



Edge Insights Service

EDGE SUSTAINABILITY: NAVIGATING STRATEGIES FOR SUCCESS

Edge computing presents unique challenges and limitations when compared to traditional data centres. This report provides practical guidance for edge operators aiming to achieve sustainability goals.



Executive Summary

While sustainability is a strategic priority for many organisations including telcos, it can be challenging to decide on the optimal approach for meeting sustainability targets, especially in emerging technology domains. Edge computing is projected to grow from a total global market size of US\$9 billion in 2020 to over US\$450 billion in 2030, at a CAGR of 49% over the 10-year period¹. It is critical that this rapid growth is developed in a manner as sustainable as possible, arguably more so than the traditional centralised cloud. In addition to pursuing sustainability strategies that have been successful in traditional data centre environments, edge operators should look to incorporate advanced edge-specific sustainability strategies, which could eventually result in annual energy savings of up to 5%, equating to ~US\$40 million in total energy costs and over 1m tCO₂eq of emissions for a large group operator (see Figure 12 and Figure 13 for a breakdown of calculations). However, there are no one-size-fits-all answers and edge operators must select the optimal advanced sustainability strategy based on their unique context. This report examines the multitude of factors that need be considered when defining your edge sustainability strategy and provides guidance based on common scenarios.

Edge computing: The next frontier of sustainability strategies

While traditional data centre sustainability has been widely covered, edge computing as a less mature technology has received less attention. Edge brings with it unique challenges and restrictions, meaning the sustainability strategies pursued must be carefully considered.

Figure 1: Evaluation of advanced edge sustainability strategy sub-categories

Edge sustainability factors	Specialised cooling applied to COTS hardware					Specialised hardware
	Precision liquid cooling		Tank/tub immersion cooling (e.g. Submer MicroPod)	Direct to chip water cooling (e.g. CoolIT RX4)	Air-conditioned cooling cabinets (e.g. EDGE 5 Micro Data Centre)	Ruggedised servers (e.g. Gecco)
	Liquid to air (e.g. Iceotope KUL RAN)	Liquid to liquid (e.g. Iceotope KUL 2)				
Legacy (NDC) site reuse	●	●	●	●	●	●
Energy use	●	●	●	●	●	●
Water consumption	●	●	●	●	●	●
Longevity of equipment	●	●	●	●	●	●
Limited maintenance	●	●	●	●	●	●
Site adaptation	●	●	●	●	●	●
COTS hardware adaptation	●	●	●	●	●	●

● Addresses factor well
● Adequately addresses factor
● Insufficiently addresses factor

Source: STL Partners

¹ STL Partners: <https://stlpartners.com/research/edge-computing-market-sizing-forecast-third-release/>

For example, many edge deployments are remote and in non-data centre (NDC) facilities which lack features such as advanced heating, ventilation and air-conditioning (HVAC) systems, uninterrupted power supply (UPS), and continuous support staff due to a lack of cost-effectiveness.

Edge operators should start by adopting the sustainability strategies widely employed in traditional data centre environments today (see Figure 4 for more detail), which are complementary, so can be implemented simultaneously. Edge operators seeking to progress their sustainability efforts should also look to adopt advanced edge sustainability strategies of which this report outlines two major categories: specialised cooling applied to common-off-the-shelf (COTS) hardware and specialised hardware. To guide decision-making, we have ranked these strategies based on their effectiveness in addressing key edge sustainability factors (see Figure 1 and Figure 10).

Importantly, edge deployments vary considerably, as do the needs of edge operators, necessitating a tailored approach to select the optimal advanced edge sustainability strategy. Providing a one-size-fits-all solution is impractical. This report provides high-level guidance to identify the most suitable solutions for the common scenarios encountered in edge computing (see Figure 11).

Table of Contents

Executive Summary.....	2
Edge computing: The next frontier of sustainability strategies.....	2
The edge compute sustainability challenge	6
Sustainability is a strategic priority for cloud computing and should be prioritised at the edge	6
Edge sustainability is overlooked.....	6
Edge versus cloud sustainability.....	7
The dark truth about PUE and why this matters more for edge.....	8
Start by pursuing 'baseline' sustainable edge strategies.....	9
Beyond the baseline: Advanced sustainability strategies for edge environments	10
Which sustainability strategies are suitable when?	15
How big an impact can the right strategy have? An example from the telco world.....	16

Table of Figures

Figure 1: Evaluation of advanced edge sustainability strategy sub-categories	2
Figure 2: Edge computing environments often lack traditional data centre features.....	7
Figure 3: Some types of edge are more likely to be basic, remote and NDC-like than others.....	8
Figure 4: Traditional data centre sustainability strategies can be implemented at the edge.....	9
Figure 5: Advanced edge sustainability strategies	10
Figure 6: Overview of precision liquid cooling	11
Figure 7: Overview of tank/tub immersion cooling.....	12
Figure 8: Overview of direct to chip water cooling.....	12
Figure 9: Overview of air-conditioned cooling cabinets	13
Figure 10: Evaluation of advanced edge sustainability strategies.....	14
Figure 11: Which sustainable edge strategies are suitable when?.....	15
Figure 12: Operator energy consumption breakdown.....	16
Figure 13: Potential annual savings for an international group operator	17

