Revision 4

Whitepaper
The Datacentre of the Future

Rolf Brink

Visit asperitas.com/about/ for more info

Asperitas IMMERSED COMPUTING®

In collaboration with

BILFINGER TEBODIN

A datacentre is not about ICT, power, cooling or security.

It is not even about scale or availability of these systems.

It is about the availability of information.

Realising this, allows an approach to the datacentre environment from a completely different angle, which opens the door to robust and simplified solutions for the biggest challenges facing the industry today.

This document describes this approach, which is called:

The Datacentre of the Future

For more information, feedback or to suggest improvements for this document, please send your suggestions or inquiries to: whitepapers@asperitas.com

- 11-07-2017-Original publication
- 13-07-2017-Second publication (corrections)
- 30-12-2021-Third publication
- February 14, 2021-Sixth publication

Asperitas,

Laarderhoogtweg 18,

1101EA Amsterdam,

The Netherlands

Estimated reading time: 15 minutes

Readability general: basic

Readability technically: simplified

Table of contents

Asperitas
IMMERSED COMPUTING®

1. From energy consumer to energy producer

The global move to cloud based infrastructures and the Internet of Things (IoT) generate high demand for datacentre capacity and high network loads. The energy demand of datacentres is rising so quickly that it is causing serious issues for energy grid operators, (renewable) energy suppliers and governments. Grid operators and energy suppliers can hardly keep up with the demand in the large datacentre hubs let alone ensure enough renewable energy generation is available where we see the demand.

Not only does this raise questions of sustainability on all levels, the demand for flexibility and high loads requires a different approach to the business model of the datacentre. With the ultimate challenge of becoming an energy neutral industry.

The key to resolving this challenge is the adoption of liquid cooling techniques in all its forms. Asperitas is committed to approaching this challenge head-on. Asperitas is taking an active role in letting go of the limitations of IT systems and datacentre infrastructures of today and finding new ways to drastically improve datacentre efficiency, helping to realise a future where the datacentre transforms from energy consumer to energy producer.
In this whitepaper, we present our vision of the datacentre of the future. A datacentre that faces the challenges of today and is ready for the opportunities of tomorrow.

“We are building the biggest thing humans will ever build. This is the dawn of our magnum opus. We have the opportunity to see the big picture, and build it responsibly.”

Patrick Flynn, Senior Director of Sustainability, Salesforce

Asperitas IMMERSED COMPUTING®

Reading note

Some technical details and scientific data in this document have been simplified. Degrees are used in Celsius where Kelvin should be mentioned, detailed and critical parts of water circuits have been omitted or simplified and variable environments are represented as being stable. This is purposely done to reduce the complexity of this document and ensure the basic understanding of the involved concepts.

2. The advantages of liquid

Most of the challenges facing datacentres today can be solved by introducing liquid cooling techniques in all its forms. So why liquid? The use of liquid as a cooling agent has numerous advantages over air.

2.1 Maximising heat removal and reuse

For datacentres to become a carbon neutral operation, energy reuse is key. Electrical energy cannot just be “consumed” or “destroyed” anymore, it must be converted into usable thermal energy which is utilised.

Water is a far more suitable medium to absorb, transport and reject thermal energy in datacentres. With a much higher heat capacity and density compared to air, water can absorb and transport almost 3500 times as much thermal energy per second.

Already **water circuits are commonly** used to allow for absorption and rejection of thermal energy throughout datacentre facilities. Because water is much easier to distribute, it is routed through chiller systems (CRACs or Computer Room Air Conditioners) to enable the cooling of the air inside the IT room. Some technologies already go further and utilise water circuits to provide cooling for IT racks or equipment on the data floor (white space). The main purpose of these systems is to capture as much heat as possible directly into a water circuit and to prevent the use of inefficient air chillers.

Air to liquid comparison

A 1 MW critical load datacentre produces 1 MJ of thermal energy every second. The amount of fluid movement required to harvest and transport 1 MJ/s of thermal energy with a rise of 10°C can be calculated.

Air has a heat capacity of 1005 J/kg°C.

