power and resumption of normal operation. If thermal storage tanks and distribution systems have been charged with warm water during the power failure event, then these will take time to be brought back down to their normal operating temperature. With Class B and Class C installations, depending upon

the duration of the cooling outage, the fabric of the data hall may well have warmed up and could take 30 minutes or longer to be brought back down to its normal temperature.

There is no specific set of classifications for continuous cooling that are applicable to liquid-cooled installations, and the UTIs current classification system is sufficiently broad to still be applicable, albeit with the omission of irrelevant components such as cooling fans. But several differences exist, that are worth highlighting:

- With full immersion and chassis-level precision immersion systems, the volume of dielectric fluid held within the enclosure provides an amount of thermal storage that is completely lacking in an air-cooled rack. However, one of the primary benefits of liquid cooling is its ability to support higher power densities. So although the dielectric fluid slows down temperature rise of the equipment, the rate at which energy is being injected into the enclosure may be very much greater and this needs to be considered when assessing the need for thermal storage in the FWS and UPS-support to the FWS and TCS circulating pumps.
- With cold-plate technology, the cold plate will have very limited thermal capacity on its own, and any interruption in the flow of cooling medium is likely to result in very rapid temperature rise. UPS-backed pumps would be an essential requirement.
- As the heat load is fully contained within the enclosure of a full and chassis-level precision immersion systems, there is no opportunity to take advantage of any thermal capacity that the room fabric might have, which would be the case with Class B air-cooled systems.
- The higher operating temperatures typically used in liquid-cooled systems reduces the reliance upon chiller plant. The dry air and hybrid-type coolers often used for heat rejection are faster to re-start on resumption of power supplies than chillers and could resume cooling almost immediately.

2.0

Iceotope Liquid Cooling Technology

2.0 Iceotope Liquid Cooling Technology

2.1 Company Profile



Key Facts

- Founded: October 2005
- Headquarters: Sheffield, UK
- Number of Employees: 50
- Core Patents: 25
- Championed by: Lenovo, Schneider Electric, Avnet, European Union
- Supported by: Ombu, Aster, Edwards Investments
- Industry Engagements: ASHRAE TC 9.9, Open Compute Project
- Awards: Deloitte UK Technology Fast 50 2019, Deloitte Technology Fast 500 EMEA 2019, Barclays Regional International Expansion Entrepreneur of the Year 2020, Sunday Times BGF 10 Green Tech to Watch 2020

History of Iceotope

- Iceotope was founded in 2005, with financial backing provided by Enterprise Ventures and South Yorkshire Investment Fund at its seed phase. 2010 saw early pilot deployments of its technology in Swiss banks.
- November 2011: Iceotope Raised £1m from investment partners of Guernsey.
- September 2012: Iceotope had its first, customer-owned, proof of concept installed at Leeds University.
- November 2013: Iceotope installed its liquid-cooled servers in the Poznan Supercomputing and Networking Center in Poland.
- December 2013: Iceotope raise \$10m of investments in a Series A round led by Aster Capital, the venture capital firm.
- March 2014: Iceotope announce the launch of PetaGen, a liquid cooling system, designed for High-Performance Computing and suitable for Supercomputing data centre facilities.
- August 2014: PetaGen is deployed to University of Huddersfield 3m Buckley Innovation Centre and Rolls Royce.

- November 2014: Iceotope announce new products aimed at HPC workloads with Lenovo at SP14.
- July 2015: Iceotope develops the first fan-less Mellanox-based InfiniBand and Ethernet high-performance network and interconnect switch for HPC.
- December 2015: Iceotope is selected to work on the European Exascale System Interconnect and Storage (ExaNeSt) project funded by the EU Horizon 2020.
- Early 2016: First installs in ASEAN.
- Early 2017: First installs in USA.
- September 2017: Iceotope become the lead partner in the EuroEXA project.
- Late 2017: Global resellership with Schneider Electric.
- October 2019: Iceotope announce a partnership with Schneider Electric and Avnet.
- December 2019: The partnership announces the industry's first integrated rack with immersed, liquid-cooled IT.
- May 2020: Iceotope announce deployment of its chassis-level precision immersion liquid cooling technology at Eco Data Center's colocation facility in Falun, Sweden
- November 2020: Iceotope announce with Lenovo deployment of its chassis-level precision immersion liquid cooling technology featuring the ThinkSystem SR670.

2.2 Iceotope Technology

2.2.1 Enabling Architecture and Technologies

Iceotope's Liquid Enabled Architecture

Iceotope's cooling architecture provides a chassis level, partial fill precision immersion cooling solution, with precision delivery technology. This solution provides the server electronics industry a clear path to achieve industry leading metrics. This includes highest heat recovery, lowest PUE, and lowest server cooling power, all while maintaining high performance computing requirements.

Additionally, this architecture drives intrinsic value with nearly silent operation, high environmental resiliency, physical security, and improved operational reliability.

At its core, dielectric fluid is recirculated within a sealed enclosure through a heat exchanger, whereby the heat is rejected via a Facilities Water System.

- **KU:L Heat Sink** The KU:L heat sink is one of the foundational pieces to the Iceotope reference solutions. It is primarily used on CPUs and GPUs, providing localized 'bathtub' containment. It receives dielectric cooling fluid as part of the 'Precision Delivery' architecture, via an internal manifold and a nozzle directed at the centre of the chip. The KU:L sink fills up and spills over onto the main system board, which then washes over and cools other

system components and support chips as the fluid returns to the pump as part of the closed loop process.

The KU:L sink is easily adaptable to all CPU and GPU form factors, but can also be applied to other high powered system ASICs.

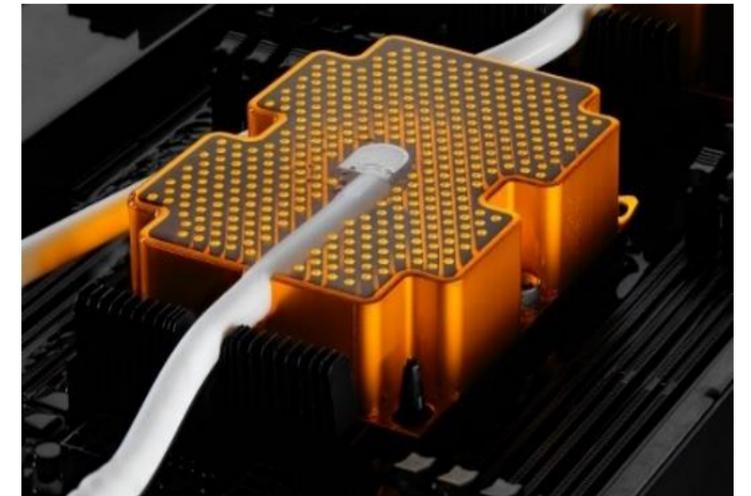


Figure 2.1: KU:L Heat Sink

- **KU:L Cold Plate:** The KU:L Cold Plate is another key piece to the Iceotope cooling technology portfolio. This cold plate is used to provide maximum cooling for HPC (High Performance Computing) applications, and can be used to compliment or in conjunction with KU:L system architecture.



Figure 2.2: KU:L Plate

2.2.2 Reference Architectures

KU:L System

Iceotope's KU:L system is the core reference platform, which targets users and applications where high volume reference form factors are used at scale. These users would typically be hyperscale or cloud providers.

This system provides the high-performance cooling needs for most industry applications and use cases, with all the aforementioned benefits.

As shown in Figures 2.3 and 2.4, the key enabling components include pumps to recirculate the dielectric fluid within the electronics system, in-system heat exchanger to reject the heat, internal manifold to deliver the cooling fluid to the key components, KU:L heat sinks to provide the localized cooling, and rack level manifold to reject the heat to facilities cooling.



Figure 2.3: 1U KU:L System Example

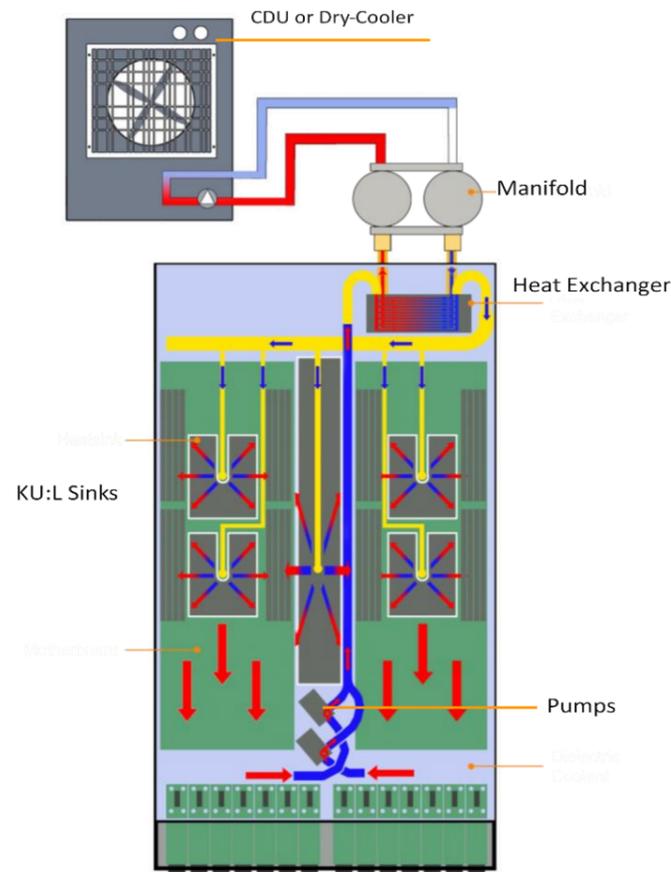


Figure 2.4: KU:L System Operating Layout

KU:L Hybrid System

The KU:L Hybrid system uses the core enabling components of the KU:L system, with the exception of replacing the KU:L heat sink with a cold plate for CPU cooling. This solution targets users and applications in the HPC (High Performance Compute) and AI (Artificial Intelligence) market segments, where the very highest cooling performance is required for the CPU.

This system solution provides a combination of primary, 'Direct to Chip' cooling to the CPU via cold plates, which removes 60-70% of system heat, while the dielectric coolant loop of the partial fill immersion system removes the rest. Figure 2.5: Typical Combined Immersion and Cold Plate technology rack layout shows an example of a 1U Hybrid Cooling solution.



Figure 2.5: Typical Combined Immersion and Cold Plate technology rack layout

Figure 2.6: Hybrid System Operating Layout shows the layout of the hybrid cooling solution, highlighting the two co-existing cooling loops. The key enabling components include pumps to recirculate the dielectric fluid within the electronics system, in-system heat exchanger to reject the heat, internal fluid manifolds for both the dielectric and water cooling loops, KU:L cold plates for CPU cooling, and rack level manifold to reject the heat to facilities cooling.

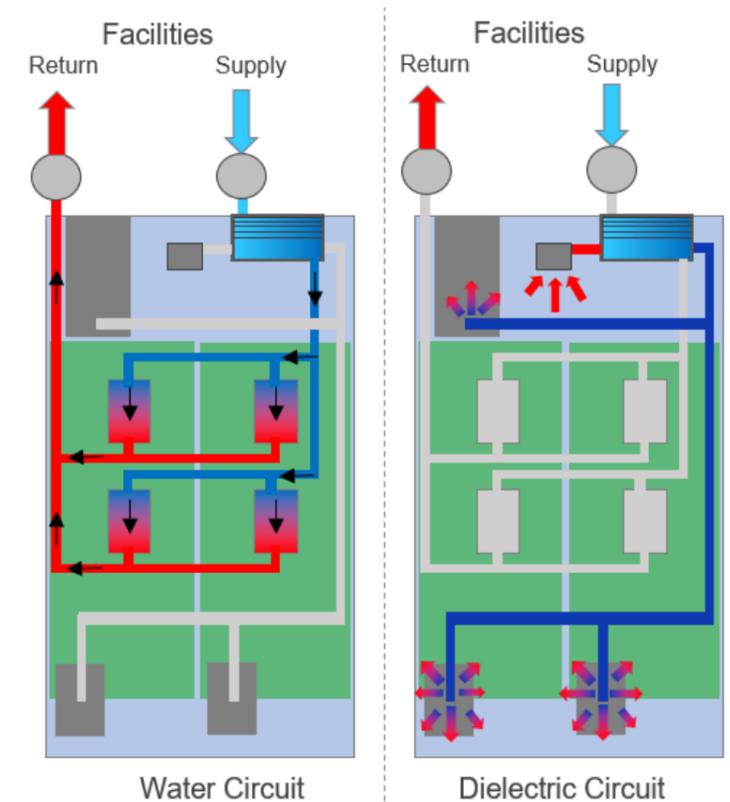


Figure 2.6: Hybrid System Operating Layout

Other Electronics Cooling Solutions

Dielectric immersion fluids offer the unique advantage of being able to cool the whole IT equipment stack, from servers, to switch, to storage boxes. Iceotope has demonstrated and deployed at-scale holistic solutions incorporating these other critical IT components.

2.2.3 Reference Solutions

Hyperscale Cloud KU:L Box

In conjunction with Schneider Electric, Iceotope has developed a prototype rack that utilises existing server mounting locations and adheres to standard equipment designs, as far as possible.

The rack houses up to eight ‘Cubbies’ that can host liquid-cooled, KU:L Box chassis-type server enclosures, plus a liquid-cooled power supply unit.



Figure 2.7: KU:L box chassis exterior view

The Chassis enclosures comprise an insulated folded steel structure, featuring side-mounting rails, front pull handles to facilitate withdrawal, a tightly fitting, hinge-mounted lid, and status indicator LEDs.



Figure 2.8: KU:L box typical internal configuration, with PHX to rear of chassis

Cooling is achieved with internal dielectric fluid which is pumped throughout the box. This dielectric is cycled through plate heat exchangers at the rear of the box which are linked to a manifold at the back of the racks.

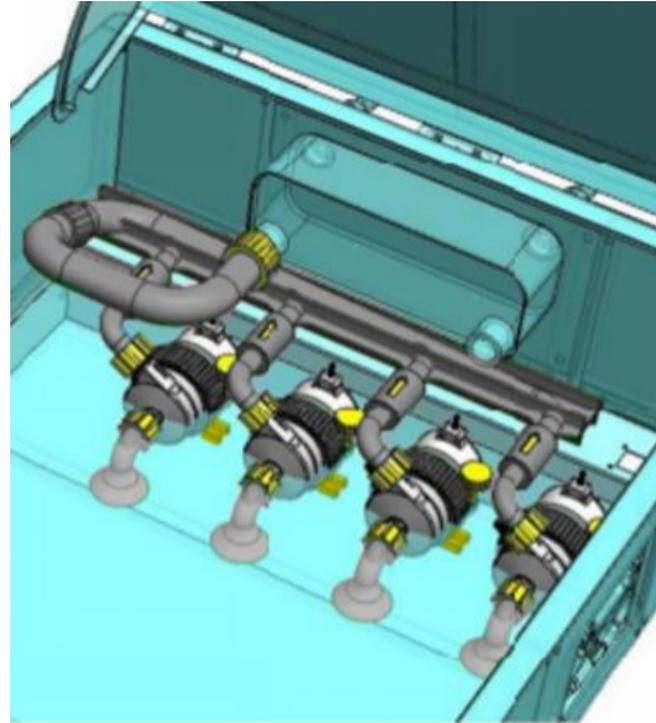


Figure 2.9: KU:L Box PHX and cooling fluid pumps

Racks are made up of individual KUL box chassis systems. Each box slides into the rack frame. Together several boxes make an entire rack. Different box configurations can be used to alter the operation, performance, and power of the rack. A standard rack layout is illustrated in Figures 2.10 and 2.11.

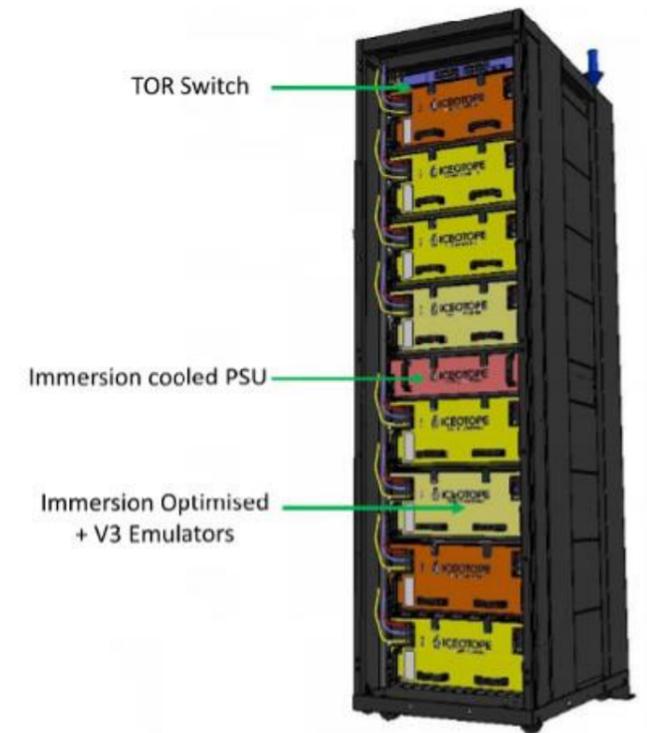


Figure 2.10: Iceotope Rack Layout

The arrangement is configured to allow individual KU:L box chassis to be slid out of the rack for access, without the need to manually isolate or drain the connection to the cooling water system and without disrupting the flow of cooling water to adjacent KU:L boxes within the rack.



Figure 2.11: Rack with KUL box extracted

Cooling water to the KU:L box chassis is supplied via a pair of manifolds that are built in to the rear of the rack. Insulated flexible hoses provide the final connection to each Cubby, terminating with self-sealing blind mate connectors that engage with the bayonet connections on the KU:L box chassis, when these are slid into position within the rack. The arrangement is configured to allow individual KU:L box chassis (Cubby) to be slid out of the rack for access, without the need to manually isolate or drain the connection to the cooling water system and without disrupting the flow of cooling water to adjacent KU:L boxes within the rack.

Cooling water from the manifolds is supplied to built-in Plate Heat Exchangers (PHX) within each of the KU:L boxes. The purpose of the PHX is to provide primary cooling to the internal immersion cooling system, whilst also providing hydraulic separation.

Each KU:L box is partially filled with a dielectric coolant, to a depth of around 30-50mm (varies according to internal IT equipment deployment). Warm dielectric fluid is drawn up from the bottom of box by miniature pumps and is circulated and cooled through the internal PHX before being delivered to heat producing IT components, which are partially, or fully immersed to maximise heat transfer.

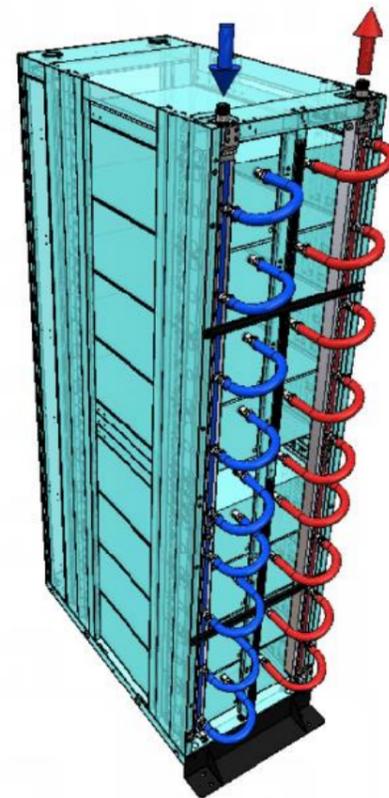


Figure 2.12: Rack Manifold layout

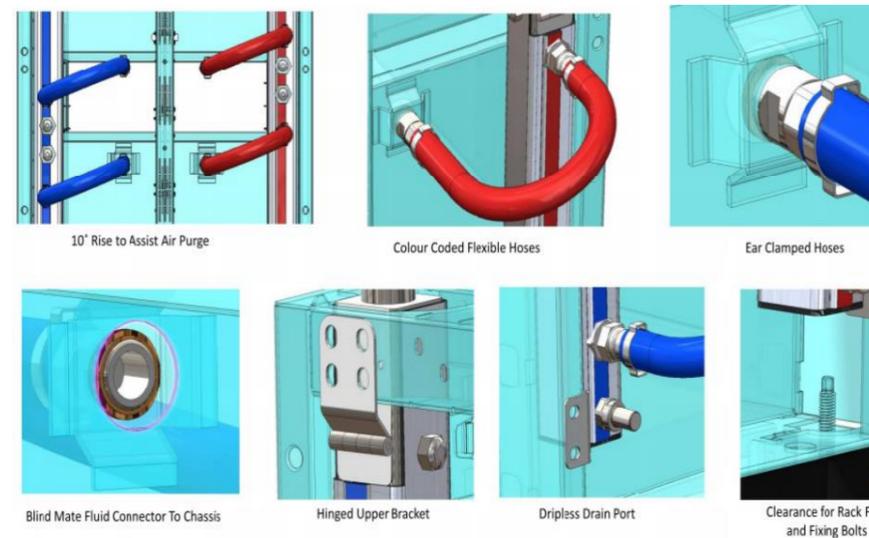


Figure 2.13: Manifold and rack connections

Each KU:L box has only a single PHX and single pair of connections to the cooling water manifolds, which are potential single points of failure. However, the type of PHX proposed by Iceotope are of stainless steel, fully brazed construction and are rated for much greater pressures than they would be subjected to under normal operating conditions, so the risk of failure is low. The dielectric fluid pumps are mounted in parallel on a common discharge manifold and provide N+1 redundancy. The pumps are designed to be 'hot swappable' without the need for tools. Pump status is monitored via the Chassis BIOS.

The internal dielectric fluid target mean temperature is 50°C, with the return temperature onto the PHX in the range of 55-57°C and a delta of 10K. The optimum operating temperatures for the facility cooling water system is 40°C flow and 46.5°C return, with a maximum water leaving temperature of <50°C. This falls within ASHRAE Liquid Cooling Class W5. Increased operating temperatures result in increased dielectric pump speed. As pump life is related to rotational speed, maintaining the leaving water temperatures at no greater than, say 47°C is beneficial.

This technology permits an increased IT load within each rack, beyond what is feasible with air-cooled racks. Iceotope reports that the KU:L BOX can host up to 60 kW/rack IT load.

Individual KU:L boxes can be configured to support three different load steps; 2kW, 4kW and 6kW, to address different IT configurations. For simplicity, Iceotope proposes to supply the same size PHX to all models of KU:L box and then compensate for the difference in cooling load by providing appropriately sized flow restrictors within the FWS connections to the PHX. These flow restrictors also provide terminal resistance at the point of connection to the manifold, which is intended to ensure that the manifolds are hydraulically self-balancing, regardless of the capacity and order in which the KU:L boxes are deployed within the rack. Connection to the FWS is via a differential pressure control valve, which maintains a constant 50kPa across the flow and return manifolds. This maintains a self-balancing arrangement.

KU:L 2

KU:L 2 is a flexible and server adaptable solution, where current 1U and 2U air-cooled servers can be converted to liquid cooling by a simple swap of air cooling infrastructure with Iceotope's liquid cooling technology.

This solution targets the emerging need for enterprise compute at edge locations, providing high performance cooling with environmental resiliency. Figure 2.14 shows an industry standard 2U air-cooled server, converted from air cooling to liquid cooling and integrated into the Iceotope KU:L 2 system. This example system has a maximum server compute power of ~2 kW with a cooling power of less than 1%.



Figure 2.14: KU:L 2 2U Air-Cooled Server Converted to Liquid Cooling.

Deployment Configuration #1 – Small Scale In-Rack HRU

This configuration involves a half rack supporting up to 4X KU:L 2 systems fully integrated with an in rack water to air HRU (Heat Rejection Unit). These systems are small and compact and therefore can be deployed anywhere. They are environmentally resilient, secure, and quiet, while enabling high-performance compute.

The target applications and users might include industrial factories, smart cities, utilities, oil and gas, healthcare, military, and mobile edge.



Figure 2.15: Deployment configuration #1 layout

Deployment Configuration #2 – In-Row Water to Air HRU

This configuration involves full height racks with up to 16X systems in each rack. These racks are paired with an in-row water to air HRU.

The target applications and users for this configuration would be small enterprise and colocation facilities, where water cooling is not readily

available, but requires the high performance and resilient cooling offered by KU:L 2.



Figure 2.16: Deployment configuration #2 layout

Deployment Configuration #3 – Edge Container with HRU

This configuration is targeted to be used at a remote edge location, where edge shelters, pods, or containers would be used to house the rack. These containment shelters would rely on a larger water to air HRU and could require auxiliary DX cooling to assist in extreme temperature environments.