The edge compute sustainability challenge

Sustainability is a strategic priority for cloud computing and should be prioritised at the edge

Data centre sustainability is becoming increasingly critical in the face of climate change and rising global energy demands. Data centres already draw approximately 1.5% of the world's electricity², a figure set to expand in tandem with the growth of cloud and edge computing. Future projections underscore the scale of the situation: by 2030, data centres could be responsible for consuming 5-10% of total electricity in some countries³.

The imperative for data centre sustainability extends beyond environmental stewardship – it's also strategic concern. Providers are increasingly held accountable for their carbon footprint, with customers, investors, and regulatory bodies demanding transparency and action. This shift places data centre efficiency and sustainability at the forefront of corporate strategy, emphasising not just cost savings but also corporate responsibility and long-term viability.

Edge sustainability is overlooked

There is much less coverage on the topic of edge compute sustainability than there is on traditional data centre sustainability. There is an expectation of significant growth in edge compute. STL Partners has projected that the total global market size will grow from US\$9 billion in 2020 to over US\$450 billion in 2030, at a CAGR of 49% over the 10-year period. It is critical that this growth is developed in a manner as sustainable as possible, arguably more so than the traditional centralised cloud. While some of the sustainable data centre strategies also apply at the edge, the nature of certain edge environments means that these strategies become less effective or even unviable. More edge-specific strategies are needed.

This picture is further complicated by the desire of edge operators to simplify their deployment and management of distributed compute, necessitating standardised, highly reliable and low-maintenance infrastructure. Finally, every edge operator will have a unique situation, reflecting their existing edge location infrastructure, operating environment, and resources. Therefore, edge operators looking to improve the sustainability of their operations must develop sustainable strategies that are specific to them and their unique requirements. Which sustainable edge strategies are suitable and when should they be used?

² <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>

³ <https://www.bloomberg.com/news/articles/2021-10-14/europe-s-data-centers-will-gobble-up-a-lot-more-electricity?leadSource=uverify%20wall>

Edge versus cloud sustainability

How do we define edge computing?

Edge computing can be defined as the processing of enterprise and/or telco network workloads at the (physical and topological) edge of the network as opposed to in a centralised data centre or the cloud. Edge computing emphasises the importance of handling data and applications closer to their source, with a view to reducing latency and improving efficiency in data transmission and processing. It also ensures data sovereignty and different levels of security. Different types of edge computing can be distinguished by the physical or logical location in which workloads are processed. This report focuses on three main types of edge computing:

- **On-premises edge computing:** On-premise edge computing refers to computing resources that reside at the customer site, for example on an IoT gateway physically on-site, an on-premises data centre, etc.
- **Network edge (mobile and access network):** Network edge refers to edge compute locations that are at sites or points of presence (PoPs) run by a telecoms operator, for example at a central office in the mobile network or at an ISP's node.
- **Regional data centre:** Regional edge refers to small carrier-neutral data centres or internet exchanges, often located near tier two and tier three cities.

Edge computing environments typically differ from traditional data centre environments due to their remote and distributed characteristics. Implementing standard data centre features, such as advanced HVAC systems, uninterruptible power systems (UPS), and continuous support staff, is not typically cost-effective in these settings due to a lack of economies of scale. Figure 2 describes some of the key differences between a data centre and non-data centre environment.

Figure 2: Edge computing environments often lack traditional data centre features

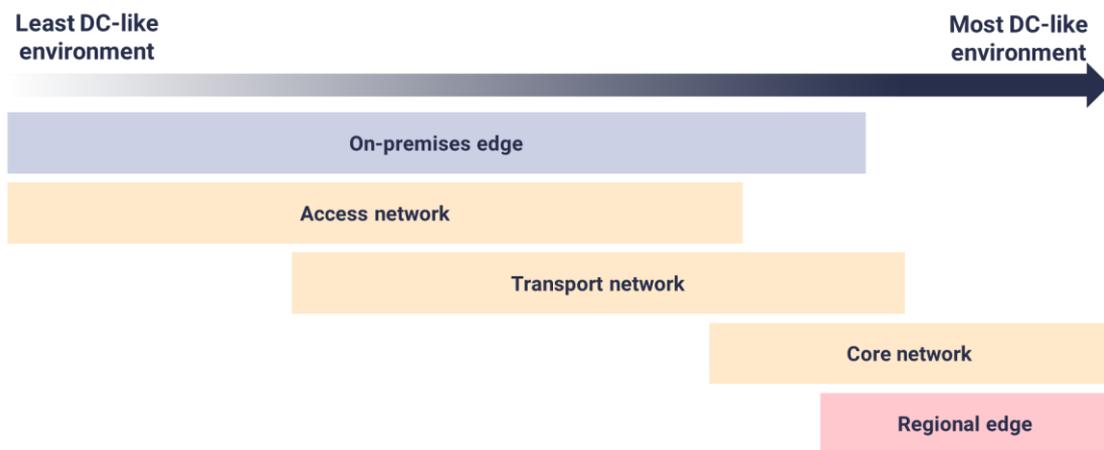
Non data centre (NDC) edge environment	Data centre (DC) environment
Often using repurposed legacy spaces	Purpose-built data centres
No data-centre-specific HVAC/UPS	Advanced HVAC/UPS
Remote sites (truck roll required, complicated access, e.g. permission for site access)	Centralised sites
No permanent support staff	24/7/365 support staff