With a density of 0.001205 kg/L, a 1°C rise carries: 1005 J/kg°C * 0.001205 kg/L = 1.2 J/L per 1°C.

Absorbing 1 MJ/s with a 10°C rise: Absorbing 10°C with one litre of air/second carries:

10°C *1.2 J/L per 1°C =12 J/s per L/s.

Therefore 1 MJ/s now requires:

1000000 J/s / (12 J/s per L/s) =83333 L/s air.

Water has a heat capacity of 4187 J/kg°C

With a density of 1 kg/L, a 1°C rise carries: 4187 J/kg°C * 1 kg/L = 4187 J/L per 1°C.

Absorbing 1 MJ/s with 10°C: Absorbing 10°C with one litre of water/second carries:

10°C *4187 J/L/s=41870 J/s per L/s.

Therefore 1 MJ/s now requires: 1000000 J/s / (20935 J/s per L/s) = 24 L/s water.

Thermal transport efficiency

Fans that create 100 L/s (approximately 200CFM) of airflow with 1cm water head of pressure difference to overcome obstructions to flow require 55W of fan power.

A pump can harvest the same amount of heat against a 2m water pressure head with a flowrate of only 0.03 L/s using around 1W – requiring less power and space and higher compute density.

Using other thermodynamic principles with the above flowrates, liquid can be transported 200 times the distance of air with the same thermal losses.

2.2 Increasing temperatures

Besides the capability to transport energy with minimal heat loss, the second decisive factor to make energy reuse viable is a sufficient temperature difference (ΔT). Since air based systems are designed to allow relatively small temperature windows for normal operation, temperatures for reuse are often too low to facilitate direct reuse scenarios.

Introducing liquid cooled IT radically changes the possibilities for datacentre cooling and energy reuse. Since liquid has a high heat capacity and a much higher density than air, it allows electronic components to be cooled more effectively. This enables higher environmental operating temperatures for IT and results in higher quality thermal energy for reuse. Especially when mixing different liquid technologies for different platforms, high temperature differences can be achieved in the water circuit. Not only does this facilitate energy reuse, it also reduces facility overhead installations.

2.3 Minimising facilities

The adoption of liquid cooled IT in datacentres allows for more effective utilisation or reduction of the datacentre footprint. This means that an existing facility can be better utilised to allow for more IT.

The higher heat capacity of liquids allows **for more dense IT environments and higher IT capacity.** With most liquid technologies, the IT itself becomes more efficient. This is caused by the reduced or eliminated dependence on air handling within the IT chassis. Individual components are cooled more effectively and can therefore be used with higher amounts of energy and closer to each other. When liquid penetrates the IT space, internal fans are reduced or completely eliminated. This can save 6-45% of the IT energy footprint in the facility. This also reduces the emergency power requirements within the facility. After all, eliminated energy consumption does not require UPS or generator sets.

It also **reduces the overhead installations** of the datacentre facility. First of all, because of the reduced volume of the fluid which requires cooling. Secondly because with most liquid technologies, much higher temperatures can be supported. This allows for further reduction of expensive cooling systems.

Fan energy consumption

In general, fan overhead under load will vary between 6% and 45%. These are the extremes under stable conditions.There are many factors which relate to the actual consumption of fans. These factors can further increase or decrease the actual overhead of fan energy.

Some examples are:

Chassis size: 1U server chassis often cause high overhead due to smaller fans, where higher chassis (2U+) are much more efficient.

Server design: Air flow resistance within the chassis influences the power consumption of fans.

Environmental temperature: Fans often work harder in higher environment temperatures to allow sufficient cooling.

Server utilisation: Cloud environments are often utilised between 30% and 50%, whereas the highest server and fan efficiency can be achieved at around 70%-80% utilisation.

2.4 Reducing maintenance

The increased heat capacity of liquid greatly **reduces thermal shock and stress** on IT components. Because it takes more energy to heat up a liquid environment, energy fluctuations (CPU throttling) will have less impact on temperature fluctuations. With Total Liquid Cooling (see next chapter), oxygen does not get in touch with IT components, which eliminates oxidation altogether.