The value of this configuration is that it brings remote monitoring, HPC/AI enabled at edge, it has physical security, ease of service and environmental resilience.

The target users are mobile/telco, communication/audio/visual and transportation.



Figure 2.17: deployment configuration #3 layout

Deployment Configuration #4 – Large Scale In-Row CDU

Configuration 4 has enterprise high performance compute scale, where several racks are cooled with an in-row CDU. Each rack can have up to 16X KU:L 2 systems and are fed by an in-row water to water CDU.

The major benefit here is high-density at scale, while being able to achieve industry leading heat recovery and energy reuse.

Target users and applications might include universities, oil and gas, government labs, scientific and research institutions.



Figure 2.18: Deployment configuration #4 layout

2.3 Comparisons with other Forms of Liquid Cooling

2.3.1 Full Immersion Tanks

“Open bath” refers to the “open” liquid-air interface and thus surface tension between the liquid and the air being a distinctive element. Open bath systems are usually tanks which contain a larger body of dielectric liquid where electronics are immersed into the bath. Multiple electronic assemblies are sharing the same liquid. This liquid can be based on single- or two-phase technology. Regardless of the term, open bath systems can be fully sealed, but are always opened from the top to service IT equipment. The water circuit for open bath immersion systems is either connected to a CDU which circulates the dielectric liquid, or to an integrated heat exchanging device which is part of the tank. For a facility interface, CDU's are usually designed for 100 kW or more, whereas an integrated heat exchanging device is usually designed for 10-100 kW cooling capacity.

As the name suggests, a feature of full immersion tanks is that the electronic equipment is fully submerged within the dielectric fluid, which is then circulated around the tank, rather than specifically across the electronic components. Whilst this ensured good heat transfer from electronic components, it requires a significantly greater volume of dielectric fluid to be provided than is the case with the Iceotope chassis approach, where fluid is continuously circulated over the components, from a shallow sump in the base of the cubby.

This difference in approach offers several distinct advantages over full immersion tanks:

- Far less dielectric fluid is required for a given cooling capacity. This has cost and environmental benefits.
- The reduction in fluid volume equates to a significant weight reduction, which for a large deployment would equate to an advantageous reduction in structural load.
- The KU:L Box cubbies each form a self-contained unit that is interchangeable with other cubbies within the rack. This is particularly advantageous for remote Edge-type deployments, where on-site server build facilities may be minimal or non-existent, as it provides ‘Plug and Play’ functionality. In the event of a server failure, a replacement Cubby can be shipped to site, pre-built, tested and already filled with dielectric fluid, where it can be quickly exchanged for the failed unit.
- At all times, the replacement server would remain secure within its sealed Cubby and would not need to be exposed to any airborne contamination. This exchange process would not involve any disturbance to adjacent operation equipment within the rack. This would not be the case with full immersion tanks, where any failed server would need to be manually lifted out of the dielectric bath and then replaced with a new component. The replacement process would also likely be carried out alongside other live equipment held in the same tank.
- The ‘Plug and Play’ functionality of KU:L Box is also advantageous for provisioning new installations. The racks can be delivered to

site pre-populated with Cubbies, which would be tested at the factory and arrive fully charged with dielectric fluid.

2.3.2 Cold Plates / Direct-to-Chip

Liquid-cooled Cold Plates are often used as a supplementary means to cool very high capacity components within an air-cooled rack. The cooling fluid could be a dielectric fluid, or de-mineralised water. It is usually delivered to the Cold Plate by hoses extended from the CDU or a local heat exchanger.

This arrangement offers a good method for addressing exceptional heat loads within the rack, such as HPC or AI chip sets, which would otherwise be beyond the capabilities of the data hall air cooling system. This can be useful, as it allows for liquid cooling to be deployed to selected racks, where it is needed, whilst the rest of the equipment within the data hall is supported by the base cooling system. However, on larger scale deployments, the reliance upon two cooling systems to serve the same rack may not be so advantageous.

The Iceotope KU:L Hybrid System addresses this issue by providing by the addition of an integrated liquid-cooled cold plate into the chassis, which provides ‘Direct to Chip’ cooling to the CPU, whilst the remainder of the equipment is cooled by the dielectric coolant loop of the partial fill immersion system.

This solution offers the high cooling capacity of a cold plate, whilst removing the need for secondary air cooling to the remainder of the server and the additional infrastructure this would require.

2.4 Interface with Facility Water System

The format of the Iceotope rack supports conventional rack layouts within the data hall. The ‘hot’ and ‘cold’ aisles are replaced by ‘Access’ and ‘Service’ designations for the front side withdrawal and access to the Cubbies and rear service connections to facility water and power supplies.

The cooling manifolds are formed from 50x50mm stainless steel square sections with 1” dia. BSP connections. The manifolds can potentially be top or bottom fed, but for the purposes of purging air from the manifolds and cubby internal heat exchangers, a top feed arrangement is preferable.

The design of the Facility Water System that serves the liquid cooled racks must take account of the following objectives:

- Maintain a constant flow of cooling water to each rack. Cooling water temperature can vary according to the season and ambient temperature, subject to limiting maximum temperature delivered to the rack.
- Be concurrently maintainable, providing dual path cooling water supply to each rack.
- Adapt to abrupt changes in rack flow rate, due to individual cubbies being withdrawn for access, without disrupting flow rate of facility water to other adjacent cubbies within the same rack.

- Allow for a phased deployed and connection of racks to the distribution, without impacting the operation of adjacent racks.
- Automatically adapt to changes in system hydronic resistance, due to racks being added or removed, or sections of the distribution being taken in or out of service.
- Minimise commissioning time and avoid the need for any interim re-commissioning arising from phased deployment.

Figure 2.19 shows a concurrently maintainable FWS system. Connection to each rack is via a differential pressure control valve + companion flow regulation valve, which together maintain the design flow rate to the Cubbies at constant 50kPa differential across the flow and return manifolds. The flow rate to each Cubby heat exchanger is regulated by a fixed orifice flow regulator in the connection to the cubby heat exchanger. These provide significant terminal resistance relative to the pressure drop of the manifold, which ensures a hydronically self-balancing arrangement, regardless of the capacity and order in which the Cubbies boxes are deployed within the rack.

If an individual Cubby is withdrawn from the rack for servicing, flow to its heat exchanger is immediately shut-off via self-sealing couplings in the connections on both the manifold and the heat exchanger. This has the effect of increasing the pressure drop across the rack to more than 50kPa, causing the Differential Pressure valve to close and thereby throttle the flow.

To ensure trouble free operation of the self-sealing couplings, it is essential that the FWS is provided with filtration to limit suspended particle size to no greater than 50µ. Failure to maintain FWS water quality may result in dirt clogging the heat exchangers or couplings. This could affect cooling capacity of the heat exchangers, or prevent the couplings from sealing correctly, leading to water leaks.

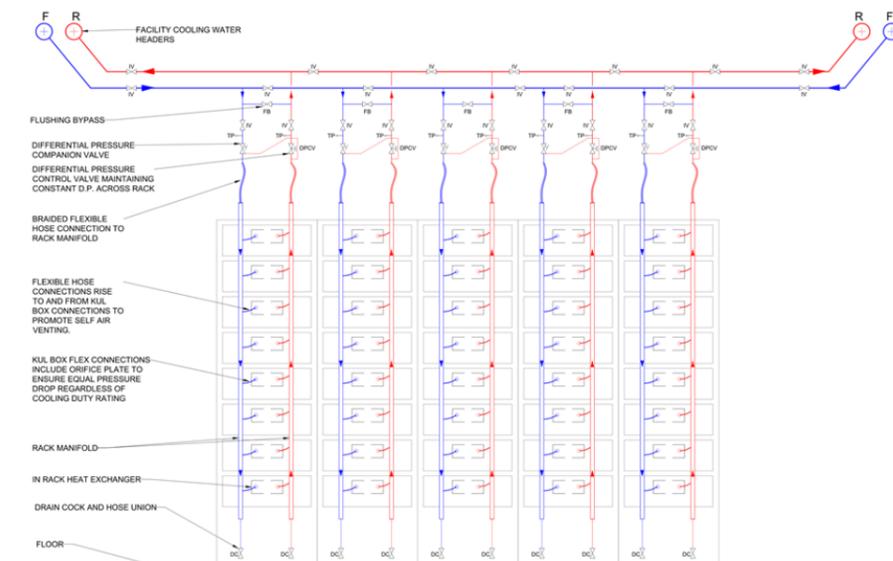


Figure 2.19: FWS Connection to Iceotope Racks

2.5 Dielectric Fluid Selection

To a large extent, the Iceotope KU:L system is 'fluid agnostic'. This enables the end user to decide which coolant best fits their use case and requirements. The design of the KU:L box chassis, the rack architecture, and material set used in the construction of the cubbies are all broadly compatible with both oil and hydrocarbon fluids, and the portfolio of compatible fluids is being steadily expanded as new fluid types emerge.

Discussion on the attributes and characteristics of the two main fluid groups is contained in Section 3.4.

3.0

The Liquid Cooled Data Centre

3.0 The Liquid-Cooled Data Centre

3.1 Characteristics of the Liquid-Cooled Data Centre

3.1.1 Building Form and Layout

As with any other data centre, in addition to the technical areas, the design of a liquid-cooled data centres will need to accommodate supporting plant and systems, including electrical switch gear and UPS, mechanical plantrooms and heat rejection, emergency standby generators, fire protection systems and fuel and water storage. In addition, there may be administration offices and ancillary/support areas including, loading dock, de-boxing and storage areas.

Overall, the internal layout may not be dissimilar to that of an air-cooled facility, but will differ in several important respects:

- The rack layout within the data halls no longer needs to respect a consistent hot aisle/cold aisle arrangement, as cooling is not derived from airflow passing down the cold aisle. Containment systems are also redundant. The most efficient layout for racks will most likely remain continuous aisles, served by straight lengths of power busbar, containment trays and pipework. However, the need to consider airflow and the relationship between racks and cooling units/Fan Walls etc. is eliminated. This may allow greater flexibility in planning rack layout to allow for say, high powered liquid-cooled racks to be deployed in clusters, surrounded by low-power data storage racks.
- The need for ventilation and limited air cooling will remain, especially if the data hall is to contain a mix of air-cooled and liquid-cooled racks. However, this may be limited to the provision of room pressurisation and humidity control, and the need to offset convective heat gains from power cables, pipework and rack chassis. The requirement for extensive CRAH Corridors is also largely eliminated, allowing for more efficient use of space.
- The ability of liquid cooling to support higher power densities reduces the number of racks required to accommodate the total ITE load, leading to small data halls. If deployed at hyperscale, this could facilitate a substantial reduction in building footprint.

3.1.2 Heat rejection plant location and footprint.

Heat rejection plant footprint can be more compact, compared to air-cooled data centre space requirements, due to higher cooling water operating temperatures, greater system simplicity and need for less equipment. Mechanical plant is more flexible and easier to locate. Depending on the facility architecture, the heat rejection plant can be installed in an external compound, remote from the main building, or at roof level.

3.1.3 Summary of Recommended Design Parameters

ASHRAE Thermal Guidelines

ASHRAE's Thermal Guidelines for Data Processing Environments by Technical Committee 9.9 (TC9.9) provides guidance and recommended design parameters for environmental conditions within data processing environments. Although more generally aimed at air-cooled environments, with a particular focus on air temperature and humidity at the server inlet, its recommendations are equally relevant to liquid cooled environments, insofar that many of the constraints relating to avoidance of condensation and static build-up etc. are applicable.

ASHRAE Liquid Cooling Guidelines

ASHRAE's Liquid Cooling Guidelines for Datacom Environment Centres (2nd Edition) also published by TC9.9, provides specific recommendations for liquid cooled environments. Table 5.1 of the publication defines Liquid Cooling Classes W1 – W5, which set the Facility Water Supply temperature ranges applicable to different types of cooling system. These are summarized in Table 1.1. and reproduced below in Table 3.1.

Liquid Cooling Class	Supply Water Temperature to CDU [°C]	Heat Rejection System	Utilisation of Waste Heat Recovery
W1	2-17	Cooling Tower + Water Chiller	Not Provided
W2	2-27	Cooling Tower + Water Chiller	Not Provided
W3	2-32	Cooling Tower + Water Chiller, Hybrid Cooler	Not Provided
W4	2-45	Dry Cooler, Hybrid Cooler or Cooling Tower	Not Provided
W5	>46	Dry Cooler, Hybrid Cooler or Cooling Tower	Provided

Table 3.1: ASHRAE Liquid Cooling Classes
(Source: Liquid Cooling Guidelines for Datacom Equipment Centers, 2nd Edition ASHRAE 2014)

The above table also includes details of the heat rejection strategies applicable to each class and the potential for WHR.

Figure 3.1. below shows the typical infrastructure arrangements that would be recommended in order to achieve different levels of supply temperature to the CDU.

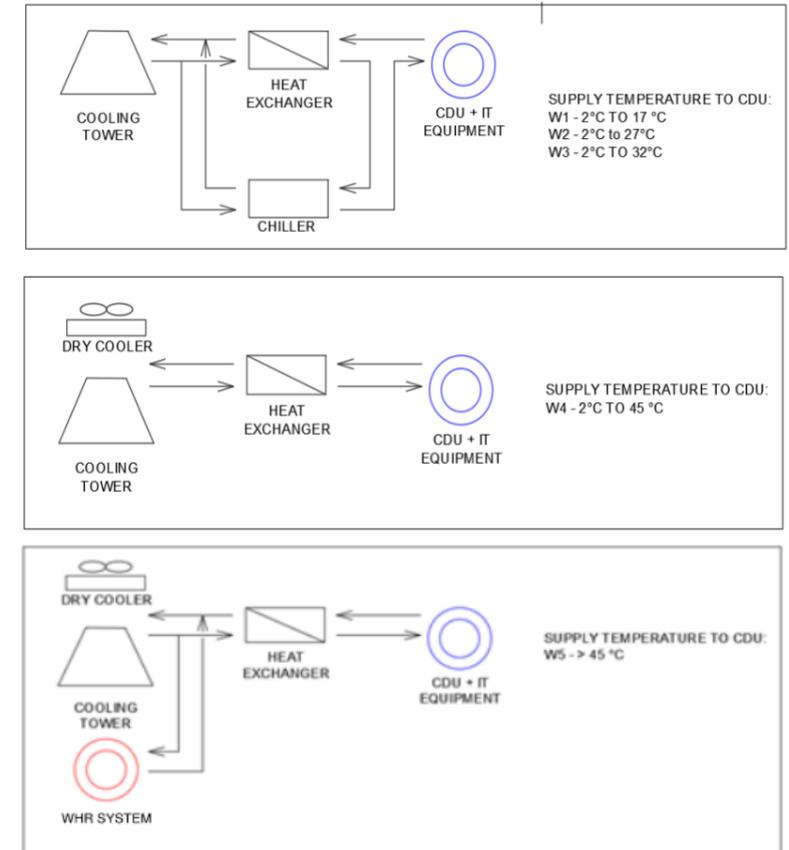


Figure 3.1: ASHRAE Liquid Cooling Classes, Typical Infrastructure Arrangements
(Source: TC9.9 Liquid Cooling Guidelines for Datacom Equipment Centers, 2nd Edition ASHRAE 2014)

An alternative arrangement for the connection of a possible Waste Heat Recovery circuit is illustrated in Figure 3.2. With this arrangement, the Waste Heat Recovery Heat Exchangers are connected in a side-stream. This allows for intermittent operation of the WHR circuit without diversion of flow within the FWS to the main heat rejection plant.

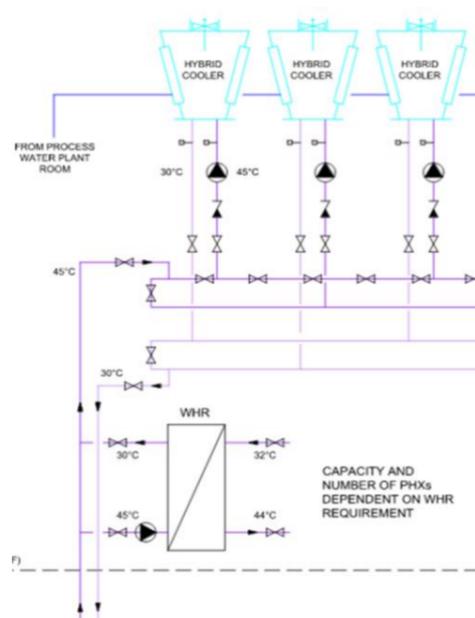


Figure 3.2: FWS Connection to Waste Heat Recovery Circuit

3.2 Design Considerations

3.2.1 Typical Temperatures, pressures, operating conditions.

As general rule, the higher the cooling water temperature, the better.

As summarized in the previous paragraphs, ASHRAE TC9.9 has defined liquid cooling classes W1 to W5, making a distinction based upon the cooling water temperature supplied to the Cooling Distribution Unit, where W1 indicates the coolest cooling water temperature (17°C) and W5 indicates the warmest cooling water temperature (>46°C).

It is advisable to keep the pressure of cooling water in the Technology Cooling and Facility Water Systems as low as possible, to reduce the risks of pipework failures and to mitigate the impact of any leaks, should these occur. ASHRAE TC 9.9. suggests locating Cooling Distribution Units (PHX & pumps package) near the data hall in order to reduce pumping pressure. CDUs can also act as a hydronic pressure break between the FWS, (which may be connected to roof-mounted plant and operate a high pressure, due to the static head on the system), and the TCS that serves liquid-cooled racks within the data hall.

Designs of liquid cooling classes W3, W4 or W5 permit warmer water supply temperatures within the TCS (up to 32°C for Class W3), which allows more efficient operation of heat rejection plant, facilitating water savings and waste heat recovery. Cooling water ΔT should be determined in consideration of the racks and heat rejection equipment limits, and preferably be set as wide as possible, to reduce pumping power consumption. Technology Cooling System water pressure from the CDU should ideally be in the order of 1 or 2 bar. On smaller distribution systems, sub-atmospheric pressure pumping may be considered.

3.2.2 Connection to Facility Power, IT and Cooling Infrastructure.

Fibre trays and power bus bars can be routed above liquid cooled racks, in the same way as in air-cooled facilities. With the Iceotope KU:L Box Rack, connections are to the rear of the rack, in what would be the hot aisle in an air-cooled data centre. Power and data connections and distribution within the rack follow OCP standards. The KU:L Box Cubbies and chassis are configured to permit tool-less access and servicing where possible.

Hydronic connections to liquid cooled racks are typically via reinforced flexible hoses, fitted with self-sealing connections at the rack. In the case of the Iceotope Rack, connection to the TCS is via manifolds to the rear of the cabinet. TCS pipework can be run at high level above the Service Aisle, or below the raised floor.

High level pipework, as shown in Figure 3.3, provides for a simple services distribution strategy within the data hall, with all services suspended at high level on a shared support frame suspended from the soffit or supported from slab level. No raised floor is needed and the racks can be installed directly on the floor slab, as is the case in many facilities that have hot aisle containment. High level pipework also facilitates air venting of the racks and TCS/Dielectric heat exchanger, via the distribution pipework. A risk with high-level pipework is that equipment and power connections are potentially exposed to any water that leaks from the pipework above.

Two strategies can be adopted to overcome this risk:

- Adopting low water pressure in the TCS circuit feeding the racks, in conjunction with the installation of a pipe in pipe containment, or drip trays. If leakage should occur, these would manifest themselves as drips or trickles of water that could be contained within the drip tray, rather than a high-pressure jet that could soak adjacent bus bars and equipment. Drip trays and containment systems should be fitted with leak detection systems to alert operating personnel to the presence of a problem.
- Adoption of sub-atmospheric pressure pumping within the TCS. These systems can exploit the hydrodynamic phenomenon known as Venturi Effect to create a sub-atmospheric pressure within the distribution pipework. Circulation pumps are arranged to create a liquid suction dynamic, rather than a liquid compression dynamic typical of pumped flows. In the event of pipework damage, the pressure of the surrounding air is greater than that of the liquid in the pipework, which impedes or prevents water spillage. These systems can be effective on smaller distribution circuits, but are not suitable for large distribution systems, as the negative circulation pressure needed to draw water around the entire distribution can cause cavitation at the pump.

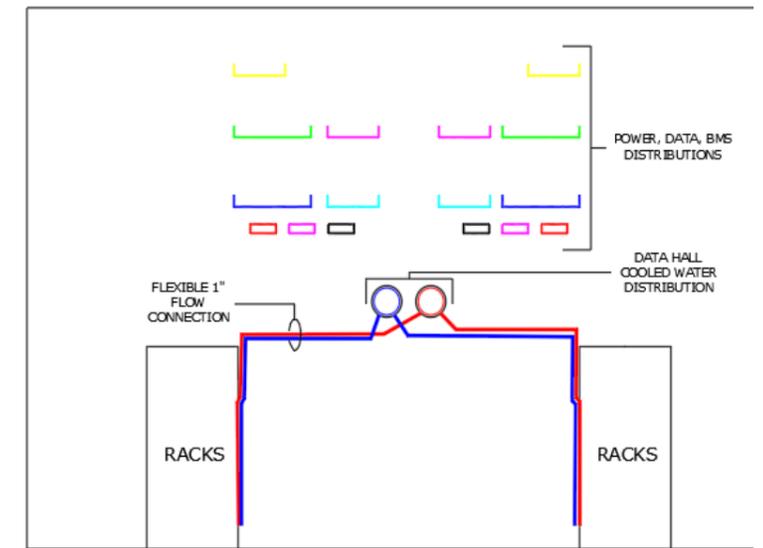


Figure 3.3: High Level Services Strategy

An alternative “Low Level Services Strategy is shown in Figure 3.4 below. This completely removes the risk of water leaking from above the racks.

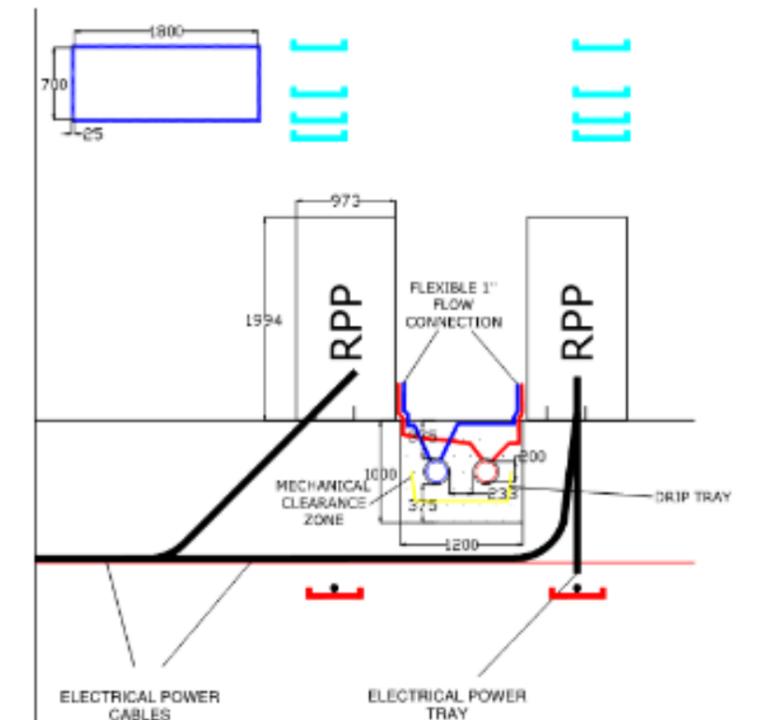


Figure 3.4: Low Level Services Strategy

3.3 Dielectric Fluids

A critical element of all liquid cooling systems is the dielectric fluid. To be effective, dielectric fluids must be electrically non-conductive, have good thermal capacity (at least for single phase applications), be

environmentally sustainable (i.e. have low or zero ODP and GWP, and be non-polluting) and be compatible with the electronic components that they will come into contact with. Dielectric Fluids should have low viscosity, present low risk of fire and be non-toxic to anyone who might come into physical contact with either the fluids themselves, or any fumes or vapours they might emit. Other desirable attributes are that they should be chemically stable, be relatively cheap to purchase and have a reasonably long service life; ideally at least equal to the expected lifespan of the IT equipment they will cool.

The dielectric fluids typically used for ITE cooling applications fall into two broad categories:

- Oils – both mineral and synthetic.
- Fluorocarbons – for both single and two-phase applications.