Source: STL Partners

Edge computing environments vary in their divergence from traditional data centres, with some types more different (from data centres) than others. The closer the edge location is to the data source, or the further out the edge extends, the more it tends to diverge from standard data centre characteristics.

However, there is still significant variability within the same type of edge, including differences in infrastructure and support. For example, on-premises edge computing could involve processing data in a remote, environmentally harsh setting like a wind turbine, or it might occur in the basement of a large corporate headquarters with 24/7 support staff and a concentration of servers. Figure 3 shows this spectrum of edge computing environments, emphasising the variety across the different types.

Figure 3: Some types of edge are more likely to be basic, remote and NDC-like than others



Source: STL Partners

The dark truth about PUE and why this matters more for edge

How you measure the effectiveness of sustainability strategies is critical. Power usage effectiveness (PUE) is a commonly adopted metric to measure the energy efficiency of data centre operations today. PUE is the ratio of the amount of power required to run the entire data centre compared to the power required to run the IT equipment.

$$\text{PUE} = \text{Total power consumed by a data centre} / \text{Total power consumed by the IT equipment}$$

A PUE ratio of 1.00 (which in reality is impossible to achieve) would be a perfectly efficient operation, in which all the power is consumed by the IT equipment and nothing else (e.g. UPS, cooling, transmission losses, etc.). However, PUE can be misleading as it doesn't accurately represent the efficiency of energy consumed by the actual silicon. For example, server fans consume power that counts as 'consumed by IT equipment'. This means that having less air conditioning, resulting in server fans running at maximum speed, could result in a lower PUE ratio for the site, but the overall power consumed to deliver the same useful workloads could actually be higher. This issue is particularly important at the edge, where 'passive' cooling (without dedicated air conditioning) is more common.

Furthermore, having large volumes of dusty, humid air pushed through servers reduces their lifespan and increases the need for repairs, which further impacts their sustainability.

A better metric is total usage effectiveness (TUE). TUE incorporates IT energy usage effectiveness (ITUE), which evaluates the efficiency of the IT equipment's energy use in relation to its computational output. As such, it accounts for the impact of support components, such as cooling fans and power supply units.

$$\text{TUE} = \text{PUE} \times \text{ITUE}$$

$$\text{ITUE} = \text{Total energy consumed by the IT equipment} / \text{Total energy consumed by the compute components of IT equipment (e.g. the CPU)}$$

By seeking to optimise TUE, edge operators can secure better economic and environmental outcomes. This will lead operators to pursue more advanced sustainability strategies at the edge.

Start by pursuing 'baseline' sustainable edge strategies

Sustainability strategies frequently employed in traditional data centres can also be effectively adapted for edge computing environments. Figure 4 outlines the advantages of these common sustainability strategies, as well as the potential obstacles they might face in edge computing deployments. These strategies are designed to be complementary, so they can be implemented simultaneously. However, the distinctive characteristics of certain edge environments pose unique challenges, which can reduce the effectiveness or feasibility of these strategies. For instance, the heterogenous nature of edge devices makes it challenging to develop lean code that is tailored to each device running in unique environments, thereby reducing the practicality of green code as a sustainability measure.