Although liquid circuits also require maintenance, due to the reduced or in some cases completely eliminated cooling installations, the overall **facility maintenance is greatly reduced.** When IT is optimised for higher thermal properties, the amount of IT can also be reduced, which again reduces maintenance and operating cost.

2.5 Flexibility

Liquid allows for **great flexibility** due to the simplified cooling infrastructure. Installations can be quickly put in place or expanded with a minimised footprint. This enables simplified datacentre designs with reduced environmental impact. These can be set up more easily with less construction and deployment time for both new installations and upgrades.

2.6 Designing the datacentre of the future

With all these advantages, liquid offers solutions that are just not attainable in any other way. This is why **liquid is the future for datacentres**. But what does this future look like? Which liquid technologies are available and what does this mean for the infrastructure of the datacentre?

In the next chapter, we outline the basic liquid technologies operating in datacentres today. After that we explore the most beneficial environment for these technologies: a hybrid temperature chain. Further on, a model of connected, distributed datacentre environments is introduced.

With the dedication to liquid, Temperature Chaining and the distributed datacentre model, the datacentre of the future transforms **from energy consumer to energy producer.** This approach will drastically reduce the carbon impact of datacentres while stimulating the energy efficiency of unrelated energy consuming industries and consumers.

3. Liquid technologies

Liquid technologies available today can be roughly divided into four different categories: cooling at the room, rack or chip level and immersion.

All shown brand names maintain very high design standards and they all use an intelligent control system which allows complete optimisation and control over the water circuit. All shown technologies can operate in the same datacentre environment.

Image 1 CRAC

Computer Room Air Conditioning (CRAC) can be water cooled. This allows for an air cooling set-up which is least economical, but in some cases unavoidable due to (non-racked) legacy systems which require cooling or due to partially liquid cooled IT systems.

Image 2 ILC by U-Systems

Indirect Liquid Cooling (ILC) involves water cooled racks with (active) rear door or in-row heat exchangers which are water cooled. The advantage of the active rear doors is that all the heat from air cooled IT is immediately absorbed by the water circuit when it leaves the rack which eliminates the need for CRACs, also in partial ILC implementations. This makes cooling systems very efficient, and supports limited efficiency on the IT itself by assisting ventilation.

See www.usystems.co.uk for more information on water cooled racks.

Visit usystems.co.uk

Image 3 DLC by Asetek

Direct Liquid Cooling (DLC) effectively cools parts of the IT with purpose built coolers which combine cold plates and pumps that are mounted directly onto the chips instead of a traditional heat sink. This generates energy efficiency on the IT side due to the reduced amount of fan energy. Although the water circuit captures all of the heat from the largest heat sources inside the chassis, this approach may still require CRAC units or combinations with ILC for rejection of thermal energy from the rest of the IT components.

See www.asetek.com for more information on Direct-to-Chip Liquid Cooling.

Total Liquid Cooling (TLC) completely immerses the IT components in liquid. There is hardly any energy loss and IT equipment is made very energy efficient, eliminating kinetic energy (fans) from being used by the IT. Since water conducts electricity, an intermediate dielectric substance is required which requires forced or convective transfer of heat. This dielectric can be oil or chemically based. The infrastructure and power advantages are maximised with this approach and the energy footprint is fully optimised.

Image 4 TLC by Asperitas

See www.asperitas.com for more information on Total Liquid Cooling

Visit asperitas.com

4. The hybrid infrastructure

Introducing water into the datacentre whitespace is most beneficial within a purpose-built set-up. This means that the focus for the design of the datacentre must be on absorbing all the thermal energy with water. This calls for a hybrid environment in which different liquid based technologies are co-existing to allow for the full range of datacentre and platform services, regardless of the type of datacentre.

4.1 Optimal technology

Since there is no such thing as one solution for all, any platform should be designed with the optimal technology for its different elements. Therefore, each part of a platform should be set up with a mix of optimised technologies. For example, storage environments are least suitable to be cooled directly by liquid due to the low energy production and the common dependency on moving parts. These can be set up in water cooled racks. High volumes of servers which require the least maintenance can best be positioned in a Total Liquid Cooling environment. Varying specialised server systems which require constant physical access are best situated in Direct Liquid Cooled environments.