A third fluid that falls outside of these two main groups is water. This is invariably used within the main cooling/heat rejection Facility Water System (FWS) but may also be found in the Technology Cooling System (TCS) loop that serves Cold Plates in 'Direct to Chip' cooling applications. Water has the advantages of being cheap, abundantly available and possessed of good heat transfer capacity. However, it is not a dielectric, so cannot be used in any application where it is likely to come into direct contact with the IT equipment. Its presence in the data hall also warrants that very careful attention is paid to preventing leaks, or adequately containing them if they should occur.

A detailed assessment of the properties of all the fluids that are currently available on the market is outside the scope of this report. However, the attributes and considerations for the two main groups are summarized in the following sections.

3.3.1 Oils

General

Oils have for a long time been used as a dielectric coolant in power transformers and other electrical switch gear. By their nature, they are only suitable for single-phase applications, which includes the Iceotope KU:L Chassis system. They are frequently used in full immersive tank systems, which operate in a similar way to the cooling loop of an oil-cooled transformer, where the heat producing components are fully immersed in the fluid, which is circulated through a radiator, or heat exchanger, either via a pump, or by natural convection.

Both mineral and synthetic oils have been used for liquid cooling applications, although mineral oils have generally fallen out of favour, due to long-term incompatibility with various materials and components normally found in ITE equipment.

Dielectric oils are typically clear and non-viscous. They have a high boiling point and are non-evaporating. Whilst this attribute provides for a minimal/zero loss rate when the enclosure is opened for service and inspection, it also means that an oily residue remains on any surface or component that has been in contact with the fluid. Any spillage or drips that occur during server maintenance must therefore be cleaned with

an approved cleaning agent. Maintenance activities that require components to be withdrawn from the enclosure generally need to be carried out above the open chassis in the case of chassis-level precision immersion systems, or on a purpose-built tray, designed to catch any fluid run-off, and convey this back into the enclosure. Similar measures must be applied when transporting used components, which means they must be cleaned, dried and then placed in a bin, or sealed bag before dispatch. For this reason, if not carefully provisioned for, maintenance tasks can be messy and potentially unpleasant.

Fire Risk and Safety

Dielectric oils and the fumes they give off do not pose significant risk to human health, so long as reasonable precautions are taken during handling. Use of protective nitrile gloves is recommended to avoid skin contact, but minor splashes on the skin that are quickly wiped off will do no harm. Of greater risk is the potential for scalding, which could arise if IT technicians were to encounter fluids at their normal operating temperatures, which could be >50°C.

Dielectric oils of the type commonly used in ITE liquid cooling applications are typically classified as 'Combustible Liquids', rather than 'Flammable'. Mineral oils have a flash point of approximately 115°C, whereas for synthetic oils it is around 200°C. In both cases, the flash point temperature is well above the normal operating temperatures of immersive liquid cooling systems, although they will eventually burn if another source of combustion is available to first start and then propagate a fire.

Environmental Impact

Both synthetic and mineral oils have zero ODP and GWP. However, neither mineral nor synthetic oils evaporate to any meaningful extent – nor are they readily biodegradable in the environment. Hence, any spillage that gets into soil or water courses may potentially result in long term pollution.

3.3.2 Fluorocarbons

General

The development of fluorocarbon dielectric fluids has its origins in the manufacturing process of electronic components. There are four main types in use for liquid cooling applications:

- Perfluorocarbons (PFCs)
- Perfluoropolyethers (PFPEs)
- Hydrofluoroethers (HFEs)
- Fluoroketones (FKs)

These fluids all possess good dielectric properties, are chemically stable and have good heat transfer properties. They will evaporate and are available in a wide range of boiling points, ranging from 50°C up to around 270°C, making them suitable for both single and two-phase cooling applications. This characteristic means that the substance will evaporate readily off of IT components when these have been removed

from the immersion tanks or chassis tray for maintenance or replacement, which gives them a significant operational advantage over oils, which will always leave a messy residue. However, this feature can also lead to high loss rates if care is not taken to contain the fluid both during operation and when undertaking maintenance. Fluids formulated for two-phase applications would be particularly prone to this, and particularly when working with full immersion systems, it may be necessary to drain the fluid from the tank, prior to the commencement of maintenance.

Fire Risk and Safety

None of the Fluorocarbons used in liquid cooling applications have a flash point – nor are they flammable or combustible. They therefore pose no fire risk. They also do not pose any significant risk to human health from either mild inhalation or contact, but both should be avoided by operatives handling the substances. Any splashes on skin or surfaces would quickly evaporate but use of protective nitrile gloves is recommended. As with oils, handling hot, single phase fluids could result in scalding.

Environmental Impact

Fluorocarbons are non-polluting of soils, ground water and water courses, in that they will evaporate into the atmosphere. They have zero ODP, but all have a GWP to varying degrees. Of the four groups of chemicals above, PFCs and PFPEs have been around the longest, and have a GWP rating that is into the thousands. This means that they will contribute to atmospheric warming, if released into the atmosphere and their chemical stability means that they will persist for a long time, potentially decades. HFEs and FKs have been developed to have much lower GWP ratings. Fluorocarbon fluids liquids should therefore always be contained in gas-tight, or substantially gas-tight enclosures to ensure minimal substance loss to the atmosphere through evaporation. At end of life, Fluorocarbon fluids should also be reclaimed and recycled or disposed of in much the same way as F-Gas refrigerants. Local Environmental Codes in some regions may impose restrictions on the use of some substances that have a high GWP rating.

3.3.3 Oil and Fluorocarbons Fluids Compared

- **Total Cost of Ownership** – Oils have a lower CAPEX and OPEX cost than Fluorocarbon fluids, as they are cheaper to manufacture and do not suffer a loss rate due to evaporation. Whereas Fluorocarbons are significantly more expensive to produce and are subject to a loss rate due to evaporation over their service life, even when good containment measures are maintained. However, the service life of oils is generally 10-20 years, with modern synthetic oils at the higher end of the range, whereas fluorocarbon fluids can have a service life of up to 30 years, if properly maintained and filtered.
- **Component Compatibility** – Early types of mineral oil caused compatibility problems with PVC cable insulation (embrittlement

and cracking) and other components such as neoprene O-rings and sealing gromets. Paper ID labels and screen-printed serial numbers on individual electronic components were known to be stripped off, with the residue from the former occasionally being ingested into heat exchangers, where they can cause blockages. The introduction of synthetic oils has overcome these issues to a large extent, but material compatibility still needs to be checked. Fluorocarbons, on the other hand, generally display very good compatibility, particularly older PFC and PFPE types. More modern HFE and FK substances have been known to react with some component and insulating materials over prolonged exposure and the use of carbon filters is recommended by some manufacturers to address this.

- **Fire Risk and Safety** – Both oils and Fluorocarbons are generally safe to handle and occasional inhalation of fumes does not pose a risk to health. However, normal handling precautions should be observed and protective gloves worn by operatives. All oils are ultimately combustible, but do not represent a significant fire risk in the context of liquid cooling applications.
- **Environmental Impact** – Oils have zero GWP and ODP but are polluting in the environment. Fluorocarbons are non-polluting and have a zero ODP rating, but do have a GWP, which can be quite high in the case of older PFC and PFPE substances, which demands that measures are put in place to limit losses to atmosphere and may place them outside of local environmental codes in some regions.
- **Operation** – The nature of dielectric oils means that they can be messy and unpleasant to work with. Any spillage will cause needs to be cleaned up and special handling methods must be deployed, to contain oil spills during maintenance. Inevitably, unless operating procedures are observed and thorough cleaning is carried out following any maintenance, the floor and other surfaces within the data hall may become contaminated with oil deposits. Anecdotal reports refer to trails of oily footprints associated with full immersion systems. Fluorocarbons on the other hand are very clean to work with and equipment removed from immersion tanks will soon dry, due to evaporation. However, this imposes the need for vigilance when it comes to maintaining containment.

4.0

Performance Comparison

4.0 Performance Comparison

4.1 Objectives of the Comparison

4.1.1 Compare Footprints

The widespread adoption of Liquid Cooled Systems could lead to significant changes in data center design. Subject to how the IT equipment is deployed, the increased rack power density permitted by Liquid Cooling could allow for significant reduction in Data Hall footprint. In most air-cooled facilities, ITE rack power density is typically in the range of 7-10kW/rack, whereas the Iceotope KU:L Box chassis-level precision immersion cooling system can support up to 60 kW/rack. The impact this could have upon rack count and data hall footprint is illustrated below in Figure 4.1: Potential Reduction in Data Hall Footprint due to Increased Power Density. This shows the layout of 4.6MW ITE data hall containing 468 air-cooled racks, of approx. 10kW each. Superimposed upon this is the notional footprint of an equivalent liquid-cooled data hall, also of 4.6MW ITE, populated by just 78 liquid-cooled racks at 60kW/rack.

In practice, it is unlikely that such an arrangement would be feasible, as a real-world facility would need to house a mixture of different rack configurations, some of which might contain relatively low-powered data storage and networking equipment. However, it is undoubtedly the case that the adoption of liquid cooling could lead to improved utilisation of space and/or allow for the required ITE capacity to be housed within a smaller building. This could be attractive to a DC Developer seeking to maximise the potential of a confined site, or to fully utilise the available site power supply by increasing power density.

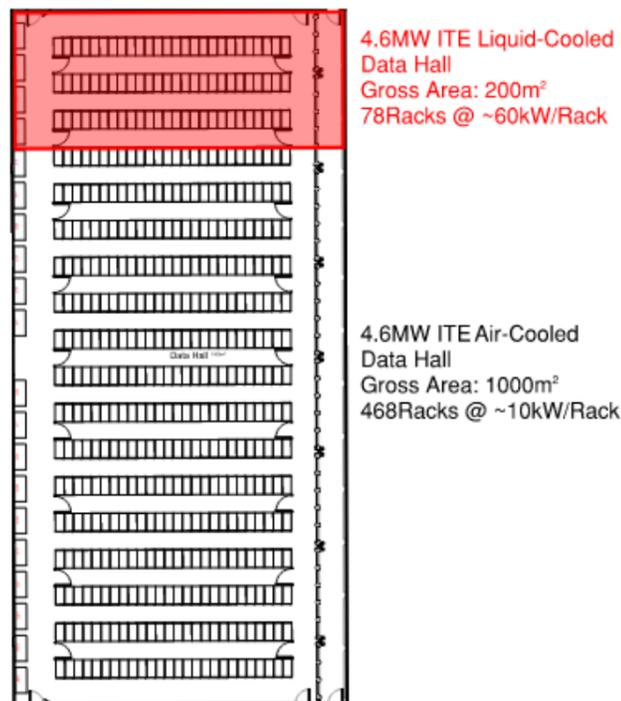


Figure 4.1: Potential Reduction in Data Hall Footprint due to Increased Power Density

4.1.2 Comparison of Annualised and Peak Energy and Water Consumption for Different Representative Locations

Energy and water requirements are strongly dependent to site specific climatic conditions. Nevertheless, the choices made at design stage will also impact the data centre power and process water demands.

Where liquid cooling systems can be adapted to the higher W4 and W5 range of cooling water temperatures described within the ASHRAE Liquid Cooling Classes (refer to Table 3.2), elevated FWS cooling water temperature can allow:

- Reduced dependency on compressors-based cooling (chillers), resulting in power savings.
- Reduced use of cooling towers in favour of adiabatic or dry cooling systems, resulting in reduced water usage and chemical water treatment.

Even in the presence of harsh site specific conditions, such as high ambient temperature and humidity encountered in tropical regions, worthwhile saving can be achieved.

4.1.3 Comparison of OPEX Cost

Using estimated annualised energy and water consumption for both a conventional air-cooled facility and that of a liquid cooled facility of equivalent capacity, it is possible to estimate annualised OPEX savings arising from the adoption of liquid cooling.

Although different locations have been assessed in the comparison study, for simplicity, common unit costs for power and water based upon prevailing UK prices have been assumed for all locations. These are shown on Table 4.1.

Summary of prices:

Utility Service	Price
Electricity	10.64p / kWh
Water	109.28p / m3

Table 4.1: Unit Costs for Power and Water

Sources:

Electricity: <https://www.gov.uk/government/collections/industrial-energy-prices#quarterly-energy-prices>

Water: <https://www.thameswater.co.uk/media-library/home/wholesale/our-charges/wholesale-tariff-document-2020-21.pdf>

4.2 Basis of Design

In order to effectively explore the various potential advantages of liquid cooling, three different scenarios are considered, each representing different possible deployments of Iceotope's Liquid Cooling Technology. Although the company names used in this study are entirely fictitious and locations have been changed to protect the identity of those involved, each scenario is based upon a real project, where Data Center Developers have considered the possible benefits of adopting liquid cooling within their facilities. It is important to note that their deliberations were not solely driven by considerations of OpEx and/or CapEx cost. Many developers are considering different ways of optimising their assets, or the potential of a site. For example, liquid cooling can sometimes offer a way of utilising trapped capacity within an existing data hall, or spare power that might be available to a campus or facility, which could not be fully exploited by retaining established cooling strategies.

- **Scenario #1** considers a Hyperscale Cloud Developer that is looking at ways of fully exploiting the power supply available to their site. Spatial constraints prevent the construction of an additional standard air-cooled data hall, as this would not fit on the plot. Boosting the capacity of their air-cooled data halls is not feasible, within the constraints of their current Plan of Record. How can they fully utilise the site supply, given the small amount of vacant space that remains on the plot?
 - **Primary Objective:** Fully utilise the capacity of the site's power supply within a new Data hall, that takes advantage of the limited space available on the plot. As an added benefit, consider the improvement to operating costs that may be obtained from the adoption of liquid cooling.
- **Scenario #2** considers a CoLo Owner that is considering how to utilise 'trapped capacity' within the existing sub-station of a high-rise CoLo Data Centre. Each of the preceding data halls have been built to a common standard and are now all fully let and occupied. However, much of the installed tenant's IT equipment is operating far below the envisaged demand. This has resulted in there being a considerable margin of surplus capacity within the

building's sub-station. Should the Owner now fit out their last data hall to the same specification as the previous halls, knowing that this will leave a lot of unused power in the sub-station, or try to do something different?

- **Primary Objective:** Utilise the trapped power available on site by increasing the Rack Power Density in the fifth and final hall of the development, with a partial adoption of liquid cooling.
- **Scenario #3** considers an Edge DC Operator based in Singapore. The local market is very price competitive and land costs are steep. How can the Owner both reduce operating cost, and exploit a smaller, cheaper site for its next build?
 - **Primary Objective:** Minimise the Total Cost of Cooling per KW of ITE Power and reduce the footprint required for the new facility.

Table 4.2 summarises the proposed ITE capacity and power density for each scenario, illustrating both the 'base' and alternative options that have been considered in the comparison.

Scenario	Cooling Strategy	Facility ITE Power [MW]	Data Halls [No.]	Data Hall ITE Power [MW]	Rack Power Density [kW/rack]
Scenario #1 – Hyperscale Cloud Developer (Richmond Virginia and San Jose California)	Base Option: Air Cooled	20MW	4	5MW	6kW/Rack
	Alternative: Air-Cooled + Liquid Cooled Suite	25MW	4+1	5MW	AC: 6kW/rack LC: 8.5 / 15 / 30kW/rack
Scenario #2 – Hi-Rise CoLo Operator (London, UK)	Base Option: Indirect Air Cooled	8MW	5	1.6MW (But four existing halls currently operating @ <1.2MW)	8kW/rack
	Alternative: Mixed Indirect Air Cooled + Liquid Cooled Racks	8.8MW	4+1	4 x 1.6MW 1 x 2.4MW	AC: 10kW/rack LC: 24kW/rack
Scenario #3 – Edge Data Center Operator (Singapore)	Base Option: Air Cooled	1MW	1	1MW	AC: 5kW/rack
	Alternative: Liquid Cooled	1MW	1	1MW	LC: 24kW/rack

Table 4.2: Basis of Design for Comparison Scenarios

4.3 Operating Thermal Envelope

For the purposes of establishing annualised energy consumption, PUE and WUE, it is important to consider the respective thermal operating envelope for each Scenario, as these will be different depending upon the market segment. Hyperscale and Cloud operators tend to follow an expanded thermal envelope to reduce operating costs. This will fall either within the TC9.9 'A1 Allowable' range, or more commonly be a bespoke envelope to suit the owner's own use case and equipment requirements. CoLo Operators however, that are hosting tenant's equipment, tend to operate their facilities within the tighter TC9.9 'Recommended' Range. The difference between the two operating envelopes is illustrated in Figures 4.2 and 4.3 and summarised below:

- **Hyperscale Operator (Air-Cooled)** – Assumed supply air condition at the server inlet of 18°C to 29°C Dry Bulb and 25% to 80% relative humidity. Supply air volume generally equates to a 12K rise across the racks.
- **CoLo Operator (Air Cooled)** – Assumed supply condition following TC9.9 Recommended Range. Supply air condition at the server inlet of 18°C to 27°C dry bulb and relative humidity range equating to -9°C dew point at the lower limit and 15°C dew point and 60% relative humidity at the upper limit.
- **Liquid Cooled or Mixed Air and Liquid Cooled Environments** – Taken as being the same as those specified for corresponding air cooled environments.

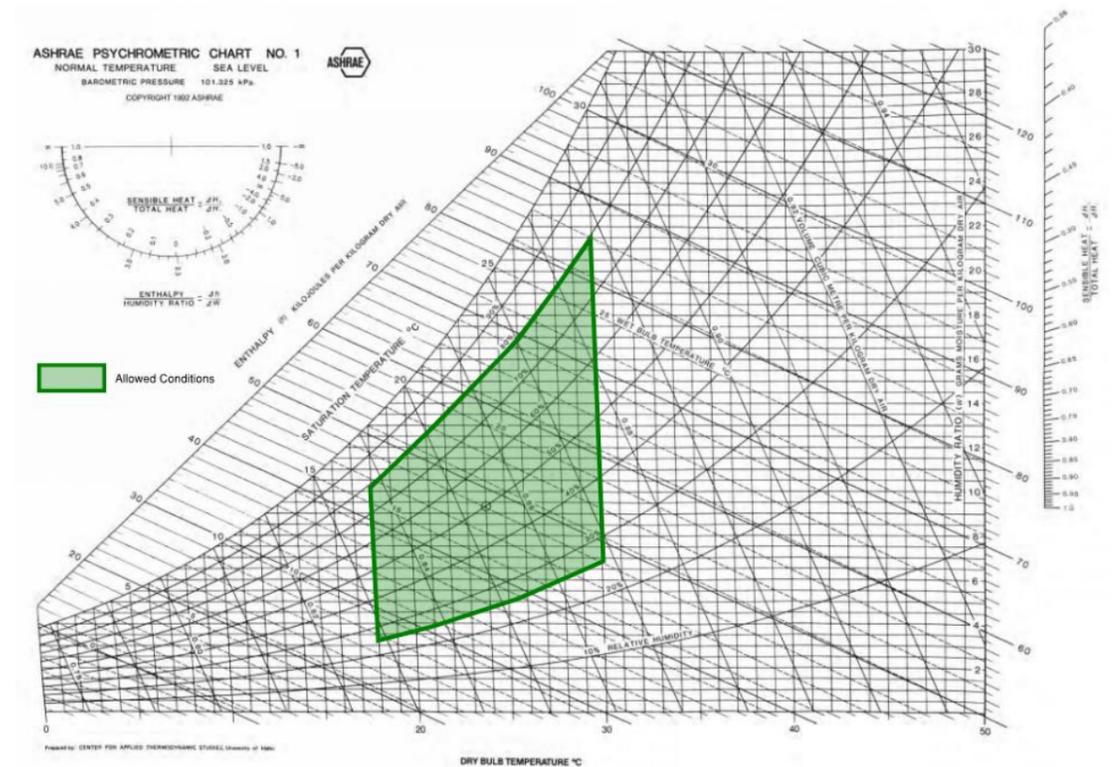


Figure 4.2: Thermal Operating Envelope – Hyperscale

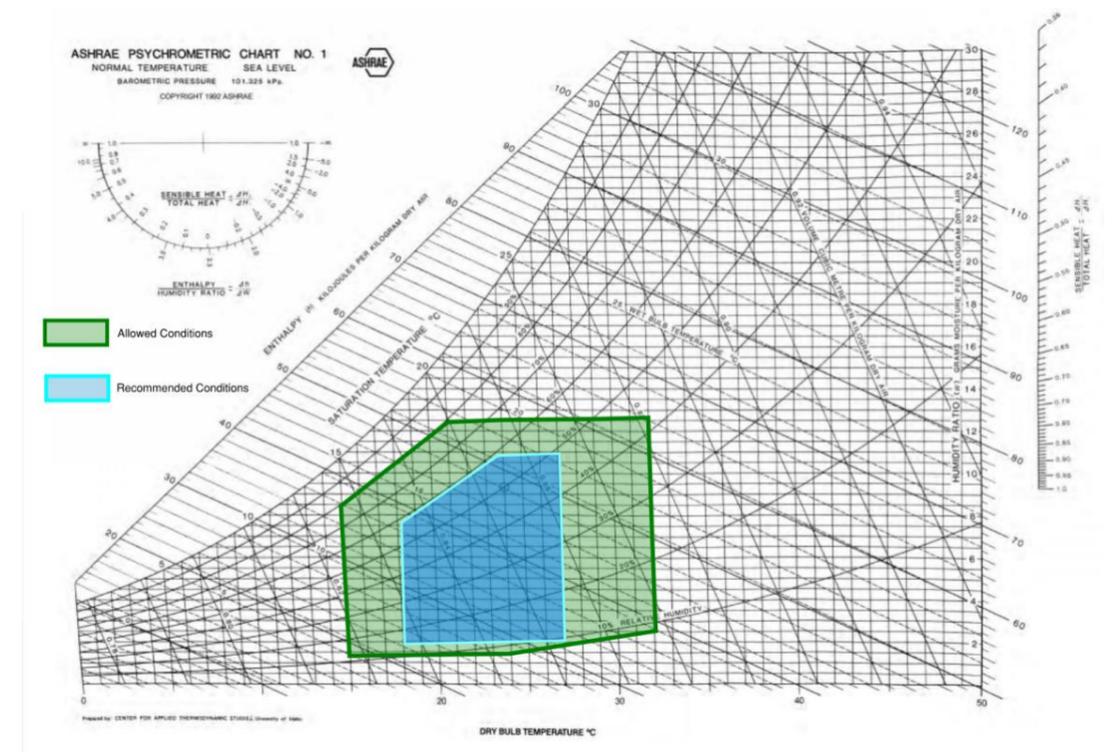


Figure 4.3: Thermal Operating Envelope - CoLo

4.4 Locations

Assumed locations for the three comparison studies are as follows

4.4.1 Eastern USA - Richmond, Virginia. Weather Data (Scenario #1a)

Richmond, VA has been considered as one of the locations for Scenario #1a. Richmond's climate is described as Cfa (Humid Subtropical) under the Koppen Geiger classification.

It presents no dry season with rainfalls homogeneously spread across the year, with cool winters and hot summers.

The considered weather reference is the 2017 ASHRAE bin file for Richmond International Airport, VA, USA.

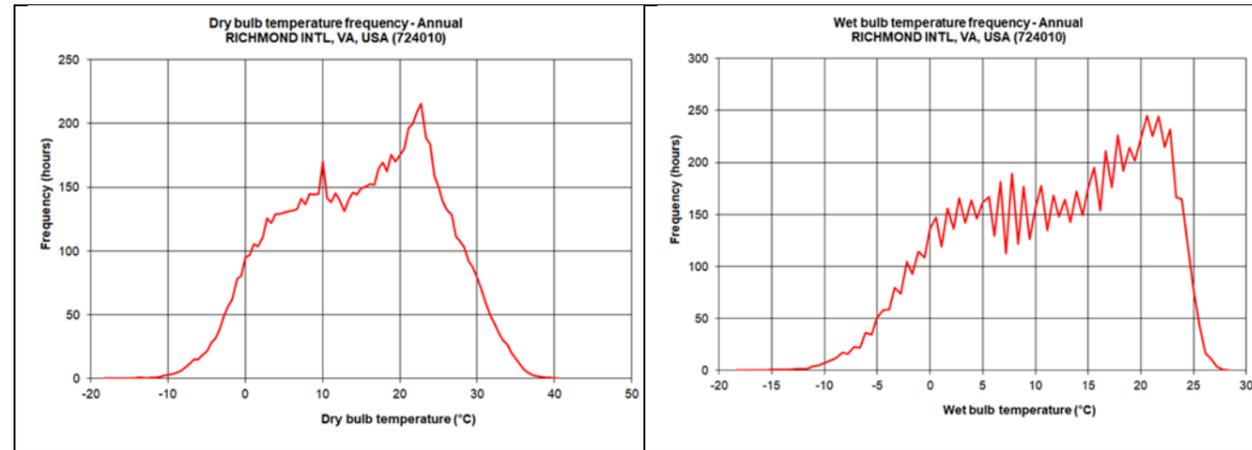


Figure 4.4: Seasonal Dry and Wet Bulb Temperatures for Richmond VA

The n=50-year Extreme Annual Design Conditions for Richmond are shown on Table 4.3.