Figure 4: Traditional data centre sustainability strategies can be implemented at the edge

Strategy	Rationale	Challenges
Renewable energy use	<ul style="list-style-type: none"> Significantly lowers carbon emissions associated with edge (IT, cooling, general operations) Long-term cost efficiency 	<ul style="list-style-type: none"> Inconsistent energy supply due to variability in energy generation, requiring robust back-up systems
Green code	<ul style="list-style-type: none"> Writing code that requires less computational power reduces energy consumption of servers, reducing the need for cooling systems Code that maximises silicon's capability, e.g. sleep modes 	<ul style="list-style-type: none"> Creating lean, reliable, and resilient code for the varied and remote operations of edge computing is challenging due to the need to maintain energy efficiency across heterogeneous devices
Modularity of IT equipment	<ul style="list-style-type: none"> Individual components can be upgraded or replaced without needing to change the entire system reducing resource consumption Adaptability/ future proofing and longevity 	<ul style="list-style-type: none"> Modular systems can be more expensive upfront Managing and upgrading individual modules can be logistically challenging in dispersed edge locations
Waste energy reuse	<ul style="list-style-type: none"> The waste heat generated by servers can be repurposed, turning a byproduct into a resource and reducing overall energy demand 	<ul style="list-style-type: none"> Smaller scale of operations makes it less economically viable For remote edge deployments there can be a lack of nearby utilisation options for waste energy
Passive air cooling	<ul style="list-style-type: none"> Only makes sense in an edge environment Avoids the high energy use associated with traditional air conditioning 	<ul style="list-style-type: none"> Need to manage heat dissipation across various devices in uncontrolled environments without relying on active cooling methods

Source: STL Partners

Beyond the baseline: Advanced sustainability strategies for edge environments

For edge operators looking to advance sustainability beyond baseline strategies, there are broadly two categories of options, described in Figure 5. More efficient cooling, in particular, is effective at reducing the energy consumption of edge operations, as cooling systems can account for up to 40% of the energy used by a typical data centre to ensure optimal server performance⁴.

Figure 5: Advanced edge sustainability strategies

	Specialised cooling applied to COTS hardware	Specialised, ruggedised hardware
Description	Using liquid/gas to remove heat from server components, which is more efficient than traditional fan-based air cooling, especially in dense computing environments	Using hardware specifically designed for edge computing, which can be more energy efficient and better suited for edge environments than servers intended for data centre deployment
Rationale	<ul style="list-style-type: none"> Reduces the cooling energy requirements Particularly effective in high-temperature or cramped edge locations 	<ul style="list-style-type: none"> Processes tasks more efficiently, reducing energy consumption and heat generation Ruggedised equipment can withstand harsh environmental conditions
Challenges	<ul style="list-style-type: none"> Higher initial installation costs Potential need for regular maintenance to prevent leaks Complexity of retrofitting existing edge locations with liquid cooling systems for some solutions 	<ul style="list-style-type: none"> Higher cost of servers Stock keeping unit (SKU) increase: <ul style="list-style-type: none"> Complexity of managing diverse inventory Operational complexity and cost Doesn't lend itself well to scaled edge deployment

Source: STL Partners

⁴ <https://dataspan.com/blog/how-much-energy-do-data-centers-use/#:~:text=For%20many%20data%20centers%2C%20the,of%20the%20center's%20total%20power.>

These categories of edge sustainability strategies can be further broken down into the subcategories presented below.

Subcategories of specialised cooling applied to COTS hardware

We have identified five subcategories of specialised cooling technologies.