Platform	Non-racked	Network	Storage	Compute	Cloud
CRAC	x				
ILC (U-Systems)		x	x	partial	partial
DLC (Asetek)			partial	x	x
TLC (Asperitas)		partial	partial	x	x

Image 5 Liquid technology for different parts of a platform

4.2 Temperature chaining

By adopting the hybrid model, systems can be connected to different parts of a cooling circuit with different temperatures. Each liquid technology has different temperature tolerances. Especially where the liquid penetrates the chassis, the stability of temperatures becomes less of a concern. Therefore, different technologies can be set up in an optimised order of tolerance to allow a multi-step increase in temperature within the cooling circuit.

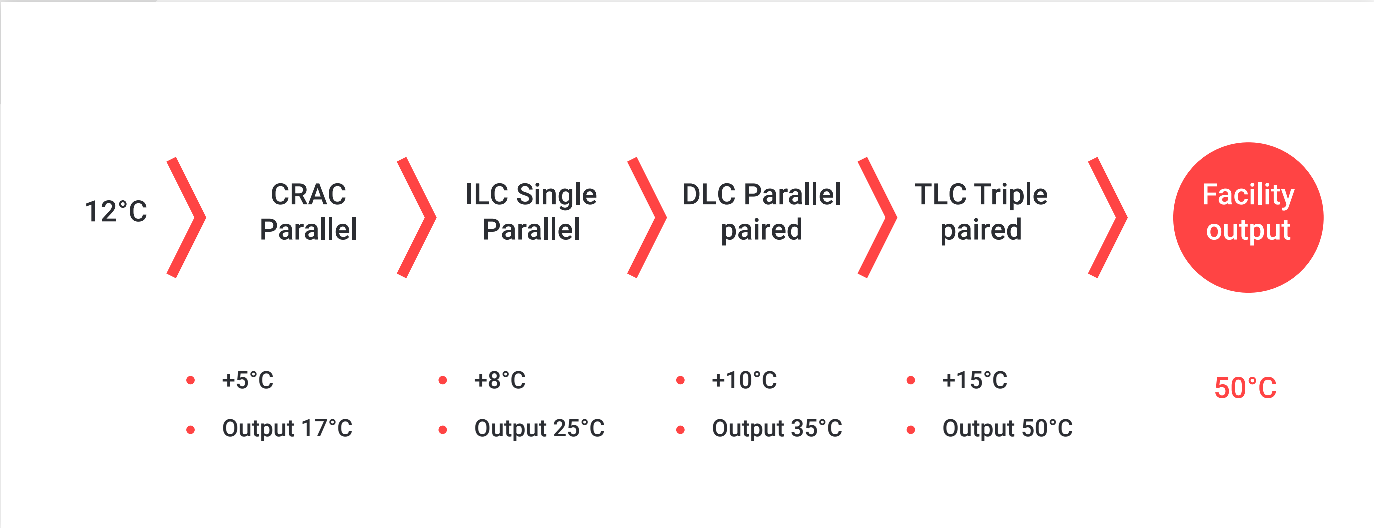

Image 6 Temperature chaining principle

This means that the water infrastructure becomes segmented. Instead of feeding each cooling setup in a parallel infrastructure, the inlets of different technologies or different parts of the infrastructure are connected to the return circuit of another part of the infrastructure. In essence, the output of a liquid cooled rack should not be routed to a cooling installation, but to a different type of liquid cooling environment. By chaining the segmented liquid circuits in larger environments, very high return temperatures can be achieved, which enables the practical and effective reusability of thermal energy and decreases investments needed to make large scale heat reuse a viable option.

Technology	Inlet range		Outlet range		Maximum
CRAC	normal	extreme	normal	extreme	delta/rack
ILC (U-Systems)	6-18°C	21°C	12-25°C	30°C	N/A
DLC (Asetek)	18-23°C	28°C	23-28°C	32°C	15°C
TLC (Asperitas)	18-40°C	55°C	22-48°C	65°C	10°C

Image 7 Temperature compatibility with different liquid technologies and maximum "practical" temperature deltas

The different liquid technologies can be applied with different temperature levels. There is a difference between normal optimised environments and more "extreme" environments where the solutions and IT equipment are more compatible or specialised for high temperature operation.