Richmond International Airport, VA, USA. Lat: 37.505N, Long: 77.320W Elev: 50m		
Extreme Annual Design Conditions	N=50-year Return Period of Extreme Values	
	Min	Max
DB Temp	-20.0°C	41.8°C
WB Temp	-20.0°C	28.5°C

Table 4.3: Extreme Annual Design Conditions - Richmond, Virginia

4.4.2 Western USA - San Jose, California. Weather Data (Scenario #1b)

San Jose, CA has been considered as one of the locations for Scenario #1b. San Jose's climate is described as Csb (Temperate Mediterranean) under the Koppen Geiger classification.

It presents a dry season in summer with rainfalls in the wintery season, with cool winters and hot summers.

The considered weather reference is the 2017 ASHRAE bin file for Mineta San Jose International Airport, CA, USA.

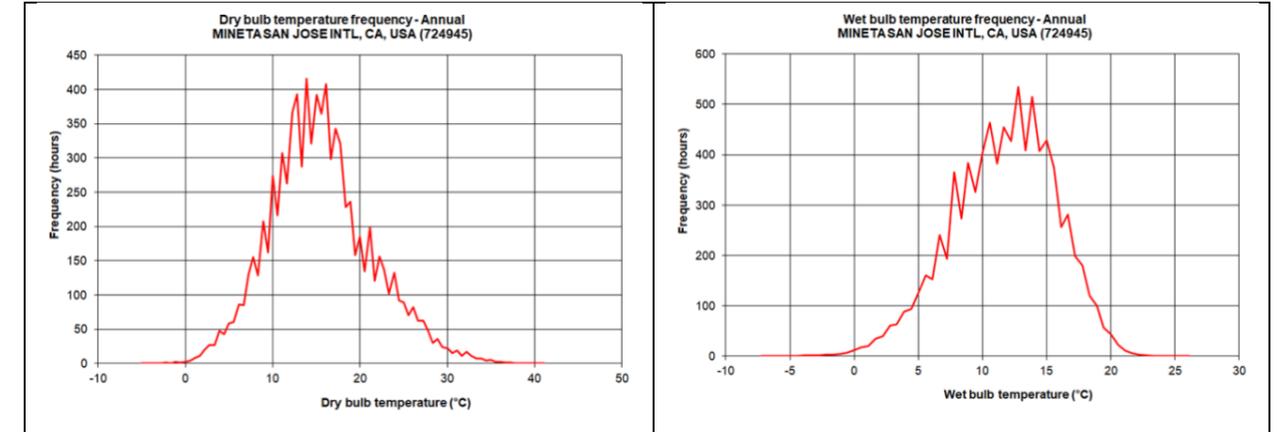


Figure 4.5:- Seasonal Dry and Wet Bulb Temperatures for San Jose, CAL.

The n=50-year Extreme Annual Design Conditions for San Jose are shown on Table 4.4.

Mineta San Jose International Airport, CA, USA. Lat: 37.359N, Long: 121.924W Elev: 51m		
Extreme Annual Design Conditions	N=50-year Return Period of Extreme Values	
	Min	Max
DB Temp	-5.4°C	41.7°C
WB Temp	-7.0°C	25.7°C

Table 4.4: Extreme Annual Design Conditions, San Jose, California

4.4.3 Northern Europe - London, UK. Weather Data (Scenario #2)

London UK has been considered as the location for Scenario #2. London's climate is described as Cfb (Oceanic) under the Koppen Geiger classification.

It presents no dry season, with rainfall homogeneously spread across the year, with cool winters and warm summers.

The considered weather reference is the 2017 ASHRAE bin file for London City Airport, UK.

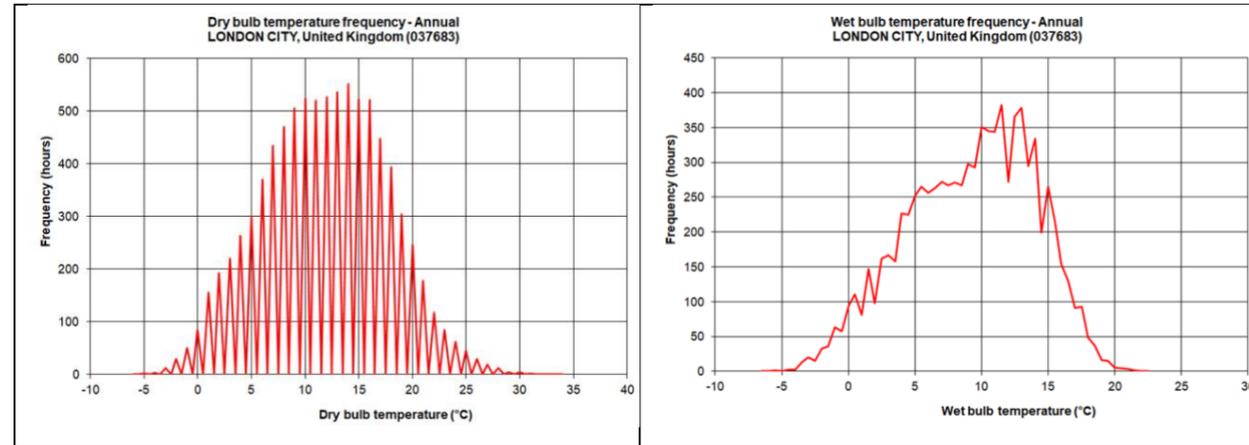


Figure 4.6: Dry and Wet Bulb Temperatures for London

The n=50-year Extreme Annual Design Conditions for London is shown on Table 4.5.

London City Airport, London, UK. Lat: 37.505N, Long: 77.320W Elev: 50m		
Extreme Annual Design Conditions	N=50-year Return Period of Extreme Values	
	Min	Max
DB Temp	-8.7°C	35.8°C
WB Temp	-8.7°C	24.1°C

Table 4.5: Extreme Annual Design Conditions - London, UK

4.4.4 Singapore - Weather Data (Scenario #3).

Singapore has been considered as the location for Scenario #3. Singapore's climate is described as Af (Tropical Rainforest) under the Koppen Geiger classification.

It presents consistent rainfalls homogeneously spread across the year, with hot winters and hot summers.

The considered weather reference is the 2017 ASHRAE bin file for Singapore Changi International Airport.

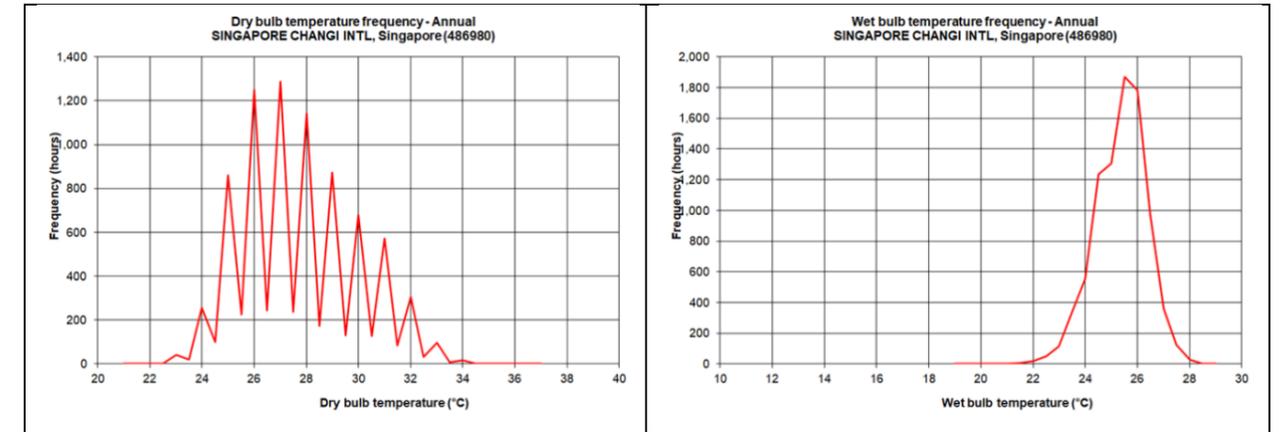


Figure 4.7: Dry and Wet Bulb Temperatures for Singapore

The n=50-year Extreme Annual Design Conditions for Singapore is shown on Table 4.6.

Singapore Changi International Airport, Singapore, Singapore. Lat: 1.350N, Long: 103.994E Elev: 7m		
Extreme Annual Design Conditions	N=50-year Return Period of Extreme Values	
	Min	Max
DB Temp	20.7°C	36.4°C
WB Temp	19.2°C	29.5°C

Table 4.6: Extreme Annual Design Conditions - Singapore

4.5 Energy Modelling

In all cases, the technology selected for the 'Base Case' air-cooled option is the best performing air-cooled technology currently available, taking account of the application and prevailing climate.

Wherever possible, annualised energy and water consumption figure for the main mechanical plant has been based upon actual plant selections obtained from vendors. Where this has not been possible, performance has been estimated based upon pro-rata the water and power consumption of different sized plant, performing under identical operating conditions to those described in each scenario.

Absorbed power for fans and pumps has been calculated based upon reasonable system resistances and estimated flow rates. Electrical losses for UPS, Transformers losses and switchgear have been assumed, based upon typical industry benchmark efficiencies.

In all cases, the same approach to calculating energy and water consumption has been adopted for both air and liquid-cooled options.

4.6 Notional Data Centre Comparison Scenario # 1 – The Hyper-Scale Developer

4.6.1 Introduction - James River Web Services



James River Web Services is a Hyper-scale Cloud provider with locations across the USA. To meet customer demand for increased capacity, they are planning to build two new Hyper-Scale Data Centres on Campus sites on the East and Western seaboard of the USA; the first in their home town of Richmond, Virginia and the second in San Jose, California. The initial Masterplan for each campus proposes two new-build data centre buildings, each housing four data halls of 5MW each.

JRWS's DCs follow a standard 'Plan of Record' design. Their preferred cooling solution features Fan Coil Walls in the Data Halls linked to external Hybrid Coolers

that can boost their heat rejection capacity during summer months by means of evaporative cooling pads on the inlet side of the heat rejection coils. In warm locations, they also add chiller trim to boost cooling capacity of the Hybrid Coolers during periods of hot weather.

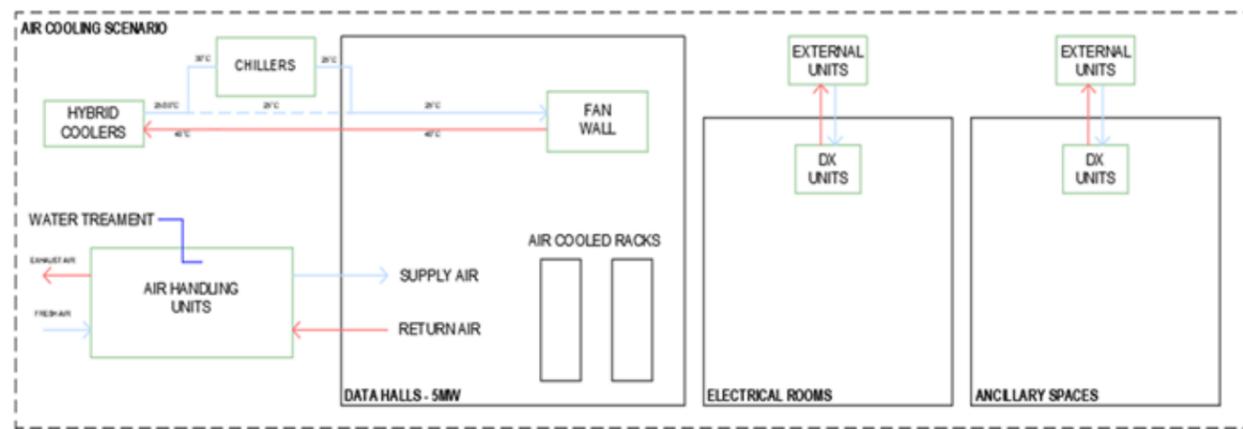


Figure 4.8: Scenario #1 - Base Option Schematic Arrangement

To maximise energy efficiency and the availability of free cooling, JRWS maintain thermal conditions in their Data Halls within Range of 18-29°C db, with a humidity range between -12°C DP and 8% rh, up to 17°C DP and 80% rh, as measured at the server inlet. Volumetric airflow rate equates to a temperature rise of 12K across the servers.

The ITE power density within each air-cooled Data Hall is 3kW per m²

The Fan Coil Wall system operates with chilled water at 25 - 40°C. For a proportion of the year, the external Hybrid Coolers can operate in 'Economiser' mode, with the evaporative pads being brought in to boost performance as the ambient temperature increases. Summertime conditions for both site locations are broadly similar, although San Jose has much milder winters. For both locations, it is apparent that to maintain acceptable supply temperatures during peak summertime weather, trim chillers are required to maintain a chilled water supply temperature of 25°C. The schematic arrangement of the system is illustrated in Figure 4.8.

Both campuses have abundant power supplies available to them, but neither can accommodate an additional data hall, based upon the footprint of the company's standard air-cooled, 'Plan of Record'. Hence, to fully utilising the additional available power, JRWS is proposing to add a 5MW high-density liquid-cooled suite to both buildings. This would be served by a dedicated Class W5 Technology Cooling System, operating at 40 - 46.5°C. An interposing Cooling Distribution Unit segregates the TCS from the Facility Water System, which operates at 38 - 45°C. Heat rejection via hybrid coolers. The relatively high operating temperature of the FWS means that trim chillers are not required for the FWS in either of the two locations and for much of the year, the Hybrid Coolers are able to operate in Economiser Mode. The schematic arrangement of the proposed liquid-cooled suite is illustrated in Figure 4.9.

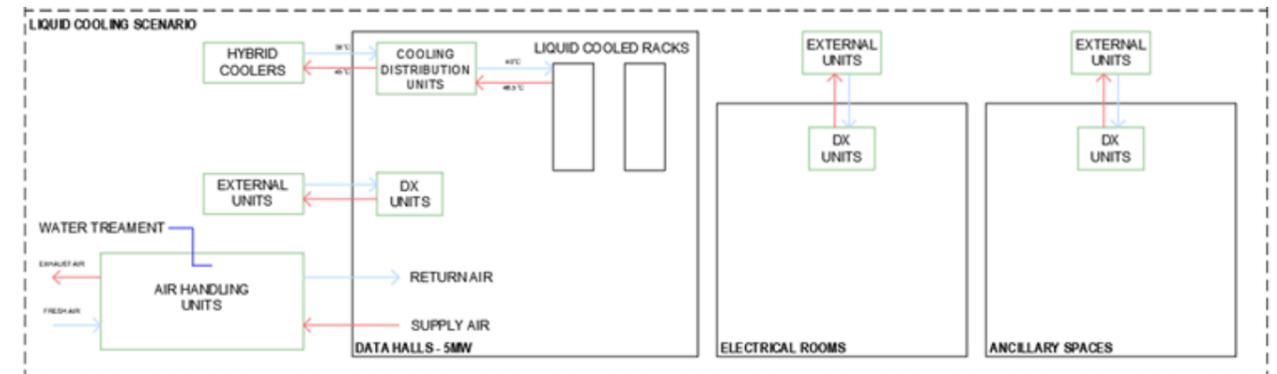


Figure 4.9:- Schematic Arrangement of Proposed Liquid-Cooled Suite

4.6.2 Hyper-Scale Performance Comparison

The objectives of this comparison scenario are to show how liquid cooling could be added to a proposed campus to allow the available site power supply to be fully utilised, when to do so using air cooling alone would either mean departing from an established 'Plan of Record', or procuring a larger site. It also demonstrates how liquid cooling at higher rack power density can be exploited to pack a lot of capacity into a relatively small amount of space.

Tables 4.7 and 4.8 compare annualised TUE and WUE values of a typical 5MW air-cooled data hall with those for an equivalent 5MW liquid-cooled data hall. The TUE value includes the losses associated with power consumption in the distribution/fan/pump power within the IT equipment as part of the calculation. The water consumption considered in this calculation is the water supplied to the hybrid cooler and pressurisation AHU. No domestic water consumption has been considered.

The air-cooled solution is assumed to be a fan wall system with primary cooling achieved by a hybrid cooler with chilled water temperatures of 25°C /40°C. In the operating hours that the hybrid cooler cannot achieve full cooling (less than 4% of the year) a chiller trims the remaining missing duty. The fan walls represent the biggest contribution to the pPUE.

The water-cooled system is based upon the Iceotope KU:L Box liquid cooling technology, which is assumed to be serving a TCS data hall circuit of 40.5°C/46°C, and plate heat exchangers to reject heat into a secondary circuit of 38°C /45°C. Because of the comparatively warm water temperatures on the secondary circuit, the system can achieve full cooling for 100% of the year with the hybrid coolers alone. In addition, the elevated water temperatures increase the external temperature that triggers evaporative cooling 'wet mode' of the hybrid coolers. Therefore, the annual water consumption of the water-cooled halls is dramatically reduced. The omission of the fan walls in this solution also significantly reduces the pPUE.

When looking at either the air-cooled or the liquid-cooled solutions, the TUE and Mech pPUE figures are very similar between the two locations. This is because in the air-cooled solution, the fan walls represent 68% of the mechanical losses which are a constant against IT load only, whilst the hybrid coolers only represent 10% of the mechanical losses. In the water-cooled solution, the hybrid cooler only represents around 20% of the mechanical losses. That means that changing the differences between climatic conditions for the two site locations does not have any appreciable impact upon the annualised PUE for either of the two systems.

However because the weather file for Richmond VA has a lot more hours in the year above 28°C (the evaporative 'wet mode' trigger point for the Hybrid coolers) the amount of water use in this location in a year is more than twice that consumed in San Jose CA.

The ITE Rack Power is the total power supplied to the racks within the data hall, whereas the ITE Power is the power supplied to the electronic IT equipment within the rack. The difference between the two figures reflects the power consumption of the server air cooling fans for the air-cooled option, and the power consumption of the dielectric fluid circulating pumps, within the Iceotope KU:L Boxes. The actual ITE power figures inform the TUE calculations, (see 1.7.6) and reflect the fact that with the liquid-cooled racks, less of the power is consumed by the fluid pumps than is used by the server fans within the air-cooled racks. This means that for a 5.0MW ITE hall, a greater proportion of the

supplied power is delivered to the IT equipment than in the equivalent air-cooled halls. For Scenario #1, when applied to the single liquid-cooled suite under consideration, the increase in available ITE power is 350kW representing a 7.8% uplift. If this were to be applied across the entire 5 Hall campus, the increase in power supplied to the IT electronics – as opposed to the rack - would amount to over 1.0MW. This is a further significant, but often overlooked, advantage of liquid cooling.

Richmond, VA, USA	Total ITE Rack Power [kW]	Total ITE Power [kW]	PUE	Mech pPUE	WUE	TUE	Mech pTUE	TWUE
Air-Cooled Data Hall	5000	4500	1.14	1.07	0.482	1.26	1.19	0.535
Liquid-Cooled Data Hall	5000	4850	1.09	1.02	0.074	1.12	1.06	0.077

Table 4.7: Comparison of PUE and WUE for Air-Cooled and Liquid-Cooled Data Halls - Richmond Virginia

San Jose, CA, USA	Total ITE Rack Power [kW]	Total ITE Power [kW]	PUE	Mech pPUE	WUE	TUE	Mech pTUE	TWUE
Air-Cooled Data Hall	5000	4500	1.14	1.07	0.231	1.26	1.19	0.257
Liquid-Cooled Data Hall	5000	4850	1.09	1.02	0.026	1.12	1.05	0.027

Table 4.8: Comparison of PUE and WUE for Air-Cooled and Liquid-Cooled Data Halls – San Jose, California

For the purposes of the assessment, the Rack Power density for the Air-Cooled option is estimated at 5.7kW/rack and 840 racks are needed to meet the design ITE Rack Power load for the hall of 5000kW, although for the air-cooled halls, only 4500kW actually reaches the server electronics, because up to 500kW is consumed by the server fans. As has been outlined in Section 2.0, the Iceotope KU:L Box Technology can comfortably host ITE rack power loads of up to 60kW/rack. For the purposes of this scenario, JRWS is planning to build a further 5.0MW data hall on a site that is too small to accommodate a 5th air-cooled hall, designed in accordance with their standard Plan of Record. By adopting liquid-cooled racks of a higher power density, the number of racks required to accommodate 5.0MW of ITE Rack Power is reduced. Even just increasing rack power density by a modest 3kW/Rack reduces the rack count by over 250 racks. The implications of this of spatial requirements of the data hall is shown on Table 4.9 and illustrated in Figure 4.10.

Data Hall Capacity [kW]	Rack Power Density [kW/rack]	Cooling Strategy	Racks [Quantity]	Data Hall Footprint [m ²]
5,000	5.7	Air Cooled	840	1,667
5,000	8.5	Air or Liquid Cooled	588	1,203
5,000	14.0	Liquid Cooled	336	741
5,000	30.0	Liquid Cooled	168	432

Table 4.9: Comparison of Spatial Requirements for Air-Cooled and Liquid-Cooled Racks

From this JRWS' ambition to increase the capacity of their two facilities by a further 5.0MW can be accomplished in a significantly smaller footprint than their standard air-cooled data halls. The results in the following section also indicate a significantly reduced Total Cost of Cooling per KW of ITE Rack Power in the new hall.