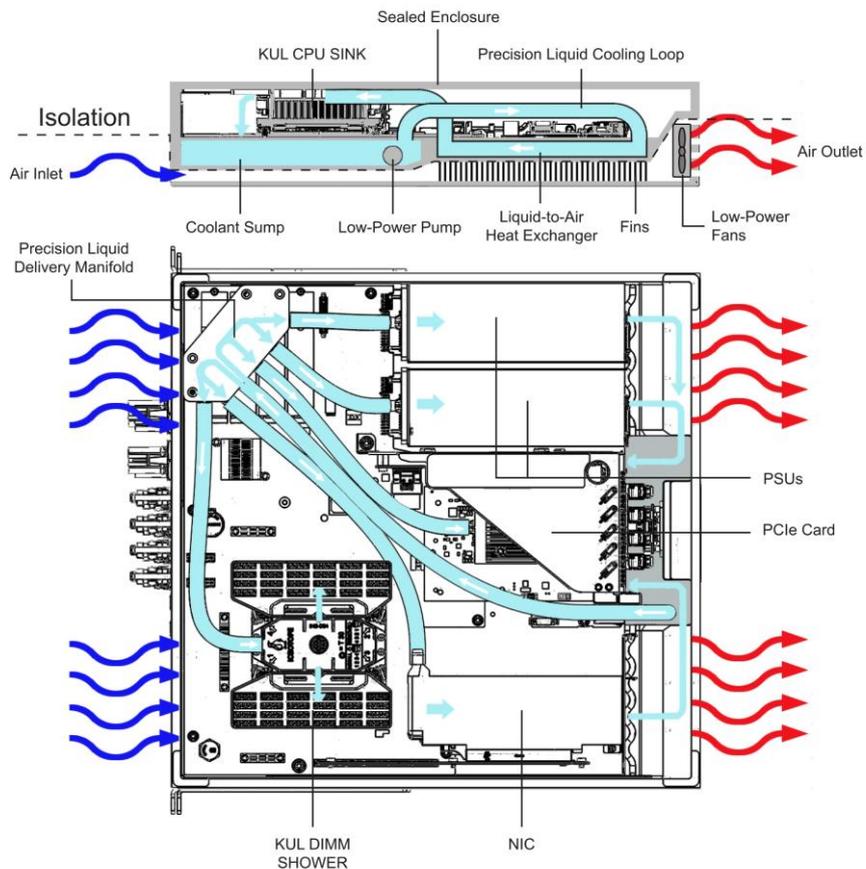
Precision liquid cooling: Liquid to air

Precision liquid cooling is a method that enhances performance, efficiency, and reliability in server cooling. It precisely delivers a minimal amount of dielectric coolant directly to the hottest components of a server first before cascading into the chassis to cool all other components. This technique eliminates the requirement for traditional air-cooling systems, enabling more flexible IT solution designs. It effectively prevents hotspots that could hinder performance, reduces the need for large cooling infrastructure, and eliminates water usage. In liquid-to-air precision liquid cooling, the liquid (which does not conduct electricity) collects the heat and then transfers the fluid to an in-chassis, in-rack, or in-row heat exchanger, where the heat is dissipated into the air (see Figure 6).

Precision liquid cooling: Liquid to liquid

Liquid-to-liquid precision liquid cooling works on the same basis as liquid to air in that the liquid coolant absorbs heat from the server components. However, instead of transferring this heat to air, it is transferred to a facility water loop where the heat is rejected using dry coolers (see Figure 6).

Figure 6: Overview of precision liquid cooling

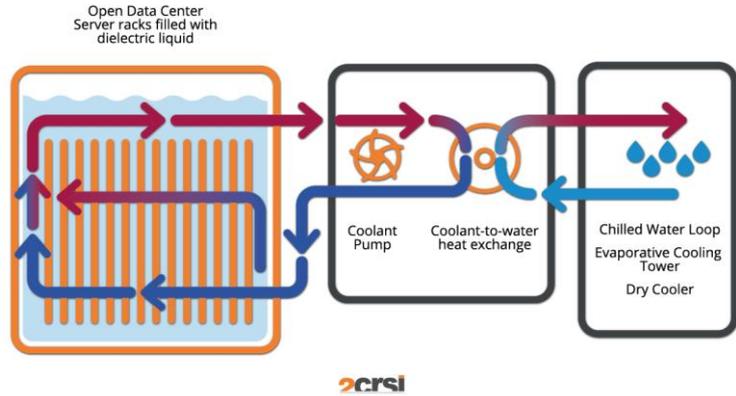


Source: Iceotope

Tank/tub immersion cooling

Single-phase immersion cooling involves submerging electronic components in a non-conductive liquid that remains liquid at all times. The coolant directly absorbs heat from the components and is then pumped to a heat exchanger. At this stage, the heat can be dissipated in two ways: either through a water-cooling circuit or by using a fan to disperse the heat into the surrounding air (Figure 7 depicts water-cooling circuit heat dissipation). However, there are three main drawbacks to this system: the special liquid used is expensive, it needs to be handled carefully due to its hazardous properties, and it is not easy to accommodate in legacy spaces.

Figure 7: Overview of tank/tub immersion cooling

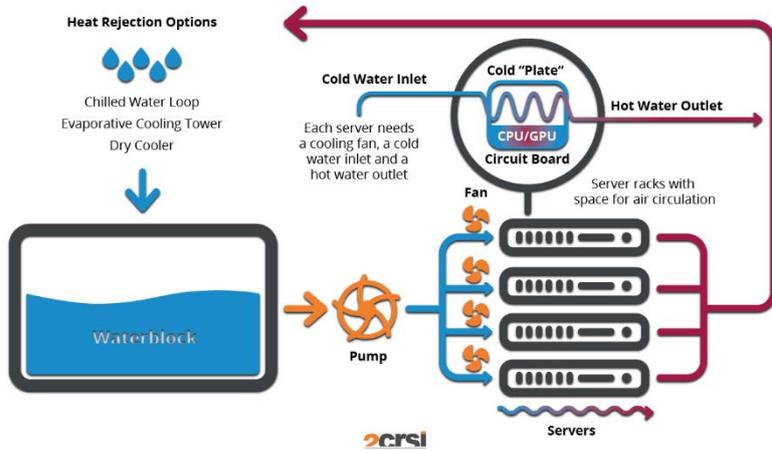


Source: 2CRSI

Direct to chip water cooling

A direct-to-chip system works by circulating a liquid, usually water, directly over the heat-generating components of the server. A cold plate, typically made of a highly conductive material like copper, is attached to these components. This plate has a special structure with channels or tubes for the coolant to flow through. As the coolant moves through these channels, it picks up heat from the server components, effectively cooling them. The now heated coolant leaves the cold plate and is either cooled down again in an external heat exchanger or removed from the system (see Figure 8).