Temperature chaining

We have already determined the required volumes for water and air to absorb and transport 1 MJ/s with a ΔT rise of 10°C: **83333 L/s air vs 24 L/s water.**

With Temperature Chaining, the same infrastructure now only requires the following amount of water:

1MJ/s with 10°C requires 24 L/s water.

The increased ΔT of 40°C can handle 4 times the amount of energy rate. Therefore, the amount of water required is now: 24/4=6L/s

This is a factor of nearly 14000 to 1.

Liquid Temperature Chaining can be implemented by adopting intermediate cooling circuits with different temperature ranges. Segmented environments can be connected with supply and return loops, mixing valves and buffer tanks to stabilise and optimise the return temperatures and volumes of each individual segment.

A major advantage of this strategy is the fact that temperature differences (DT) within a cooling circuit can be drastically increased. This reduces the volume of liquid required in a facility and reduces the cooling overhead installations.

After all, it is much more efficient to cool a large DT in a small volume of water than a small DT in a large volume of water.

4.3 Temperature chaining examples

In the following paragraphs, we outline examples for hybrid environments. Each part of a digital platform has different requirements and each datacentre has different challenges. These examples only provide insight into optimised liquid infrastructures to explain the concepts of Temperature Chaining and how different liquid technologies can fit into this concept. For simplification purposes, there are no redundant scenarios outlined. Return loops, buffer tanks and intermediate pumps to deal with volumetric and pressure aspects within different stages are not detailed, neither are actual cooling installations. There may be a chiller requirement when heat reuse is not an option, although chiller efficiency is highly improved with the Temperature Chaining scenarios below.

4.3.1 Open circuit heat reuse infrastructure

The open circuit heat reuse infrastructure is the most sustainable infrastructure by far. In this situation, the datacentre receives water of a certain temperature and all the heat generated by the IT equipment is delivered to another user with this water circuit. This means that the facility does not only reject the heat, but also the water which contains the heat to allow an external party to transport and use the warmed-up liquid. This results in a complete lack of cooling installations and the datacentre effectively acts like a large water heater. Water flows into the datacentre and comes out at high temperatures.

The ILC racks in this setup effectively function as air handlers which maintain the entire room temperature and absorb all thermal energy leakage from the DLC and TLC environments.