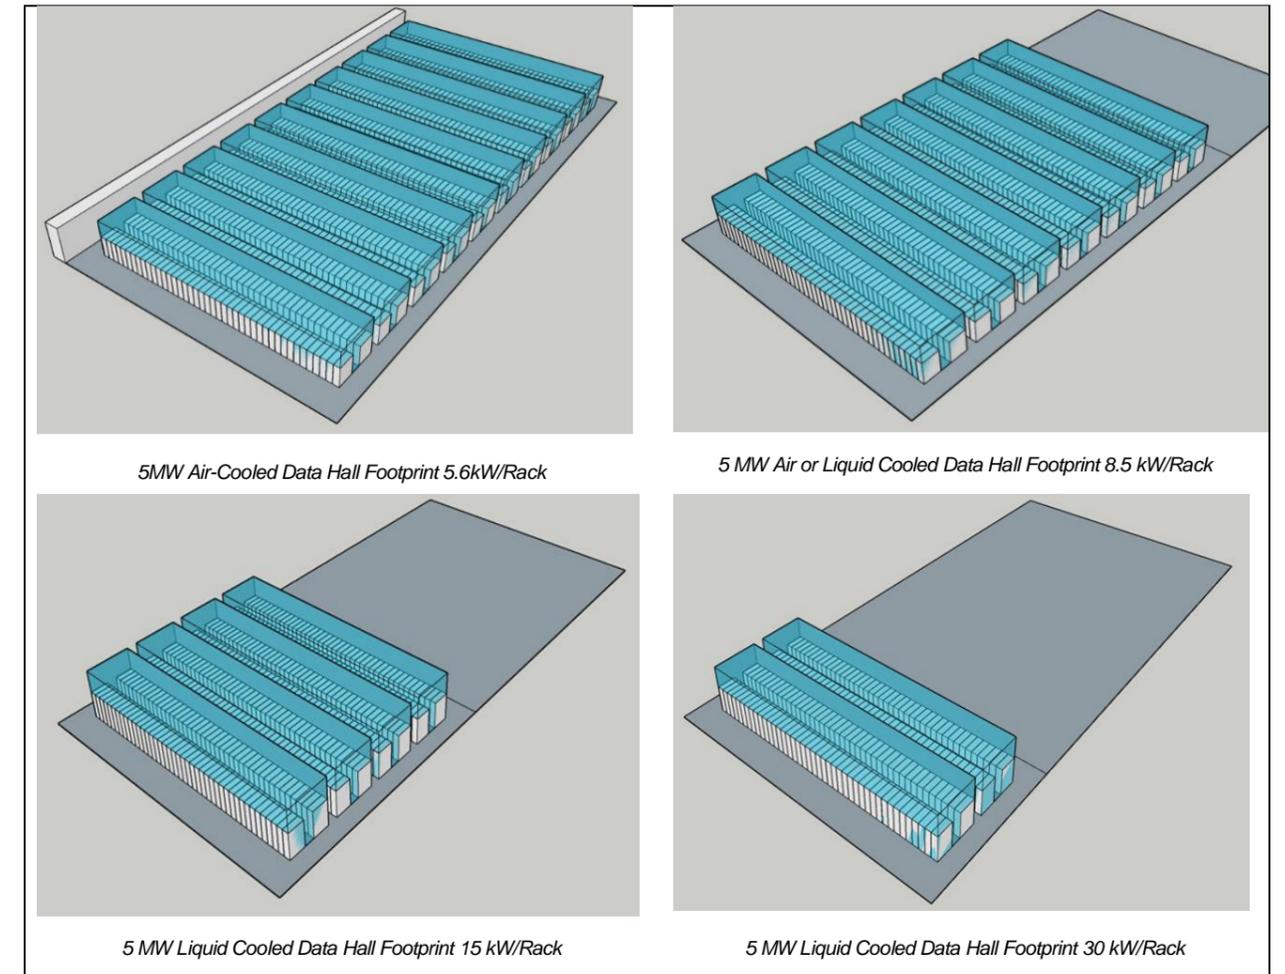


Figure 4.10: Comparison of Spatial Requirements at Different Rack Power Densities

4.6.3 Hyper-Scale OPEX Comparison

Based upon the unit costs provided in Table 4.1, the OPEX cost comparison between the air-cooled and liquid-cooled data halls is as follows:

Scenario #1 Richmond Virginia	ITE Rack Power [kW]	Total non-IT Power Demand [kWh/year]	Total IT Power Demand [kWh/year]	Total Power Demand [kWh/year]	Total water usage [l/year]
Air Cooled	5,000	6,021,715	43,800,000	49,821,715	21,091,582
Liquid Cooled	5,000	3,970,423	43,800,000	47,770,423	3,255,202

Table 4.10: OPEX Power and Water Consumption Comparison for Richmond, Virginia

Scenario #1 Richmond Virginia	electricity price [£/kWh]	water price [£/m ³]	Expected Power Supply Cost [£/year]	Expected Water Supply Cost [£/year]	Expected OPEX [£/year]	Expected IT Specific OPEX [£/kW/year]
Air Cooled	£0.1064	£1.09	£ 5,301,030	£ 23,049	£ 5,324,079	£ 1,065
Liquid Cooled	£0.1064	£1.09	£ 5,082,773	£ 3,557	£ 5,086,330	£ 1,017

Table 4.11: OPEX Cost Comparison for Richmond, Virginia

Scenario 1 San Jose California	ITE Rack Power [kW]	Total non-IT Power Demand [kWh/year]	Total IT Power Demand [kWh/year]	Total Power Demand [kWh/year]	Total water usage [l/year]
Air Cooled	5,000	5,945,054	43,800,000	49,745,054	10,136,234
Liquid Cooled	5,000	3,856,116	43,800,000	47,656,116	1,149,766

Table 4.12: OPEX Power and Water Consumption Comparison for San Jose, California

Scenario 1 San Jose California	electricity price [£/kWh]	water price [£/m ³]	Expected Power Supply Cost [£/year]	Expected Water Supply Cost [£/year]	Expected OPEX [£/year]	Expected IT Specific OPEX [£/kW/year]
Air Cooled	£0.1064	£1.09	£5,292,874	£11,077	£5,303,951	£1,061
Liquid Cooled	£0.1064	£1.09	£5,070,611	£1,256	£5,071,867	£1,014

Table 4.13: OPEX Cost Comparison for San Jose, California

4.7 Notional Data Centre Comparison Scenario #2 – The Co-Location Developer

4.7.1 Introduction – Docklands Hosting

DOCKLANDS HOSTING

Docklands Hosting is a Co-Location and Hosting provider that operates several sites within the circumference of the M25 motorway around London, UK.

Data halls in their multi-storey East London site are fully let, except for one remaining fallow hall, which Docklands is proposing to fit out. Previous data hall fit outs have each supported 1.6MW ITE at a power density of 4.8kW per m². Each hall accommodates 200

racks at Cooling is provided by Indirect Air Cooling (IAC) units mounted on external plant gantries. Thermal conditions in the data halls are maintained within the ASHRAE TC9.9 2015 Recommended Range; this being a supply condition at the server inlet of 18-27°C db, with a humidity range between -9°C DP and 15°C DP and 60% rh. Volumetric airflow rate equates to a temperature rise of 12K across the servers.

Most of the halls within the facility have been fitted-out are now fully let. However, Docklands Hosting has noted that much of the ITE rack power capacity contracted by their existing customers has failed to materialise and that their existing halls, although fully populated with customers racks, are all operating some way below their design capacity. This has resulted in significant stranded capacity within the site sub-station, the capacity and availability of which Docklands Hosting has invested in. Docklands Hosting is therefore considering ways in which the fitting-out of their final fallow data hall could make use of this stranded capacity. Upgrading the hall's UPS and electrical switchgear can be accomplished, and the generator capacity for the site is already designed to match the design ITE capacity of the building. However, increasing the capacity of air-cooling supplied to the data hall is not feasible, as further IAC units cannot be accommodated on the external plant gantry.

The schematic arrangement of the indirect air-cooling system that serves the existing data halls is shown in Figure 4.11

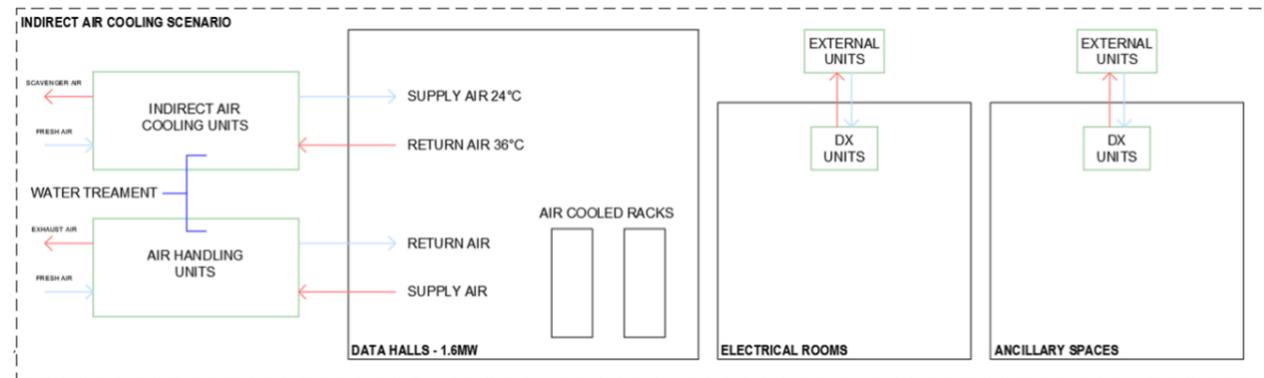


Figure 4.11: Schematic Arrangement of 1.6MW Air-Cooled Data Hall, supported by IAC Units

To access the trapped capacity held within the sub-station, Docklands Hosting is therefore considering a mixed deployment of air-cooled and high-density liquid cooled racks within the same data hall. A further factor that has encouraged them to review this option is a number of customer RfP's that they have received in recent weeks from Edge Providers, who are seeking to develop a HPC+AI deployment and have made enquiries about Dockland Hosting's ability to support high-density liquid-cooled racks.

The new system would comprise an FWS operating at 38-43°C, which would be served by Hybrid Coolers at roof level. The FWS would be linked to a TCS operating at 40-46.5°C, via an on-floor CDU. The overall rack count within the data hall would remain the same, at 200 racks, of which 40 racks would be converted to liquid cooling. This will allow the capacity of the IAC units to be focused on the remaining 160 air-cooled racks, allowing their capacity to be boosted to 10kW/rack. The 40 racks converted to liquid cooling would each be rated at 20kW, allowing the capacity of the data hall to be increased to 2.4MW. The performance of the Liquid-cooled racks has been based upon Iceotope KU:L Box technology.

The schematic arrangement of the proposed modified data hall is illustrated in Figure 4.12

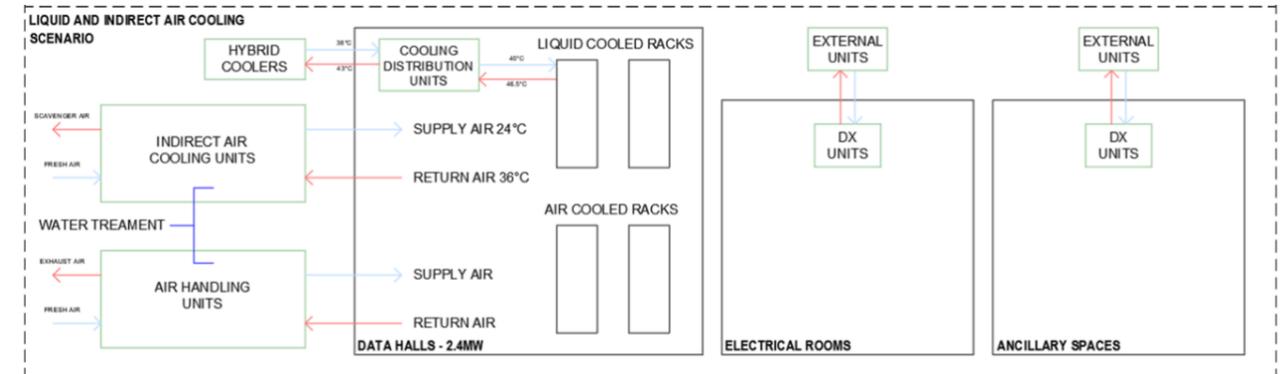


Figure 4.12: Schematic Arrangement of 2.4MW Mixed Air-Cooled and Liquid-Cooled Data Hall

4.7.2 Co-Location Performance Comparison

Objective of this comparison is to show how the capacity of an existing air-cooled Data Hall can be boosted by a limited introduction of liquid-cooled racks and how this can be done in a way that improves overall PUE and WUE.

Design parameters for the base scheme 1.6MW air-cooled data hall and the 2.4MW mixed air & water-cooled data hall are shown in Table 4.14. Notional Rack layouts for the two options are shown in Figure 4.13. Liquid-cooled racks have been shown on the ends of each row, on the far side of the data hall from the external plant gantry. This is to allow the air-cooling capacity of the IAC units to be concentrated on the air-cooled racks. In practice, liquid-cooled racks could be grouped differently, as they do not need to be arranged to follow established 'hot' and 'cold' aisle arrangements within the room.

Table 4.15 compares annualised TUE and WUE for the two schemes. The TUE value includes the losses associated with power consumption in the distribution/fan/pump power within the IT equipment as part of the calculation. The water consumption considered in this calculation is the water supplied to the IAC units and hybrid cooler. No domestic water consumption has been considered.

The indirect air-cooling system is a modular type @ N+1, providing cooling air to the data hall at 24-26°C, with a return temperature of 36-38°C. Each unit uses ambient air to cool process air, via an interposing air-to-air heat exchanger. During periods of warm weather, the cooling capacity of the unit is boosted by water sprays on the scavenger side of the heat exchanger to provide evaporative cooling. The water sprays are enabled when the ambient dry bulb air temperature reaches or exceeds 15°C. Water consumption has been reduced by means of integral DX Trim Cooling.

A fresh air supply is provided to the data hall by a dedicated AHU, for the purposes of pressurisation and humidity control. The AHU includes an integral DX cooling coil, to provide trim cooling and de-humidification when the ambient dry bulb temperature exceeds 25°C. No winter humidification is required. Winter trim heating is via heat reclaim.

Electrical rooms are provided with dedicated DX cooling units.

Fit-Out Option	Total ITE Rack Power [kW]	Data Hall Area [m ²]	Rack Count [Quantity]	IT Power Density [kW/m ²]	Rack Power Density [kW/Rack]
Base - Indirect Air Cooled	1,600	336	200 Air Cooled	4.8	8kW/Rack (Air-Cooled)
Indirect Air-Cooled + Liquid Cooling	2,400	336	160 Air Cooled 40 Liquid Cooled	7.14	10kW/Rack (Air-Cooled) 20kW/Rack (Liquid-Cooled)

Table 4.14: Scenario #2 - Data Hall and Rack Design Parameters

Scenario #2, London, UK	Total ITE Rack Power [kW]	Total ITE Power [kW]	PUE	Mech pPUE	WUE	TUE	Mech pTUE	TWUE
Air-Cooled Data Hall	1,600	1,440	1.1539	1.10	0.19	1.28	1.22	0.2057
Air + Liquid-Cooled Data Hall	2,400	2,216	1.1439	1.09	0.014	1.24	1.18	0.015

Table 4.15 – Scenario #2 - Comparison of PUE and WUE for Air-Cooled and Mixed Air & Water-Cooled Data Halls

Table 4.15 also indicates the amount of power delivered to the IT electronics (Total ITE Power) and compares this with the rated ITE capacity of the data hall (Total ITE Rack Power). The difference between these two values accounts for the in-rack parasitic loads present; the server cooling fans in the case of the air-cooled racks, and the dielectric fluid pumps inside the KU:L Box cubbies, for the liquid-cooled racks. For the mixed air and liquid-cooled data hall, the proportion of ITE power to ITE rack power is slightly greater than for the base-option air-cooled scheme. As per the earlier comments, this reflects the lower power draw/rack of the dielectric fluid pumps, compared with the server fans.

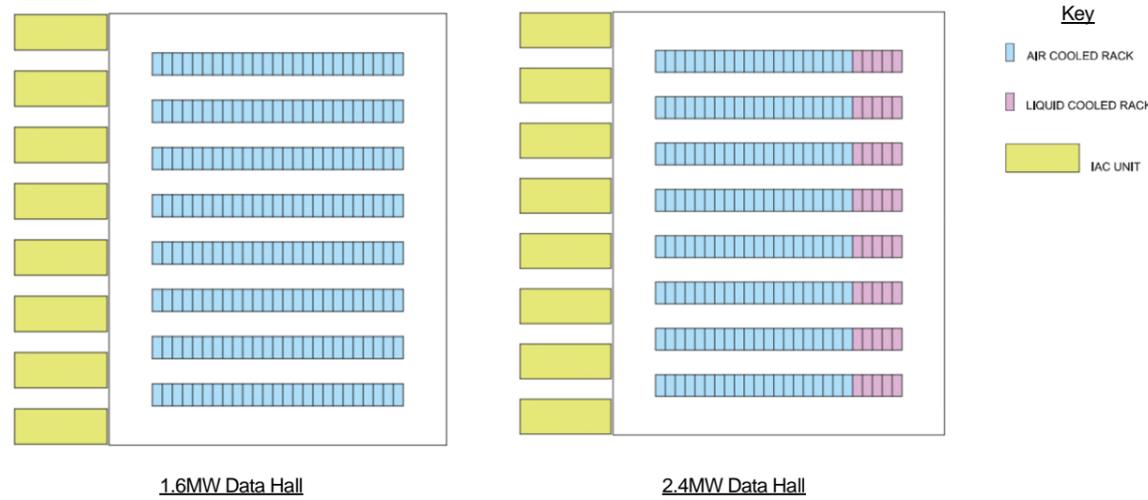


Figure 4.13 - Rack Layout for 1.6MW and 2.4MW Data Halls

Figure 4.14 and Figure 4.15 illustrates the breakdown of energy use for the two options under consideration. For both cases, although very energy and water efficient, the IAC units are unsurprisingly the largest contributors to energy consumption. This could be improved modestly, by increasing the amount of evaporative cooling available, either by reducing top-up DX trim, or bringing on the evaporative sprays earlier. However, this would come at the cost of increased WUE.

The introduction of liquid cooling incurs an energy penalty for FWS and TCS pumping and for Hybrid-Cooler Fan Power, although together these only account for 13% of the overall power consumption of the air + liquid-cooled option, despite the fact that the liquid-cooled component accounts for one third of the ITE load present in the hall. The mechanical plant power associated with liquid cooling is almost entirely made up of FWS and TCS pumping power. Given the large temperature difference available between the FWS and ambient air, fan power consumed by the Hybrid-Cooler Fans is quite small in comparison to other loads.

As can be seen from Table 4.15 the Base Option is already very energy efficient, achieving a Mechanical partial PUE of 1.10. However, by supplementing this system with liquid cooling, it is possible to improve this still further to 1.09, reducing the 'Total Cost of Cooling per KW of ITE Power' by over 23%. At the same time, the total ITE power delivered to the servers is increased from 1,440kW to 2,244kW.

Had the necessary power been available on site, an even greater reduction in the Cost of Cooling per KW of ITE Power could of course have been achieved by using liquid cooling across all 200 racks in the new hall.

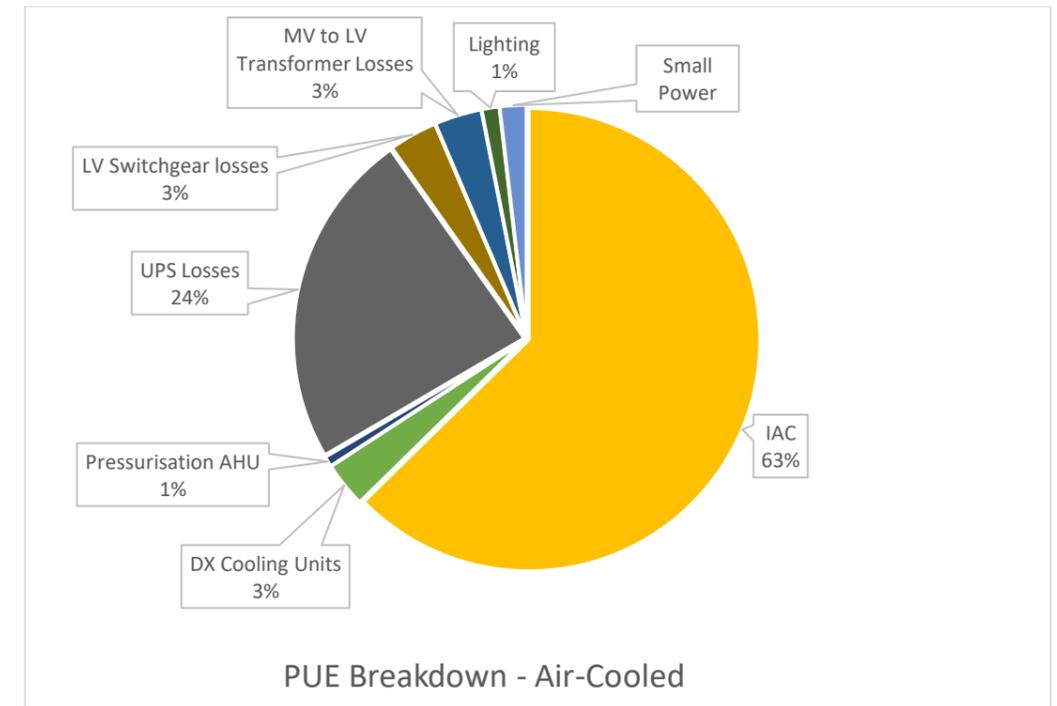


Figure 4.14 - Scenario #2 - Breakdown of Infrastructure Energy Use for Base Option

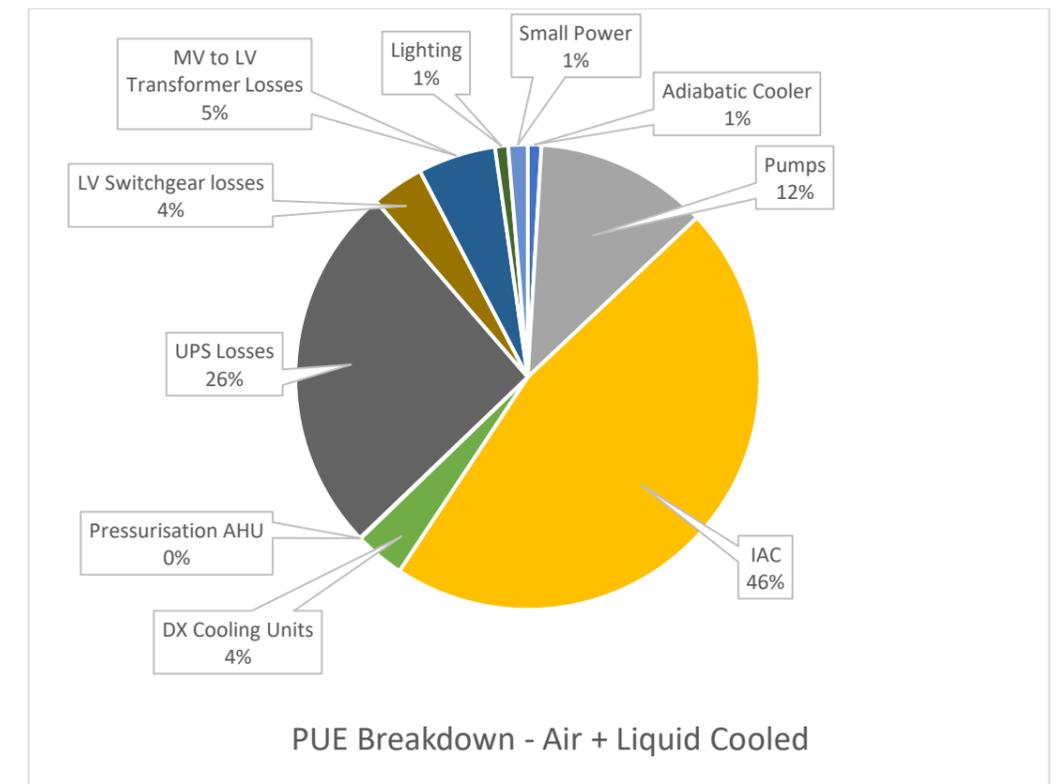


Figure 4.15 - Scenario #2 - Breakdown of Infrastructure Energy Use for Air and Liquid-Cooled Option

4.7.3 Co-Location OPEX Comparison

Based upon the unit costs provided in Table 4.1, the OPEX cost comparison between the base-scheme air-cooled data hall and the mixed air + liquid-cooled data hall are shown on Table 4.16 and Table 4.17.

The inclusion of the additional 800kW of ITE power increases OPEX cost by over £862k, but in terms of unit cost per IT kW per annum, operating costs are actually reduced from £1,183.10 per year, to £1,148.09, a saving of approximately 3%.

Based upon the additional Scenario 2 London UK	ITE Rack Power [kW]	Total non-IT Power Demand [kWh/year]	Total IT Power Demand [kWh/year]	Total Power Demand [kWh/year]	Total Water Usage [l/year]
Air Cooled	1,600	2,373,348	15,417,600	17,790,948	2,854,400
Mixed Air+ Liquid	2,400	3,257,636	22,635,840	25,893,476	314,224

Table 4.16: Scenario #2 - OPEX Power and Water Consumption Comparison for London UK

Scenario 2 London UK	electricity price [£/kWh]	water price [£/m3]	Expected Power Supply Cost [£/year]	Expected Water Supply Cost [£/year]	Expected OPEX [£/year]	Expected IT Specific OPEX [£/kW/year]
Air Cooled	£0.1064	£1.09	£1,892,957	£3,119	£1,892,960	£1,183
Mixed Air+ Liquid	£0.1064	£1.09	£2,755,066	£343	£2,755,409	£1,148

Table 4.17: Scenario #2 - OPEX Cost Comparison for London, UK

It is important to note that the Total Cost of Cooling per KW of ITE Power could have been reduced significantly further, by switching a greater proportion of the air-cooled racks to liquid cooling.

4.8 Notional Data Centre Comparison #3 Tropical Region Performance

4.8.1 Introduction – GREENDATA

GreenData

GreenData is a CoLo and Edge Operator serving the AsiaPac Region. They are proposing to build a cluster of small 1.0MW ITE Data Centers in Singapore. Their previous Data Center builds in the region have been air-cooled and utilised water-cooled chillers + open circuit cooling towers and chilled water CRAHs @ N+2 within the room. Their standard modular-build arrangement provides a data hall of approx. 415m², accommodating 132 racks at a power density of 2.4kW/m² and an average rack power density of 7.5kW/rack. Thermal conditions in the data halls are maintain within the ASHRAE TC9.9 2015 Recommended Range; this being a supply condition at the server inlet of 18-27°C db, with a humidity range between -9°C DP and 15°C DP and 60% rh. Volumetric airflow rate equates to a temperature rise of 12K across the servers.

The schematic arrangement of Greendata’s standard air-cooled modular build data centre is shown in Figure 4.16.