Figure 8: Overview of direct to chip water cooling

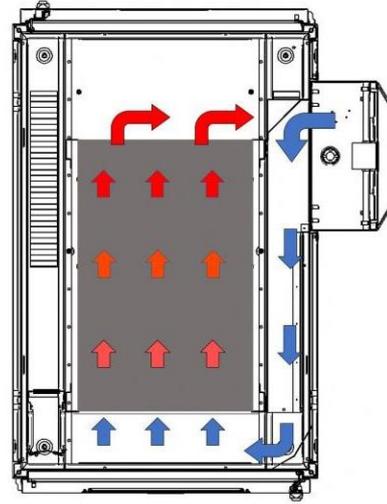


Source: 2CRSI

Air-conditioned cooling cabinets

An air-conditioned cooling cabinet for servers effectively acts as a self-contained, micro-data centre. It operates by enclosing the equipment and using an air conditioning unit. The system cools the air within the enclosure by removing the heat generated by the equipment. It recirculates the cooled air back over the equipment, maintaining a consistent temperature without the need for external air, which helps to keep the equipment at optimal operating temperatures (see Figure 9).

Figure 9: Overview of air-conditioned cooling cabinets



Source: EDPeurope

Specialised hardware: ruggedised servers

Ruggedised servers specialised for edge computing are designed to withstand harsh environmental conditions. They are compact and energy efficient, making them suitable for remote and challenging edge locations where standard servers may not be practical. These servers are built to be durable and reliable, which ensures consistent performance in rugged environments.

We set out seven criteria for comparing advanced edge strategies

The solutions outlined above have different advantages depending on the specific characteristics of the edge environment. To guide decision making, we have ranked these solutions based on their effectiveness in addressing key edge sustainability factors (see Figure 10). These sustainability considerations are particularly relevant to edge computing deployments and play a crucial role in determining the overall environmental impact and efficiency of these solutions.

Figure 10: Evaluation of advanced edge sustainability strategies

Edge sustainability factors	Specialised cooling applied to COTS hardware					Specialised hardware
	Precision liquid cooling		Tank/tub immersion cooling (e.g. Submer MicroPod)	Direct to chip water cooling (e.g. CoolIT RX4)	Air-conditioned cooling cabinets (e.g. EDGE 5 Micro Data Centre)	Ruggedised servers (e.g. Gecco)
	Liquid to air (e.g. Iceotope KUL RAN)	Liquid to liquid (e.g. Iceotope KUL2)				
Legacy (NDC) site reuse	●	●	●	●	●	●
Energy use	●	●	●	●	●	●
Water consumption	●	●	●	●	●	●
Longevity of equipment	●	●	●	●	●	●
Limited maintenance	●	●	●	●	●	●
Site adaptation	●	●	●	●	●	●
COTS hardware adaptation	●	●	●	●	●	●

● Addresses factor well
● Adequately addresses factor
● Insufficiently addresses factor

Source: STL Partners

The criteria used for this ranking are:

- Legacy (NDC) site reuse: Using pre-existing buildings, which were not originally designed to house IT infrastructure, for edge computing deployments – for example repurposing old office buildings or telecommunications switch centres as locations for servers and other computing equipment. Reusing existing structures can reduce the environmental impact compared to constructing new, purpose-built facilities by minimising the resource consumption and waste generation associated with building construction.
- Energy use: Pertains to the amount and efficiency of power consumption in edge computing environments. Efficient energy use not only reduces environmental impact but can also lead to cost savings over time.
- Water consumption: The amount of water consumed and the efficiency of water utilisation in edge computing facilities. Water is often used in cooling systems to manage the temperature of computing equipment, which is critical for maintaining performance and longevity. In terms of sustainability, the goal is to minimise water consumption to reduce the environmental footprint edge operation.

- Longevity of equipment: Equipment with a longer lifespan reduces the need for manufacturing, shipping, and disposal. This cuts down on waste, raw material use, and energy for production and recycling, thereby decreasing the scope 3 emissions linked to the supply chain.
- Limited maintenance: Equipment requiring less maintenance decreases the frequency of technician visits, or 'truck rolls', which reduces scope 1 emissions from the company's vehicle use. This benefit is especially significant at remote edge sites.
- Site adaptation: The ability to operate without installing new infrastructure, like air conditioning, specialised power systems, or additional construction, reduces embedded emissions.
- COTS hardware adaptation: Ideally, advanced sustainability technologies should allow operators to easily use any COTS server with the solution. This includes the reuse of existing servers, thereby extending their life cycle and reducing electronic waste. Adapting servers (involving disassembly, modifications and re-assembly) increases costs, plus some aspects of inventory management.

Which sustainability strategies are suitable when?