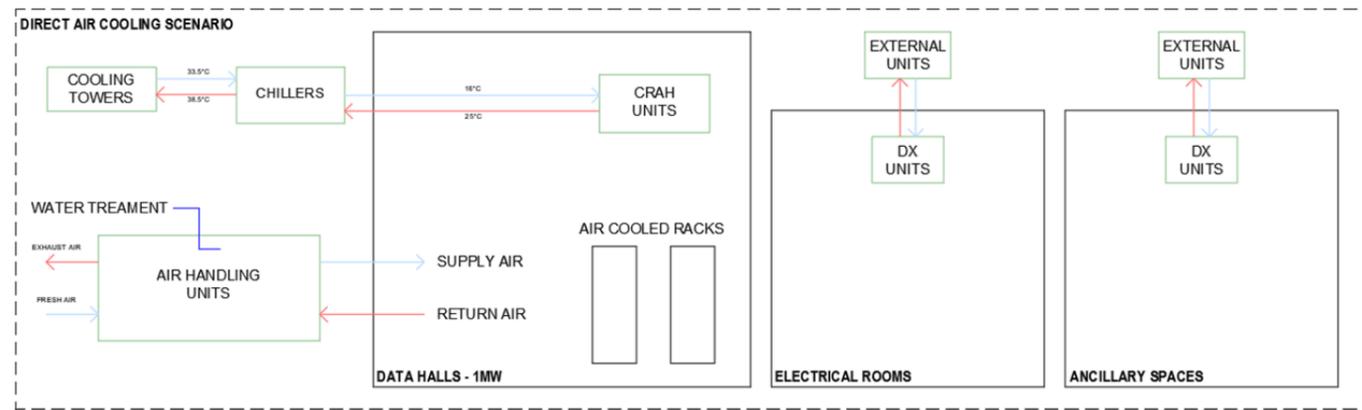


Figure 4.16: Schematic Arrangement of Air-Cooled 1.0MW Base Option, supported by Water-Cooled Chillers and CRAHs

Increased demand to host high-burn Edge applications, coupled with rising competition and the high cost of suitable sites in Singapore, has encouraged GreenData to investigate the use of liquid cooling.

The alternative scheme would comprise an FWS operating at 38-43°C, which would be served by Hybrid Coolers at roof level. The FWS would be linked to a TCS operating at 40-46.5°C, via an on-floor CDU. Liquid cooling to the racks is supplemented by air-cooled DX CRAC units to offset room gains. Racks would contain Iceotope KU:L Box liquid cooling technology, and each be rated at nominal 20kW. Rack count within the data hall would be reduced to 50 racks. This would allow a considerable reduction in data hall footprint.

The schematic arrangement of the alternative liquid-cooled scheme is shown in Figure 4.17

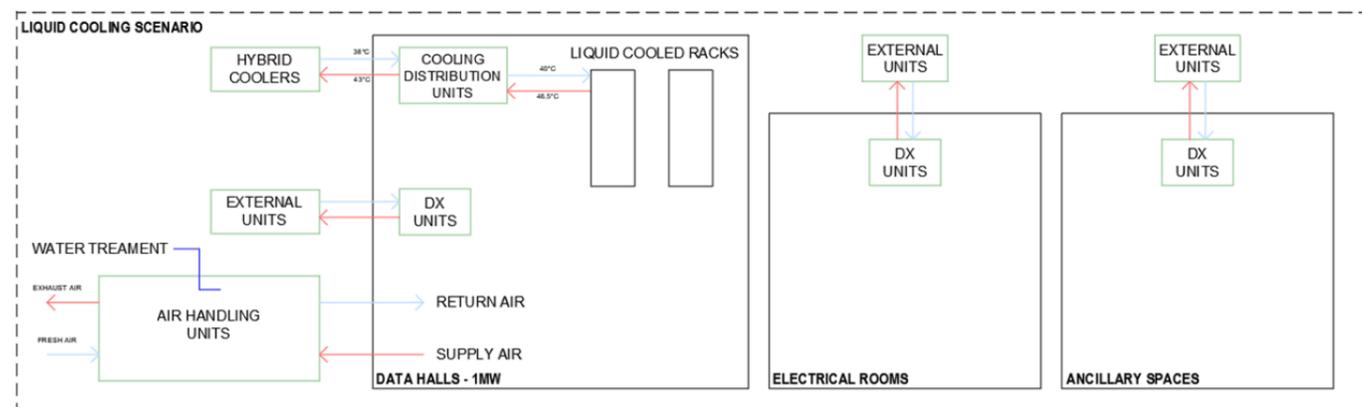


Figure 4.17: Schematic Arrangement of Liquid-Cooled 1.0MW Option, supported by Hybrid-Coolers.

4.8.2 Tropical Region Performance Comparison

Objectives of this comparison are to compare the PUE and WUE of the air-cooled and liquid cooled data halls under tropical weather conditions and also to show the respective data hall footprints, thereby demonstrating the space saving to be had from adopting liquid cooling.

Design parameters for the two schemes are shown on Table 4.18. While maintaining the same Total ITE Rack Power to the hall of 1.0MW, the introduction of Liquid cooling technology allows for a significant reduction in data hall footprint, compared to the base option modular air-cooled scheme. The scale of this reduction is illustrated in Figure 4.20.

Fit-Out Option	Total ITE Rack Power [kW]	Data Hall Area [m ²]	Rack Count [Quantity]	IT Power Density [kW/m ²]	Rack Power Density [kW/Rack]
Base – Water-Cooled Chillers + CRAHs	1,000	415	132 Air Cooled	2.4kW/m ²	7.5kW/Rack (Air-Cooled)
Liquid Cooling + DX CRACs	1,000	290	50 Liquid Cooled	3.45kW/m ²	20kW/Rack (Liquid-Cooled)

Table 4.18: Scenario #3 - Data Hall and Rack Design Parameters

Singapore, Singapore	Total ITE Rack Power [kW]	Total ITE Power [kW]	PUE	Mech pPUE	WUE	TUE	Mech pTUE	TWUE
Base Option – Water-Cooled Chillers + CRAHs	1000	900	1.56	1.50	1.519	1.73	1.67	1.6876
Liquid Cooling + DX CRACs	1000	970	1.12	1.06	0.071	1.15	1.10	0.0728

Table 4.19: Scenario #3 - Comparison of PUE and WUE for Air-Cooled and Mixed Air & Water-Cooled Data Halls

As with the previous scenarios, Table 4.19 indicates the amount of power delivered to the IT electronics (Total ITE Power) and the rated ITE capacity of the data hall (Total ITE Rack Power). Given that the options under consideration are either all air-cooled, or all liquid-cooled data hall, the difference between the two is quite marked, with only 3% of the ITE rack power being consumed by the parasitic load of the dielectric coolant pumps in the KU:L Box cubbies, compared with 10% by the server fans within the air-cooled racks.

The PUE for the liquid-cooled option is considerably improved over that of the base option. This is due to the omission of the compressor input power associated with the water-cooled chillers of the air-cooled option. Despite the high ambient temperature, no compressor-based cooling is needed for the liquid-cooled option, due to the high operating temperature of the FWS, which can be maintained by the hybrid coolers alone. The WUE is also greatly improved, as the hybrid coolers only need to operate in wet, evaporative cooling mode for a limited period of the year, whereas the open-circuit cooling towers need to operate in evaporative cooling continuously.

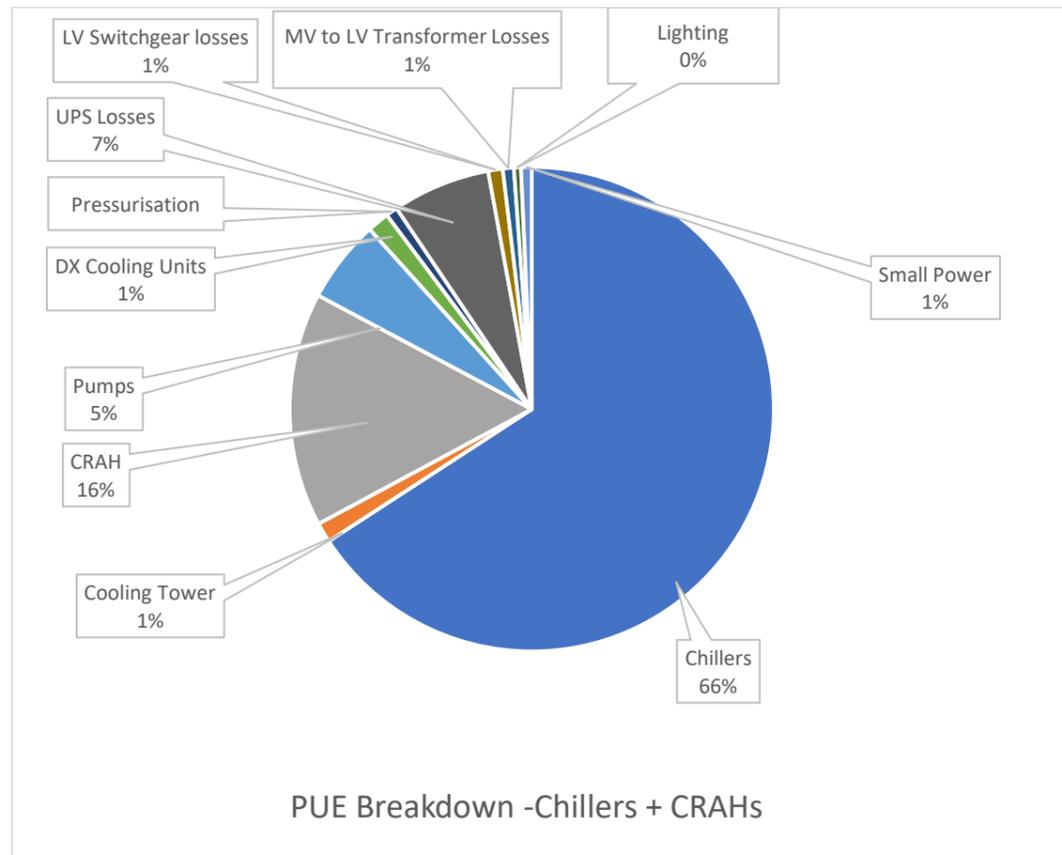


Figure 4.18 - Scenario #3 - Breakdown of Infrastructure Energy Usage - Base Option

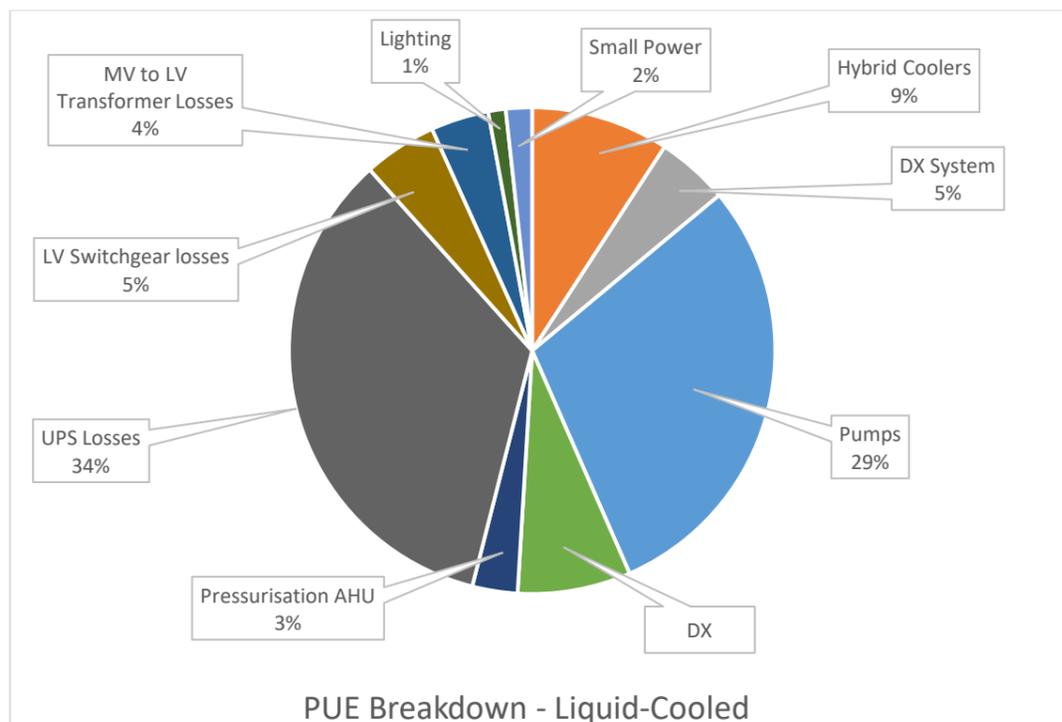


Figure 4.19 - Scenario #3 - Breakdown of Infrastructure Energy Usage - Liquid-Cooled Option

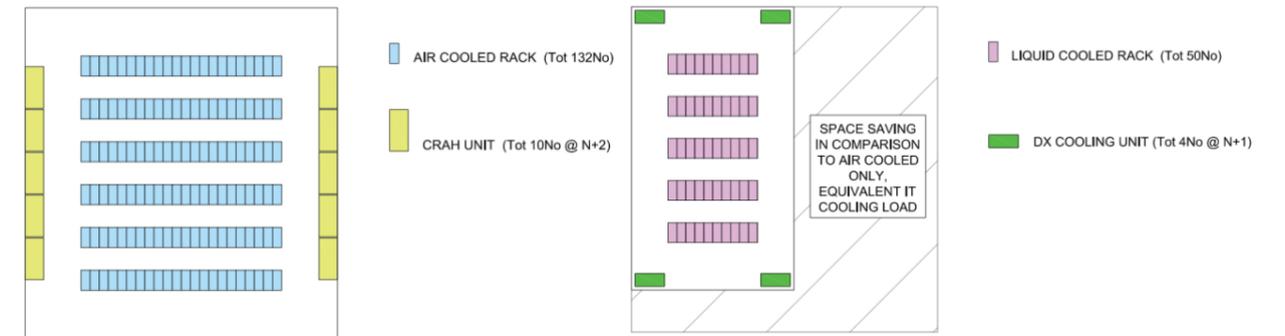


Figure 4.20: Footprint Comparison of 1.0MW Air-Cooled and Liquid-Cooled Data Halls

One of GreenData’s primary objectives was to achieve a reduction in data hall footprint, without compromising ITE capacity. The scale of reduction possible from the adoption of liquid cooling is illustrated in Figure 4.20. The 50 x 20kW Racks required for the liquid-cooled option need approx. 290m², compared with 415m² needed to accommodate 132 air-cooled racks, plus the associated CRAHs, a reduction of around 30%. The footprint could of course have been reduced even further by increasing the new rack density by up to 60kW/rack.

The footprint saving shown above does not take account of the reduction in mechanical cooling plant, as the water-cooled chillers + cooling towers required to support the base option, would take up a lot more space than the hybrid coolers associated with the liquid-cooled option. There would also be a reduction in process water storage tank size and water treatment plant.

4.8.3 Co-Location OPEX Comparison

Based upon the unit costs provided in Table 4.1, the OPEX cost comparison between the base-scheme air-cooled scheme and alternative liquid-cooled option are shown on Table 4.20 and Table 4.21.

Scenario #3, Singapore	ITE Rack Power [kW]	Total non-IT Power Demand [kWh/year]	Total ITE Power Demand [kWh/year]	Total Power Demand [kWh/year]	Total water usage [l/year]
Air-Cooled	1,000	5,352,326	9,636,000	14,988,326	14,635,165
Liquid-Cooled	1,000	1,018,856	8,760,000	9,778,856	618,836

Table 4.20: Scenario #3 – OPEX Power and Water Consumption Comparison for Singapore

Scenario #3, Singapore	electricity price [€/kWh]	water price [€/m ³]	Expected Power Supply Cost [€/year]	Expected Water Supply Cost [€/year]	Expected OPEX [€/year]	Expected IT Specific OPEX [€/kW/year]
Air Cooled	€0.1064	€1.09	€1,594,757	€15,993	€1,610,751	€1,611
Mixed Air+ Liquid	€0.1064	€1.09	€1,040,470	€676	€1,041,147	€1.041

Table 4.21: Scenario #3 – OPEX Cost Comparison for Singapore

5.0

Conclusion

5.0 Conclusions

5.1 Summary of Comparison Results

In every case, the specific climate conditions at each of the proposed locations have a major influence on the performance of the cooling system. Nevertheless, the KU:L Box liquid cooling technology consistently outperforms conventional air-cooling, in terms of both PUE and WUE. This is principally due to the much higher operating temperature of the FWS, compared to the cooling mediums used for the air-cooled solutions. In all air-cooled cases, considerable energy and water is consumed to arrive at a supply air condition that falls within the required thermal envelope. The need for this is completely avoided with liquid cooling. Even in tropical climates, the operating temperature of the FWS is high enough for the Hybrid coolers to operate in economiser free cooling mode for much of the time, and under peak ambient conditions, sufficient capacity can be maintained by reverting to 'wet' evaporative cooling mode. There is no requirement for energy-intensive compressor-based trim-cooling.

A further consistent benefit is the reduction in rack-count and data hall area that can be achieved through higher ITE rack power density. This is quite clearly shown in all three Scenarios.

5.1.1 Summary of results by region

- Scenario #1- JRWS:** The study considers two hyperscale air-cooled data centres, where lack of space to build a larger facility following the established 'Plan of Record' design prohibits full utilisation of the power supplies available to each site. By adopting liquid cooling, a further 5MW data hall can be added to each campus, using less space than would be required for an equivalent air-cooled data hall, and which is capable of superior energy and water efficiency.
- One of the facilities is in Richmond, VA, USA, where the climate is humid subtropical; the other data centre is in San Jose, CA, USA, where the climate is temperate Mediterranean type. In both locations the improvement in efficiency from the adoption of liquid cooling results in a substantial saving in energy and water costs.
- Scenario #2 – Docklands Hosting:** The studies considers an extension to an existing facility, where some of the air-cooled racks within the space are converted to liquid cooling, targeting increased ITE capacity within the data hall. Overall rack count of 200 racks remains unchanged.
- The facility is located in London, UK which has an Oceanic climate. The study shows that by converting 40 of the 200 racks to liquid cooling, the capacity of the remaining 160 air-cooled racks can be boosted from 8kW to 10kW/rack and the 40 liquid cooled racks could each support 20kW/rack. Overall, the ITE capacity of the data hall would be increased from 1.6MW to 2.4MW. Cooling energy will increase to support the additional load, but the Total Cost of Cooling per kW of ITE power reduces substantially, as does the PUE. Had there been additional power capacity available on site, it would of course have been possible to generate even greater savings by adopting liquid cooling across all 200 racks in the new hall.
- Scenario #3 – GreenData:** The study considers facilities built in Singapore, where the climate is typical is Tropical Rainforest, and compares a 1.0MW ITE Air Cooled Data Hall housing 132 racks, to a Liquid Cooled facility of equivalent ITE capacity.
- The study shows that by adopting liquid cooled racks of higher density, the total rack count could be reduced to 50 racks. This would facilitate a considerable reduction in data hall footprint, which would be an important improvement in the design, considering the cost and availability of development land within the region. Furthermore, moving away from an energy-intensive chilled water-based cooling system, would facilitate improvements in both PUE and WUE.

5.1.2 OPEX Cost Saving Arising from the Adoption of Liquid Cooling

In each case, the adoption of liquid cooling systems brings improvements in OPEX costs. This is illustrated more clearly when expressed in terms of the relative costs of cooling. The following tables set out the comparable costs for each scenario, comparing the Base Case Air-Cooled Options with the Liquid-Cooled alternative solutions.

Table 5.1 shows the OPEX costs for the Base Case Air-Cooled solutions, expressed in terms of Total Cost per kW of ITE power. Table 5.2 shows the same information for the liquid-cooled alternative solutions.

Base Case	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Application	5MW Hyperscale		2.4MW Co-Lo	1MW Edge DC
Location	Richmond VA	San Jose CA	London UK	Singapore
Tech Comparison	Fan Wall with Hybrid Cooler		IEC	Chiller + CRAH
Cost of Power for Cooling	£1,106,742	£1,098,586	£550,785	£755,899
Cost of Water for Cooling	£23,049	£11,077	£ 3,119	£ 15,993
Total Cost of Cooling	£1,129,791	£1,109,663	£553,904	£771,893
Cost of Power for Computing	£4,194,288	£4,194,288	£1,342,172	£838,858
Total Cost of Power	£5,301,030	£5,292,874	£1,892,957	£1,594,757
Total OPEX Cost	£ 5,324,079	£ 5,303,951	£ 1,896,076	£ 1,610,750
Total ITE Rack Power [kW]	5,000	5,000	1,600	1,000
Total ITE Power [kW]	4,500	4,500	1,440	900
Total Cost per KW of ITE Rack Power	£ 1,065	£ 1,061	£ 1,185	£ 1,611
Total Cost per KW of ITE Power	£ 1,183	£ 1,179	£ 1,317	£ 1,790
Total Cost of Cooling per KW of ITE Power	£ 251	£ 247	£ 385	£ 858

Table 5.1: Total OPEX Cost of Cooling per kW of ITE Power - Base Case

Liquid Cooling	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Cost of Power for Cooling	£562,263	£562,263	£663,514	£136,368
Cost of Water for Cooling	£ 3,557	£ 1,256	£ 343	£ 676
Total Cost of Cooling	£565,820	£563,519	£663,858	£137,044
Cost of Power for Computing	£4,520,510	£4,520,510	£2,091,552	£904,102
Total Cost of Power	£ 5,082,773	£5,082,773	£ 2,755,066	£ 1,040,470
Total OPEX Cost	£ 5,086,330	£ 5,084,029	£ 2,755,409	£ 1,041,146
Total ITE Rack Power [kW]	5,000	5,000	2,400	1,000
Total ITE Power [kW]	4,850	4,850	2,244	970
Total Cost per KW of ITE Rack Power	£ 1,017	£ 1,017	£ 1,148	£ 1,041
Total Cost per KW of ITE Power	£ 1,049	£ 1,048	£ 1,228	£ 1,073
Total Cost of Cooling per KW of ITE Power	£ 117	£ 116	£ 296	£ 141

Table 5.2: Total OPEX Cost of Cooling per kW of ITE Power – Liquid Cooling

Impact	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Water useage per KW of ITE Power	-85.7%	-89.5%	-92.9%	-96.1%
Power useage per KW of ITE Power	-11.0%	-10.9%	-6.6%	-39.5%
Total Cost per KW of Rack Power	-4.5%	-4.1%	-3.1%	-35.4%
Total Cost per KW of ITE Power	-11.4%	-11.1%	-6.7%	-40.0%
Total Cost of Cooling per KW of ITE Power	-53.5%	-52.9%	-23.1%	-83.5%
Increased ITE Capacity per Hall	*7.8%	*7.8%	55.8%	*7.8%
Footprint Required for additional Hall	n/a	n/a	n/a	-74.1%

*Note: The 7.8% increase in ITE Capacity per Hall in Scenarios 1a, 1b & 3 derives entirely from the reduced power draw of the on-board pumps in the liquid cooled design, compared to that of the server fans in the air-cooled base case. Substantial further increases could have been achieved by increasing the rack density and number of racks in the liquid-cooled design.

Table 5.3: Percentage Improvement in Performance between Base Case and Liquid Cooling Alternative

Table 5.3 illustrates the considerable impact of adopting liquid cooling, in terms of percentage improvement in key OPEX parameters. When considered in terms of the Cost of Cooling per kW of ITE Power, for Scenario #1, the reduction is in the order of 53% of the Base Case.

For Scenario #2, where only part of the data hall is to be converted to liquid cooling, the scheme would still achieve a reduction in OPEX cost in the order of 23%, compared to the Base Case. This is because the ITE Rack power within the data hall has increased by 50%, from 1.6MW to 2.4MW, but the cost of cooling (water and power) has only increased by around 20%, from £553,904 p.a. to £663,858 p.a. Had more power been available and a greater proportion of the racks been converted from air-cooled to liquid-cooled, an even greater saving could have been achieved.

For Scenario #3, the improvement is even more dramatic, with the adoption of liquid cooling generating a saving in the cost of cooling of over 83%, relative to the Base Case. This is because the Base Case utilised a chilled water system, which given climatic conditions in the region, is particularly energy and water intensive.

Note that for all three scenarios, the total cost of ITE power has also gone down. This reflects the increased efficiency of liquid cooling at rack level, when compared to the air-cooled Base Case. The dielectric fluid pumps in the liquid cooled servers do not consume as much power as the PSU fans in the air-cooled servers. So, although the overall ITE rack power for both options is the same, a greater proportion (approximately 7% of the total rack power) is available to be used by the electronic components – the ITE Power – with liquid cooling. This is shown by the increase in ITE Compute power for each case. Hence, in the case of Scenario #2, the 50% increase in ITE Rack Power in the data hall, translates into an increase of 55.8% in ITE power.