The variety and complexity of edge computing environments mean that the needs of edge operators differ greatly. This calls for a tailored approach to meet specific business and operational requirements in each unique edge setting. Providing a one-size-fits-all solution is impractical. Figure 11 provides high-level guidance to identify the most suitable solutions for the common scenarios encountered in edge computing.

Figure 11: Which sustainable edge strategies are suitable when?

Common edge computing scenarios	Specialised cooling applied to COTS hardware					Specialised HW
	Precision liquid cooling		Tank/tub immersion cooling (e.g. Submer MicroPod)	Direct to chip water cooling (e.g. CoolIT RX4)	Air-conditioned cooling cabinets (e.g. EDGE 5 Micro Data Centre)	Ruggedised servers (e.g. Gecco)
	Liquid to air (e.g. Iceotope KUL RAN)	Liquid to liquid (e.g. Iceotope KUL 2)				
Scenario 1 (extreme edge with limited/no enclos. in harsh environments)	✓✓	✓		✓	✓	✓✓✓
Scenario 2 (on-premise edge at satellite office)	✓✓✓	✓✓	✓	✓	✓	✓✓
Scenario 3 (Access network edge e.g. RAN)	✓✓✓	✓✓	✓✓	✓	✓	✓
Scenario 4 (Transport network)	✓✓	✓✓✓	✓✓	✓✓	✓✓	✓
Scenario 5 (regional edge data centre)		✓✓✓		✓✓	✓✓	

✓✓✓ Optimal solution in given scenario
 ✓✓ Satisfactory solution in given scenario
 ✓ Sub-optimal solution in given scenario

Source: STL Partners

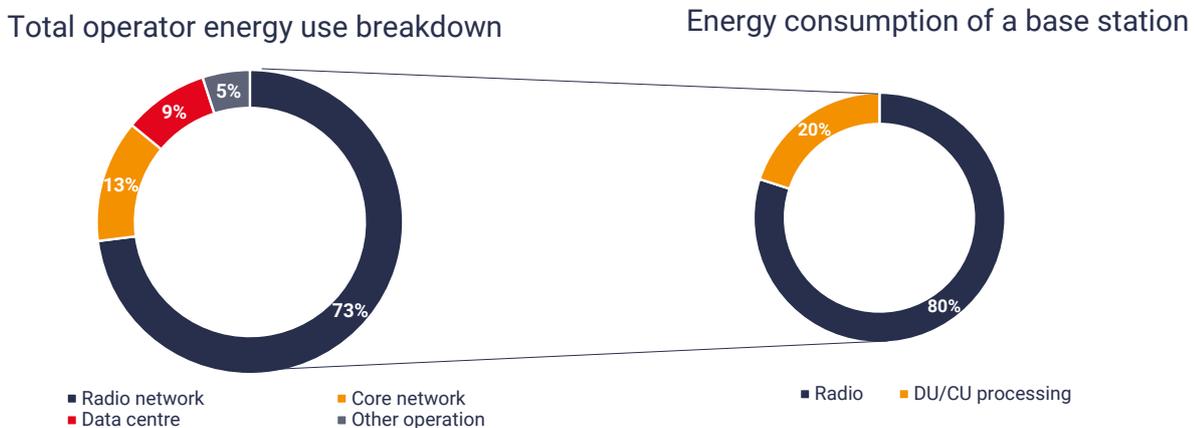
The edge computing scenarios shown above are:

- Scenario 1 (Extreme edge with limited/no enclosure in harsh environments): For example, processing units mounted directly on top of an offshore wind turbine.
- Scenario 2 (On-premise edge at satellite office): A local retail store may have its own servers on site to handle inventory management and sales transactions quickly.
- Scenario 3 (Access network edge, e.g. RAN): Edge computing resources deployed for RAN (CU/DU) workloads.
- Scenario 4 (Transport network): Edge computing resources deployed at transport network sites, for example aggregation points.
- Scenario 5 (Regional edge data centre): Purpose-built data centres that are strategically placed in specific regions to serve as regional hubs for computing and data processing needs.

The selection of sustainable edge computing strategies is a decision that hinges on a nuanced evaluation of the specific edge environment and its associated sustainability factors. As discussed in this report, the diverse landscape of edge computing scenarios necessitates tailored solutions. Whether it's accommodating legacy sites, optimising energy usage, minimising water consumption, ensuring equipment longevity, reducing maintenance needs, or streamlining site adaptation, the sustainability considerations are multifaceted. Ultimately, the selection process boils down to identifying the solutions that deliver the lowest total cost of ownership (TCO) while effectively meeting the specific needs of the edge environment. By prioritising sustainability in conjunction with operational efficiency, edge operators can make strides towards a greener and more cost-effective edge computing landscape.

How big an impact can the right strategy have? An example from the telco world

Figure 12: Operator energy consumption breakdown



Source: GSMA 2021, NGMN 2021

According to a 2021 GSMA study⁵, approximately 75% of an operator's total energy consumption is attributable to the radio access network (RAN). Delving deeper, within a base station nearly 80% of the energy is expended by the radio, according to the next generation mobile networks (NGMN) Alliance⁶. The remaining 20% is used for the distributed unit (DU) and centralised unit (CU) processing (see Figure 12).

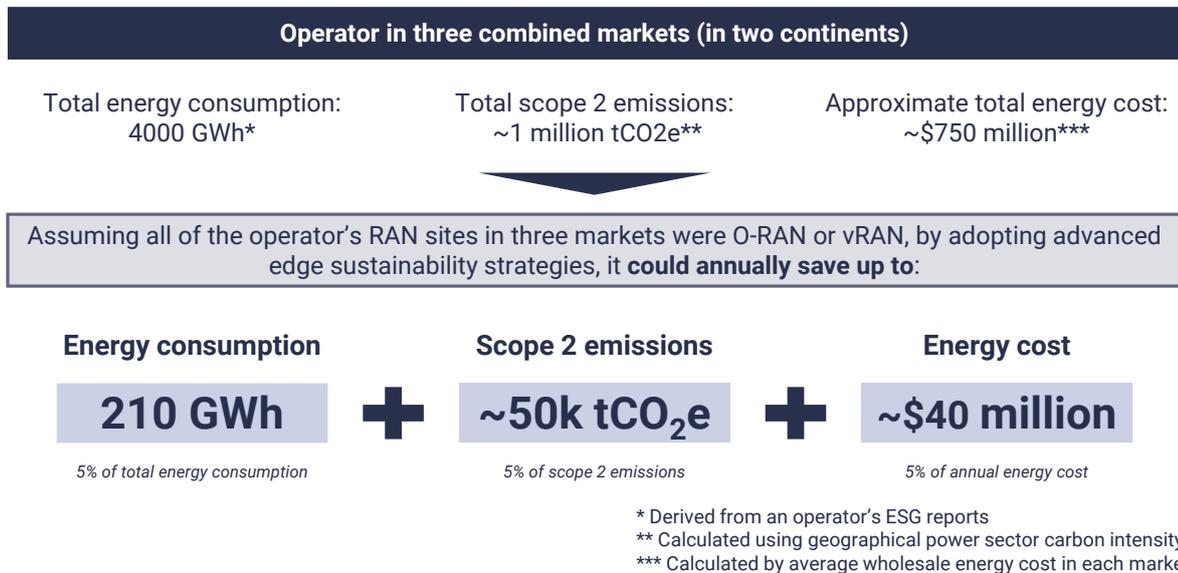
Adopting Open RAN or virtualised RAN (vRAN) is an effective route to eventually decreasing the energy consumption of the RAN through advanced features like sleep modes, load balancing, and other innovative technologies. Find out more in our recent [article](#).

In an Open RAN or vRAN deployment, the DU/CU workloads are processed on COTS hardware, which can also be made more efficient using the aforementioned advanced edge sustainability strategies.

If an operator exclusively deployed Open RAN or vRAN, we estimate that 10-15% of total energy consumption would be attributed to RAN workloads executed on COTS hardware. By embracing the advanced edge sustainability strategies, operators have the potential to diminish the energy consumption on edge-located COTS hardware by up to 40%⁷ (compared to an air-cooled server). This equates to an overall reduction in annual energy usage and associated emissions of 3-5% for a typical operator.

To project these savings figure, we'll use an international group operator as an example. This operator is present in several markets (in two continents) with a combined revenue of around US\$30bn. Figure 13 summaries the potential benefit for this operator from adopting advanced sustainable edge strategies.

Figure 13: Potential annual savings for an international group operator



⁵ GSMA 2021- Going green: benchmarking the energy efficiency of mobile

⁶ NGMN 2021- Green Future Networks, Network Energy Efficiency

⁷ Source: Iceotope

It's important to note that these figures are illustrative and have been derived by STL through some broad assumptions. In reality, there are a multitude of factors that would impact the actual savings and emissions outcome. Regardless, these are still significant savings and highlight the fact that the topic of advanced sustainable edge strategies should not be overlooked. This example has used data from countries with relatively low-carbon grid emissions, meaning that for many other operators, especially in markets with more carbon-intensive electricity grids, could reap even greater benefits. Finally, these figures only represent location-based scope 2 emission reductions (through renewable supplier contracts, many operators' market-based reported scope 2 emissions are lower or zero). While scope 3 emissions are harder to quantify, some of the advanced cooling technologies discussed are expected to extend hardware lifetimes and impact scope 3 too. As mentioned previously, advanced edge sustainability strategies can also reduce scope 1 emissions, for example by reducing maintenance and associated site visits. These benefits have not been modelled in this example.

PARTNERS



Research



Consulting



Events