5.1.3 Commentary on Building Form and Plot Area

The most evident effect brought by the greater power density achievable with High Density Liquid Cooling Technology is the reduction in data hall space requirements, relative to the air-cooled equivalent.

Scenarios #1 and #3 make a direct comparison between air and liquid cooling and illustrate the reduction in rack count and data hall size that could be achieved by adopting an alternative design, based upon liquid cooling. Scenario #3 considers an existing data hall of fixed dimensions and rack capacity but shows how the ITE capacity of the room could be boosted to take maximise utilisation of the available power supply.

The reduction in building footprint would not be solely due to the rack count within the data hall. By its nature, air-cooling plant takes up considerable space. Liquid cooling systems will also need plant space for pump sets, CDUs and heat rejection plant. However, this would be more compact and there would be fewer constraints regarding its location in relation to the racks.

For Scenario #3, where the proposed layout considers the respective space requirements for in-room cooling units, the projected reduction in data hall footprint is in the order of 30%, compared to the Base Case. Detailed design of the heat rejection plant and chillers was outside the scope of the study. However, the Hybrid Coolers needed to support the

liquid-cooled option would occupy less external plant space than the chillers and cooling towers needed for the Base Case air-cooling option. Given that a primary objective was to reduce the size of the of the plot, then this would be an additional important benefit of the switch to liquid cooling.

5.1.4 Commentary on Environmental Impact

The greatest environmental impact of any data centre arises from the carbon emissions that are associated with the power that it will consume over its life. The more energy efficient the facility is, the less carbon it will produce, assuming consistent ITE capacity. Even where power is supplied from low-carbon, or renewable sources, the need for energy efficiency remains paramount, as a greater proportion of the power supplied to the facility can be directed towards its core purpose; supporting the data processing equipment, rather than being consumed by the supporting infrastructure.

Power Usage Effectiveness (PUE) and Total Power Usage Effectiveness (TUE) are lower for the liquid cooling options in all three scenarios. The clearest way to demonstrate the impact this would have upon carbon emissions is to express the reduction in terms of kg CO₂ per kW of ITE power. Actual CO₂ emissions for each scenario would be dependent upon the carbon content (kgCO₂/kWh) of the grid power supply available at each location and the contribution of any low-carbon or renewable power sources on site. In practice, these factors could vary considerably from one location to another, making direct comparisons between the three scenarios difficult. However, by applying a universal grid power carbon conversion factor to all four locations, it is possible to arrive at an indicative reduction in CO₂ emissions. These are shown in Table 5.4 below.

Base Case	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Application	5MW Hyperscale		2.4MW Co-Lo	1MW Edge DC
Location	Richmond VA	San Jose CA	London UK	Singapore
Tech Comparison	Fan Wall with Hybrid Cooler		IEC	Chiller + CRAH
Total Power Demand [kWh/year]	49,821,715	49,745,054	17,790,948	14,988,326
CO2 Emissions [kg/year]	11,615,435	11,597,562	4,147,782	3,494,378
Total ITE Power [kW]	4,500	4,500	1,440	900
kg of CO2 per kW of ITE Power [kg/kW per year]	2,581	2,577	2,880	3,883
Liquid Cooling	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Total Power Demand [kWh/year]	47,770,423	47,656,116	25,893,476	9,778,856
CO2 Emissions [kg/year]	11,137,196	11,110,547	6,036,805	2,279,842
Total ITE Power [kW]	4,850	4,850	2,244	970
kg of CO2 per kW of ITE Power [kg/kW per year]	2,296	2,291	2,690	2,350
Impact	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Percentage Reduction in CO2 Emissions per kW of ITE Power [kg of CO2/kW per year]	11.0%	11.1%	6.6%	39.5%

Note: Emissions estimate based upon data published by BEIS (Dept. for Business, Energy & Industrial Strategy). For 2020, UK Grid generation represents 0.23314 kg of CO₂/kWh.

Table 5.4: Percentage Reduction in Carbon Emissions per kW of ITE Power

Table 5.4 shows that all Scenarios considered would achieve a significant reduction in specific carbon emissions (i.e. kgCO₂/kW of ITE power, per year). Scenarios #1a and 1b follow the same technical design and have similar annual energy consumption. Both support 5MW of ITE Power. By adopting liquid cooling, carbon emissions could be reduced by around 11% per kW of ITE power, per year.

In real terms, the carbon emissions associated with Scenario #2 would increase from 4,147.782kg/year to 6,036,805kg/year. This is because converting a proportion of the racks within the hall to liquid cooling would allow the ITE capacity to be increased from 1.6MW to 2.4MW. However, the improvement in overall energy efficiency brought by the introduction of liquid cooling means that the specific carbon emissions are still reduced by more than 6%.

Scenario #3 shows the greatest reduction in specific carbon emissions. The liquid-cooled option would not require the energy intensive chiller plant used by the Base option, which means that overall carbon emissions for a 1.0MW data hall would reduce from 3,494,378 kg/year to 2,279,842 kg/year. This equates to a specific reduction of almost 40%.

For all Scenarios, Water Usage Effectiveness (WUE) for the liquid cooled solutions also reduces, compared to the air-cooled Base Options. However, because this measure is based upon the ITE Rack Power, it fails to recognise the full scale of the reduction in water consumption that can be achieved by switching to liquid cooling. The most effective way to quantify this is to calculate the litres of water consumed per annum for every kW of ITE Power that is processed. The results are summarised in Table 5.5, showing that the reductions in water consumption based upon this measure range from 86% in San Jose CA to 96% in Singapore.

Litres of water consumed per year, per kW of ITE Power	Scenario 1a	Scenario 1b	Scenario 2	Scenario 3
Air-Cooled	4,687	2,252	1,982	16,261
Liquid-Cooled	671	237	140	638
Reduction	86%	89%	93%	96%

Table 5.5: Percentage Reduction in Annual Water Consumption per kW of ITE Power

Taken alongside the reductions in Carbon Emissions per kW of ITE Power shown in Table 5.4, these figures show the very considerable reduction in environmental impact that can be achieved through the adoption of liquid cooling.

5.2 Conclusions

There are some general conclusions that can be drawn from the study:

- Liquid cooling technology is applicable in a variety of diverse environmental conditions. Warm cooling water temperatures allows its application in hot and wet climates.
- The Increased power density that liquid cooling can support, compared to most conventional air-cooled deployments, allows for reduced data halls space requirements or, alternatively, permits substantially increased ITE power to be supported within a data hall of the same size.
- The warm operating water temperatures typical of this technology allows the use simple heat rejection plant layouts and equipment combinations. Generally dry, or hybrid coolers would be sufficient to cool the FWS return. Evaporative cooling would only be needed under peak summertime conditions and by adopting ASHRAE Cooling Classes W4 or W5, the need for trim chillers is avoided.

- The improvement in rack-level energy efficiency, due to the omission of PSU cooling fans improves TUE, meaning that for a given ITE power supplied to the rack, a greater proportion is available to support the data processing equipment, than would be the case with an air-cooled server.
- The total cost of cooling per KW of ITE Power can be substantially reduced through the adoption of liquid cooling, with savings of up to 83% achieved in the 4 scenarios.
- Given the reduced need for evaporative cooling, annual water consumption per kW of ITE Power can also be dramatically reduced, with savings ranging from 86% to 96% in the 4 scenarios.
- Liquid cooling has a significant beneficial impact upon carbon emissions, as demonstrated in Tables 5.4 and 5.5. However, careful consideration must be given to selection of dielectric cooling fluids. Some fluorocarbon fluids have high GWP ratings and are chemically stable, so would cause lasting damage if allowed to evaporate to atmosphere. Some of the older oils that have been used in liquid cooling applications would be environmentally damaging, if allowed to leak into soil or water-causes.

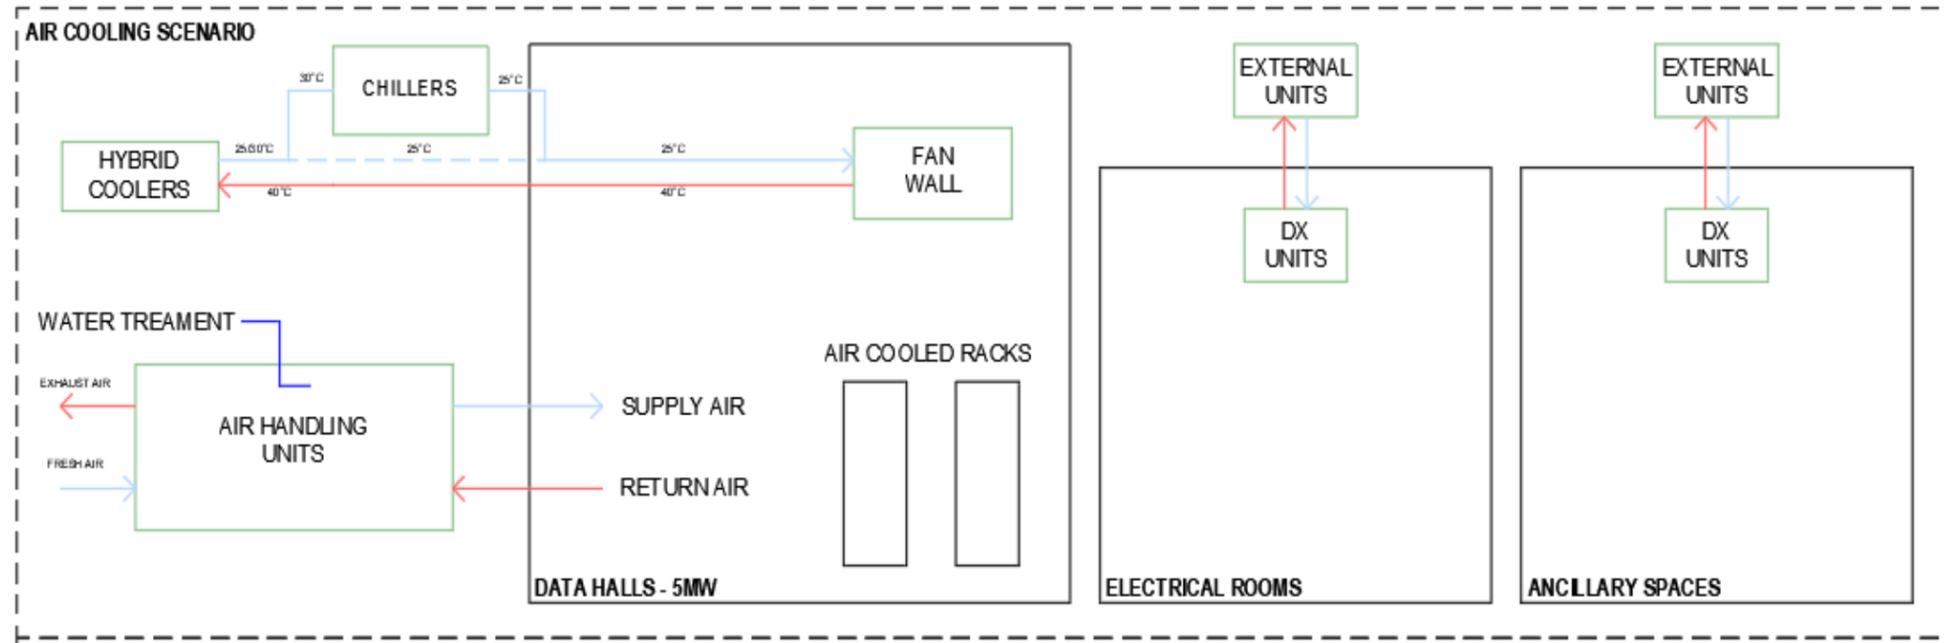
5.2.1 Future Proofing for Liquid Cooling.

The most evident future challenges facing the technology sector, as widely recognised by scientific and engineering communities, will be how to satisfy the increasing demand for data and connectivity, whilst at the same time mitigating the impact that data centres have upon the environment. Liquid cooling technology can be an effective and appealing response to those challenges:

- Liquid cooling can support the higher power densities being seen across the industry, but most especially in Edge and IoT applications.
- For Hyperscale and Cloud operators, liquid cooling offers a potential means of deploying large amounts of ITE power onto relatively small plots of land and in multi-storey buildings that do not need to accommodate large air handling plant on the roof or around the perimeter. The first generation of Hyperscale data centres, particularly those constructed in North America, have typically comprised groups of identical buildings on a common campus. The buildings are invariably of single-storey construction and are optimised for cost efficiency. They follow an established template design. However, a common feature of these campus facilities is that they occupy huge plots of development land. Suitable sites that meet all the requirements of a Hyperscale development; flat landscape, access to power, water, and drainage infrastructure, etc., are less readily available in Northern Europe and parts of APAC, and land costs are very much greater. In some locations, suitable sites of the scale required do not exist at all. This is forcing a re-evaluation of data centre designs, with multi-storey constructions starting to emerge. Liquid cooling facilitates this transition.
- Liquid cooling permits significantly more efficient use of energy and water, by virtue of lower PUE, TUE and WUE. The higher temperatures at which it can operate allow it to dispense with the need for expensive and energy intensive mechanical trim cooling. Evaporative cooling may only be required under peak summertime conditions. Higher operating temperatures also make liquid cooled facilities much more viable as a source of waste heat, that can be harvested and reused.

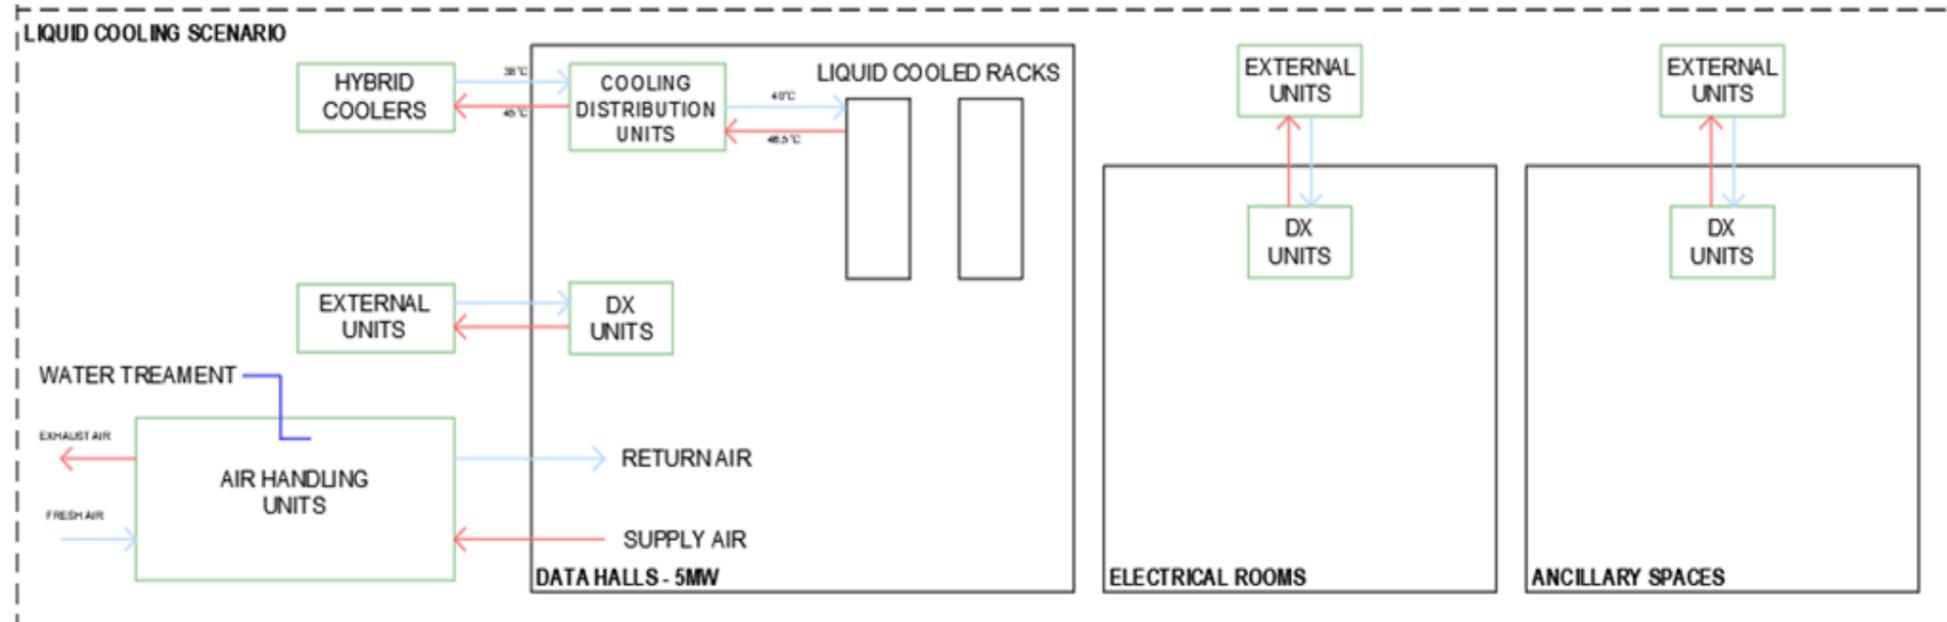
Appendices

Appendix A Scenario #1a – Richmond Virginia



Scenario #1a – Base Option Data Hall Schematic Arrangement

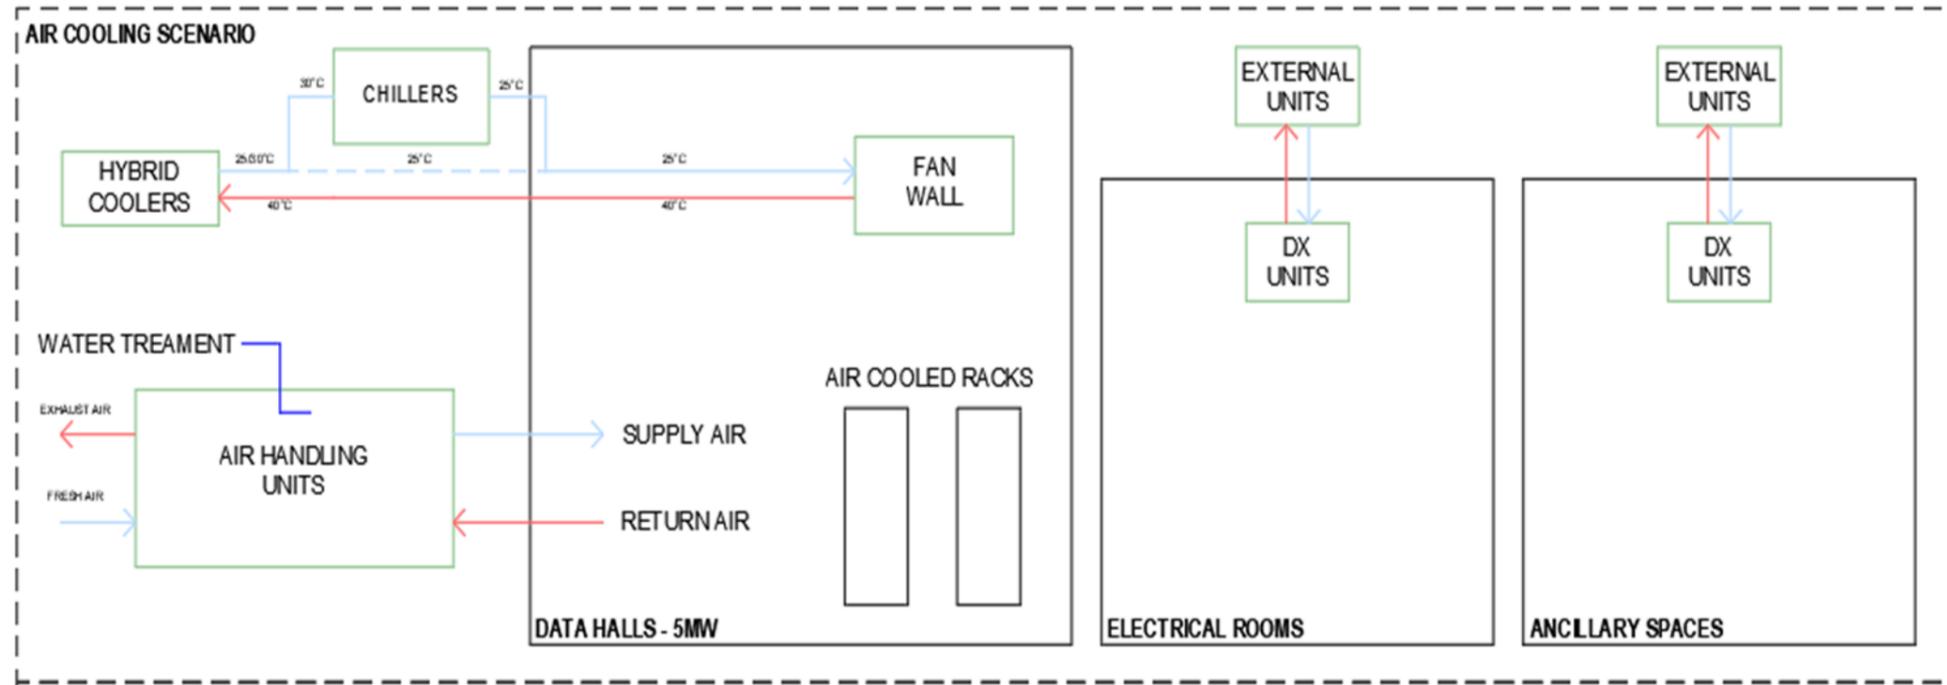
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Cooling	Chillers	44,491	0.7%
		Hybrid Coolers	314,132	5.2%
		Fan Coil Walls	2,200,512	36.5%
		RO System	62,481	1.0%
		Pumps	294,797	4.9%
Mechanical	Electrical Room Cooling	ER Cooling Units	105,231	1.7%
		Pressurisation AHU	212,624	3.5%
Electrical	Electrical Losses	UPS Losses (kWh)	1,576,800	26.2%
		LV Distribution losses (kWh)	660,983	11.0%
		MV to LV Transformer Losses (kWh)	403,573	6.7%
		Lighting	Lighting (kWh)	146,091
IT Load		Energy Demand [kWh]	39,420,000	
IT Fans		Energy Demand [kWh]	4380000	
Total Water Usage (litres)			21091582	
Total non IT load [kWh]			6021715	
Total IT load [kWh]			43800000	
Annualised PUE			1.14	
Mechanical pPUE			1.07	
Annualised WUE			0.482	
Annualised TUE			1.26	
Mechanical pTUE			1.19	
Annualised TWUE			0.535	



Scenario #1a – Schematic Arrangement of Proposed Liquid-Cooled Suite

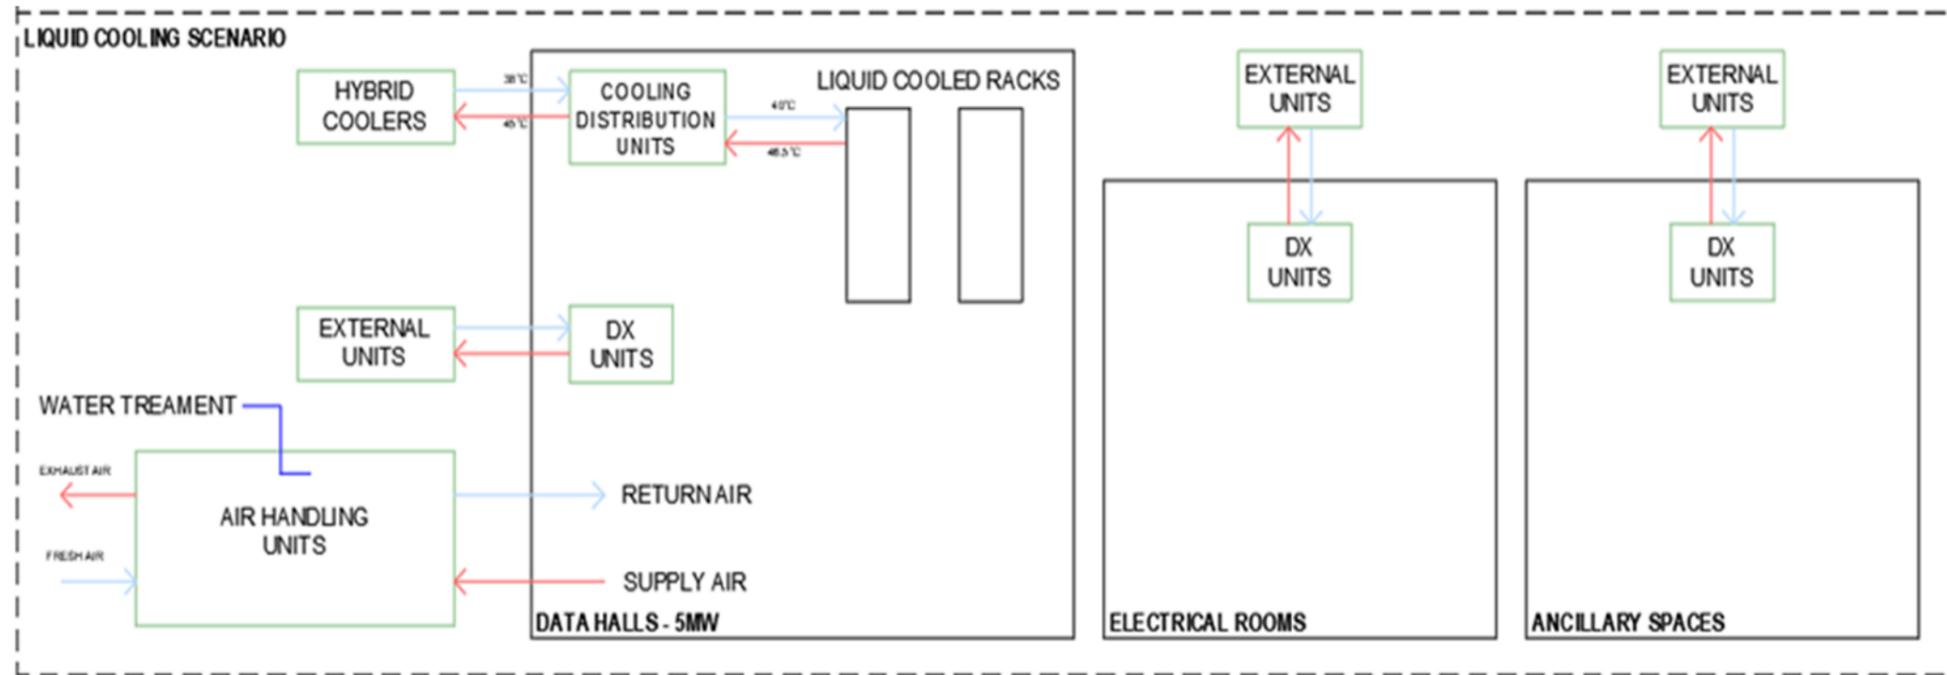
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Coolig	Chillers	0	0.0%
		Hybrid Coolers	228,154	5.7%
		Data Hall Pumps	158,863	4.0%
		RO Plant	8,972	0.2%
		Cooling Pumps	317,726	8.0%
Mechanical	Electrical Room Cooling	ER Cooling Units	107,569	2.7%
		Pressurisation AHU	212,624	5.4%
Electrical	Electrical Losses	UPS Losses (kWh)	1,699,440	42.8%
		LV Distribution losses (kWh)	676,541	17.0%
		MV to LV Transformer Losses (kWh)	414,443	10.4%
		Lighting	Lighting (kWh)	146,091
IT Load		Energy Demand [kWh]	42,486,000	
IT Pumps		Energy Demand [kWh]	1314000	
Total Water Usage (litres)			3255202	
Total non IT load [kWh]			3970423	
Total IT load [kWh]			43800000	
Annualised PUE			1.09	
Mechanical pPUE			1.02	
Annualised WUE			0.074	
Annualised TUE			1.12	
Mechanical pTUE			1.06	
Annualised TWUE			0.077	

Appendix B Scenario #1b – San Jose



Scenario #1b – Base Option Data Hall Schematic Arrangement

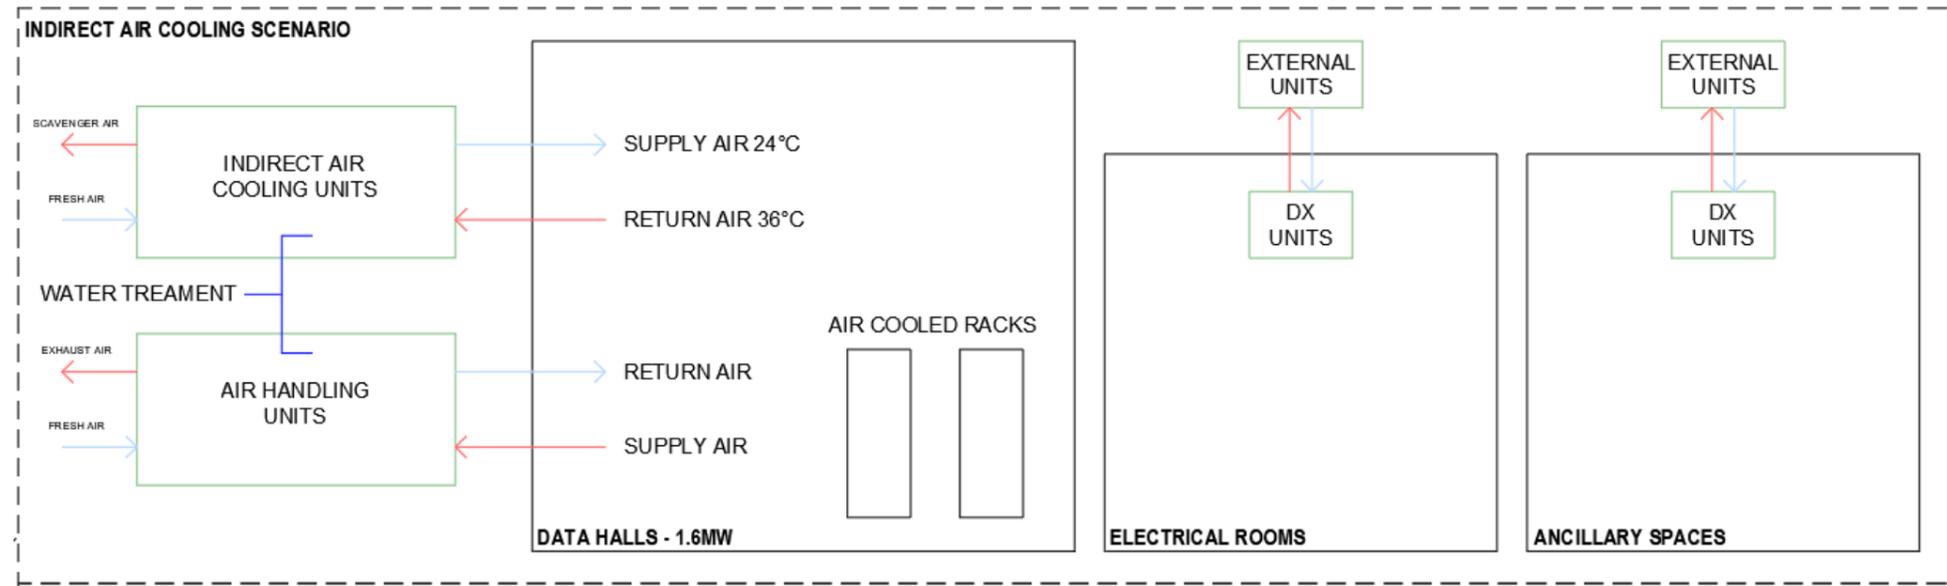
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Cooling	Chillers	86	0.0%
		Hybrid Coolers	386,572	6.5%
		Fan Coil Walls	2,200,512	37.0%
		RO System	29,880	0.5%
		Pumps	294,798	5.0%
Mechanical	Electrical Room Cooling	ER Cooling Units	87,843	1.5%
	Pressure/Humidity	Pressurisation AHU	157,889	2.7%
Electrical	Electrical Losses	UPS Losses (kWh)	1,576,800	26.5%
		LV Distribution losses (kWh)	661,003	11.1%
		MV to LV Transformer Losses (kWh)	403,581	6.8%
	Lighting	Lighting (kWh)	146,091	2.5%
IT Load		Energy Demand [kWh]	39,420,000	
IT Fans		Energy Demand [kWh]	4380000	
Total Water Usage (litres)				10136234
Total non IT load [kWh]				5945054
Total IT load [kWh]				43800000
Annualised PUE				1.14
Mechanical pPUE				1.07
Annualised WUE				0.231
Annualised TUE				1.26
Mechanical pTUE				1.19
Annualised TWUE				0.257



Scenario #1b – Schematic Arrangement of Proposed Liquid-Cooled Suite

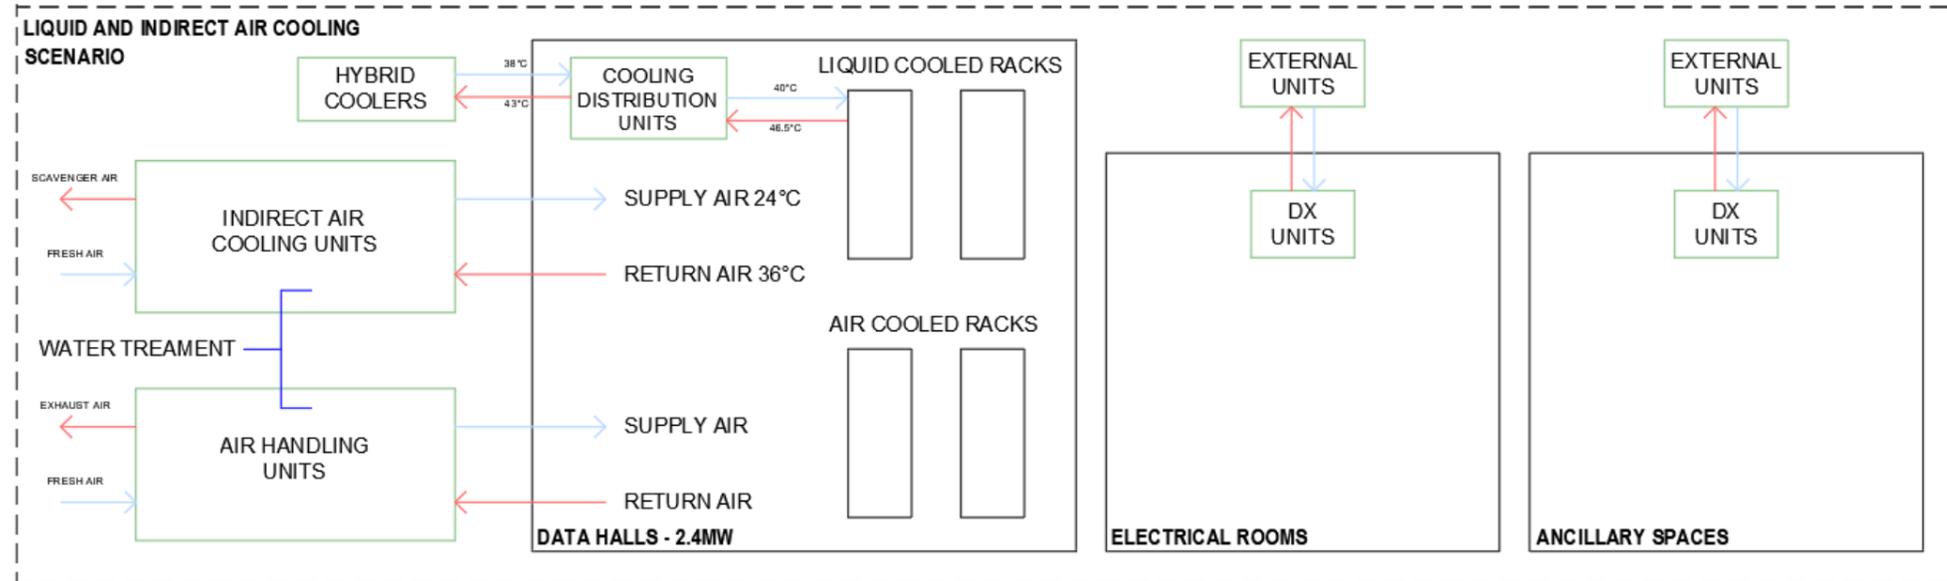
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Cooling	Chillers	0	0.0%
		Hybrid Coolers	198,012	5.1%
		Data Hall Pumps	158,863	4.1%
		RO Plant	2,921	0.1%
		Cooling Pumps	317,726	8.2%
Mechanical	Electrical Room Cooling	ER Cooling Units	85,619	2.2%
	Pressure/Humidity	Pressurisation AHU	157,889	4.1%
Electrical	Electrical Losses	UPS Losses (kWh)	1,699,440	44.1%
		LV Distribution losses (kWh)	675,705	17.5%
		MV to LV Transformer Losses (kWh)	413,851	10.7%
Lighting	Lighting (kWh)	146,091	3.8%	
IT Load		Energy Demand [kWh]	42,486,000	
IT Fans		Energy Demand [kWh]	1314000	
Total Water Usage (litres)				1149766
Total non IT load [kWh]				3856116
Total IT load [kWh]				43800000
Annualised PUE				1.09
Mechanical pPUE				1.02
Annualised WUE				0.026
Annualised TUE				1.12
Mechanical pTUE				1.05
Annualised TWUE				0.027

Appendix C Scenario #2 – London, UK



Scenario #2: Schematic Arrangement of 1.6MW Air-Cooled Data Hall, supported by IAC Units

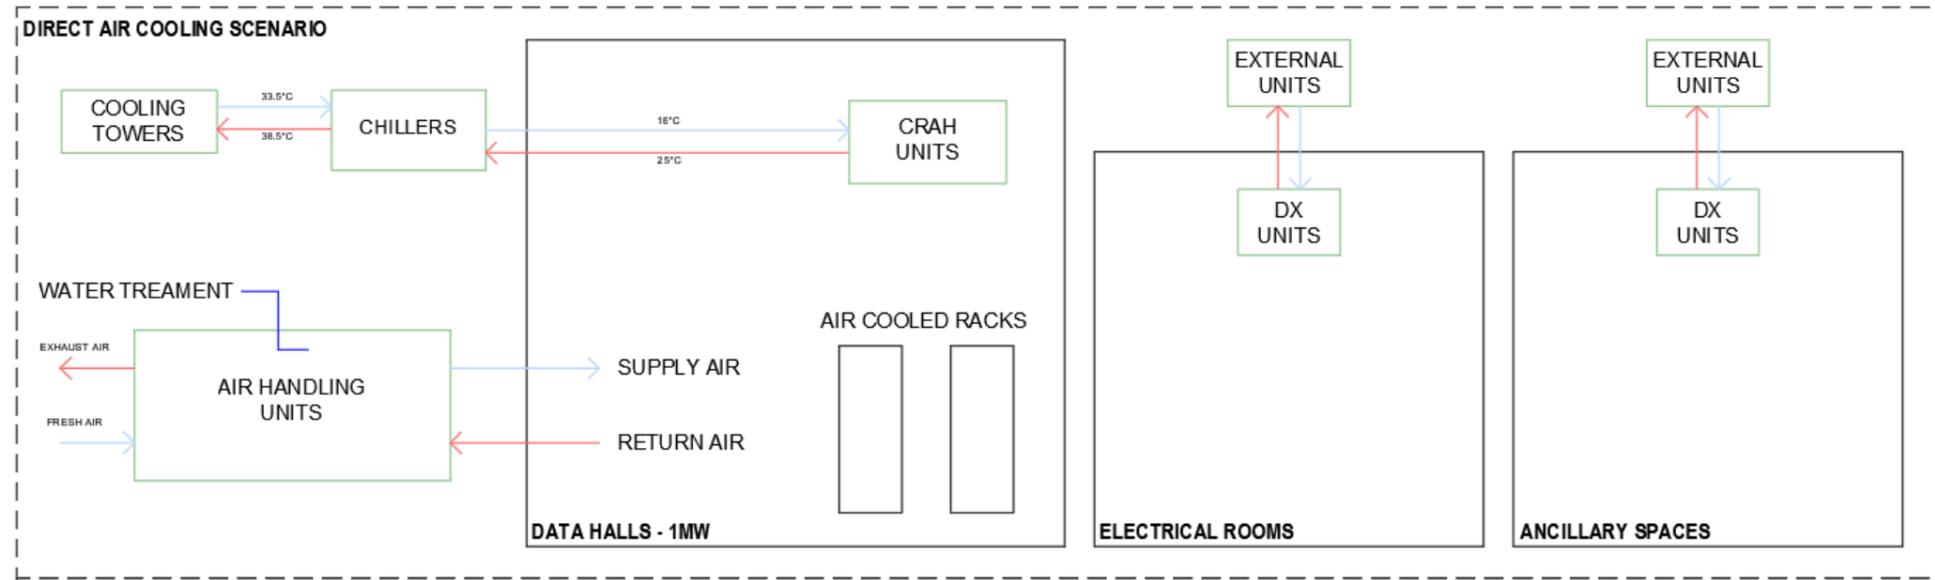
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Cooling			
		IAC	1,487,046	62.7%
Mechanical	Electrical Room Cooling	DX Cooling Units	75,288	3.2%
		Pressurisation AHU	17,522	0.7%
Electrical	Electrical Losses	UPS Losses	560,640	23.6%
		LV Switchgear losses	80,774	3.4%
		MV to LV Transformer Losses	78,493	3.3%
		Lighting	29,434	1.2%
		Small power	44,150	1.9%
IT Load		Energy Demand	14,016,000	
IT Fans		Energy Demand	1401600	
Total Water Usage (litres)				2854400
Total non IT load [kWh]				2373348
Total IT load [kWh]				15417600
Annualised PUE				1.15394
Mechanical pPUE				1.10
Annualised WUE				0.185
Annualised TUE				1.28
Mechanical pTUE				1.22
Annualised TWUE				0.206



Scenario #2: Schematic Arrangement of 2.4MW Mixed Air-Cooled and Liquid-Cooled Data Hall

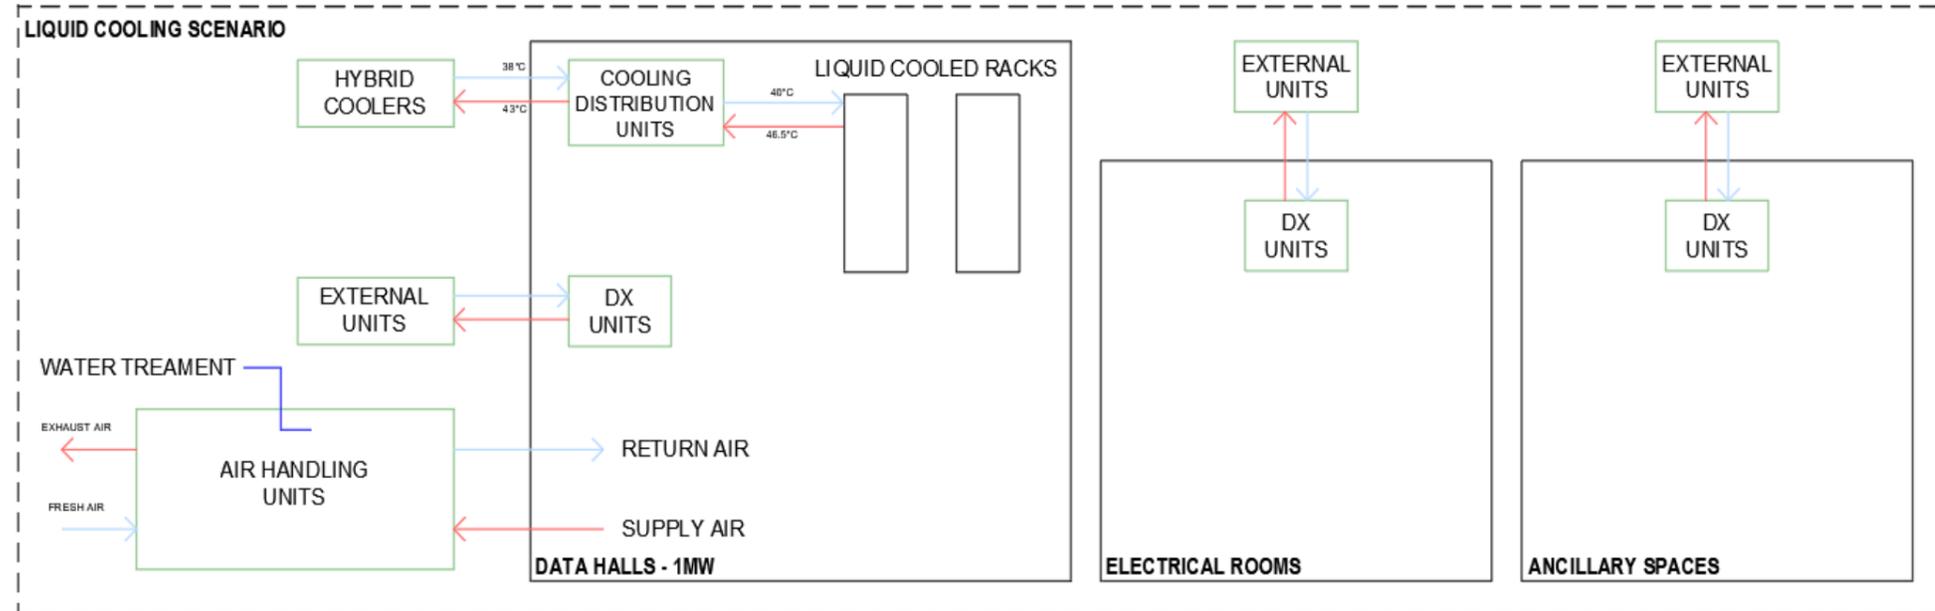
Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Cooling	Adiabatic Cooler	30,760	0.9%
		Pumps	392,023	12.0%
		IAC	1,509,352	46.3%
				0.0%
Mechanical	Electrical Room Cooling	DX Cooling Units	113,129	3.5%
		Pressurisation AHU	2,017	0.1%
Electrical	Electrical Losses	UPS Losses	840,960	25.8%
		LV Switchgear losses	119,363	3.7%
		MV to LV Transformer Losses	176,448	5.4%
		Lighting	29,434	0.9%
		Small Power	44,150	1.4%
IT Load		Energy Demand	21,024,000	
IT Pumps		Energy Demand	210240	
IT Fan		Energy Demand	1401600	
Total Water Usage (litres)				314267
Total non IT load [kWh]				3257636
Total IT load [kWh]				22635840
Annualised PUE				1.14391
Mechanical pPUE				1.09
Annualised WUE				0.014
Annualised TUE				1.24
Mechanical pTUE				1.18
Annualised TWUE				0.015

Appendix D Scenario #3 - Singapore



Schematic Arrangement of Air-Cooled 1.0MW Base Option, supported by Water-Cooled Chillers and CRAHs

Section	System	Subsystem	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Coolig	Chillers	3,523,272	65.8%
		Cooling Tower	70,992	1.3%
		CRAH	836,776	15.6%
		Pumps	296,391	5.5%
		DX Cooling Units	78,458	1.5%
Electrical	Pressure/Humidity	Pressurisation AHU	41,465	0.8%
	Electrical Losses	UPS Losses	350,400	6.5%
		LV Switchgear losses	52,088	1.0%
		MV to LV Transformer Losses	39,249	0.7%
	Lighting	Lighting	25,295	0.5%
	Small Power	Small Power	37,942	0.7%
IT Load	Energy Demand	8,760,000		
IT Pumps	Energy Demand	876000		
Total Water Usage (litres)			14635165	
Total non IT load [kWh]			5352326	
Total IT load [kWh]			9636000	
Annualised PUE			1.56	
Mechanical pPUE			1.50	
Annualised WUE			1.519	
Annualised TUE			1.73	
Mechanical pTUE			1.67	
Annualised TWUE			1.688	



Schematic Arrangement of Liquid-Cooled 1.0MW Option, supported by Hybrid-Coolers.

Section	System	Sub-System	Yearly Energy Demand [kWh]	Fraction against total Consumption
Mechanical	Data Hall Coolig	Hybrid Coolers	93,531	9%
		DX System	48,831	5%
		Pumps	300,137	29%
		DX Cooling Units	76,666	8%
Electrical	Pressure/Humidity	Pressurisation AHU	30,660	3%
	Electrical Losses	UPS Losses	350,400	34%
		LV Switchgear losses	49,817	5%
		MV to LV Transformer Losses	39,249	4%
	Lighting	Lighting	11,826	1%
Small Power	Small Power	17,739	2%	
IT Load	Energy Demand	8,497,200		
IT Pumps	Energy Demand	262800		
Total Water Usage (litres)			618836	
Total non IT load [kWh]			1018856	
Total IT load [kWh]			8760000	
Annualised PUE			1.12	
Mechanical pPUE			1.06	
Annualised WUE			0.071	
Annualised TUE			1.15	
Mechanical pTUE			1.10	
Annualised TWUE			0.073	