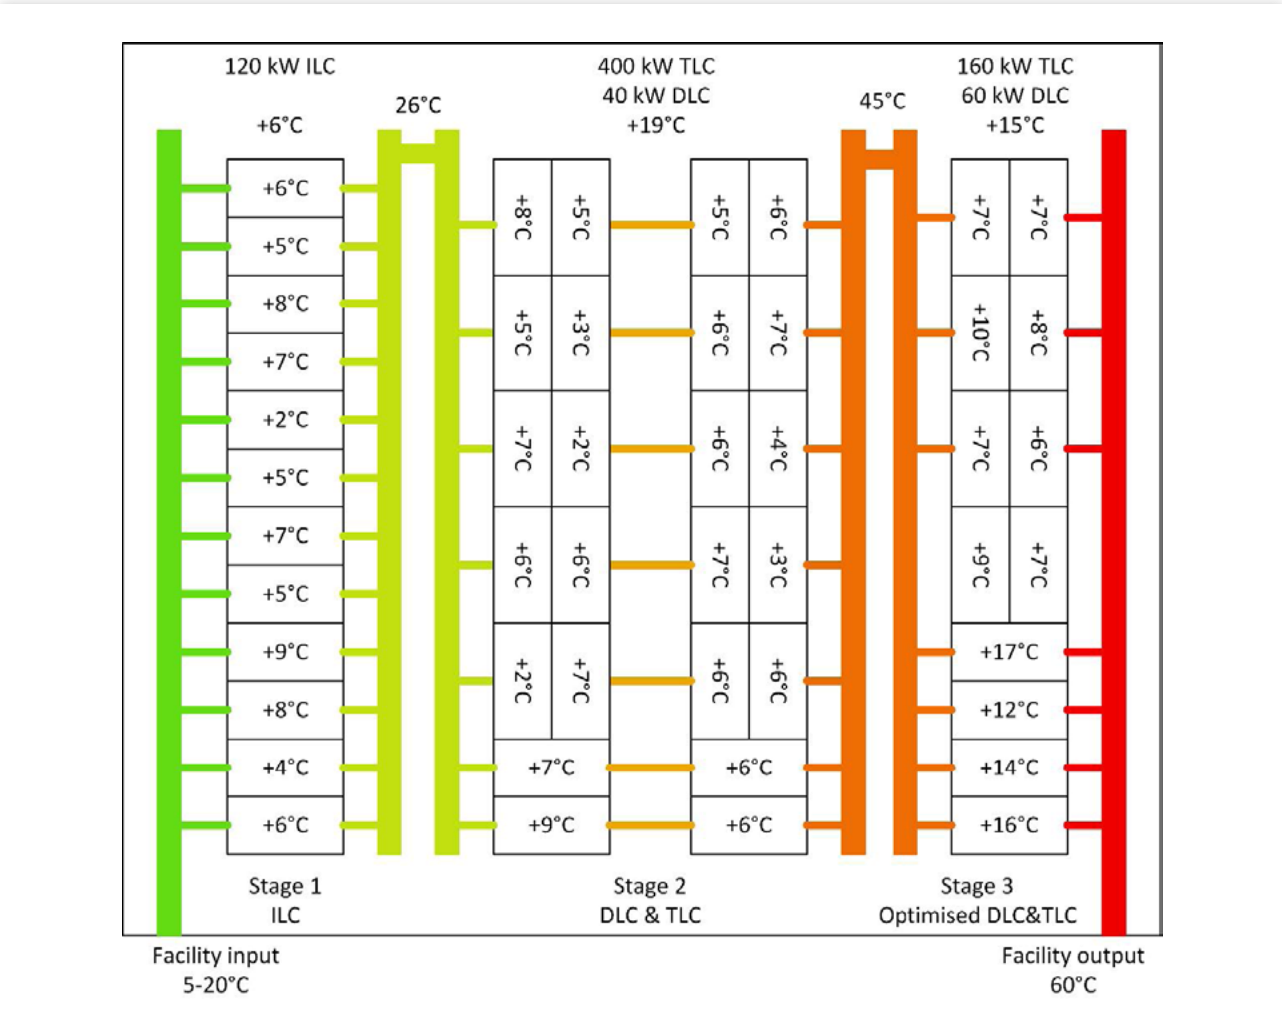

Image 8 Open Circuit heat reuse schematic (temperatures are random)

4.3.2 Micro infrastructure

In smaller footprints, the same result can be achieved by creating a small water circuit with a mixing valve and buffer tank. This allows the output of the liquid installation to be routed back to the cooling input to gradually increase the cooling circuit and achieve a constant high output temperature.

The advantage of this approach is the compatibility with variable input temperatures which are common with dry cooling installations.

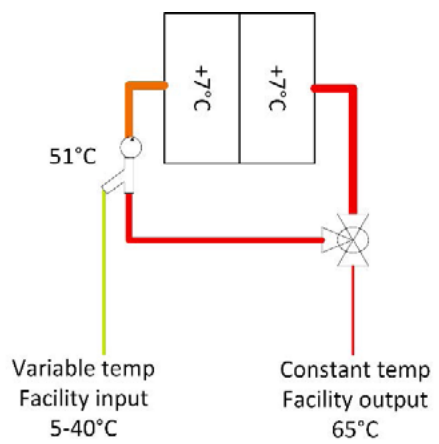

Image 9 Micro datacentre reuse schematic

Image 10 Micro reuse circuit at Asperitas Tech Centre

4.3.3 Chiller based infrastructure

Chiller based infrastructures allow the same circuit to be easily expanded by putting systems in line and thus making existing chillers more efficient and effective. It also allows for reverse compatibility with air cooled installations by implementing more traditional CRAC units in the same water circuit.

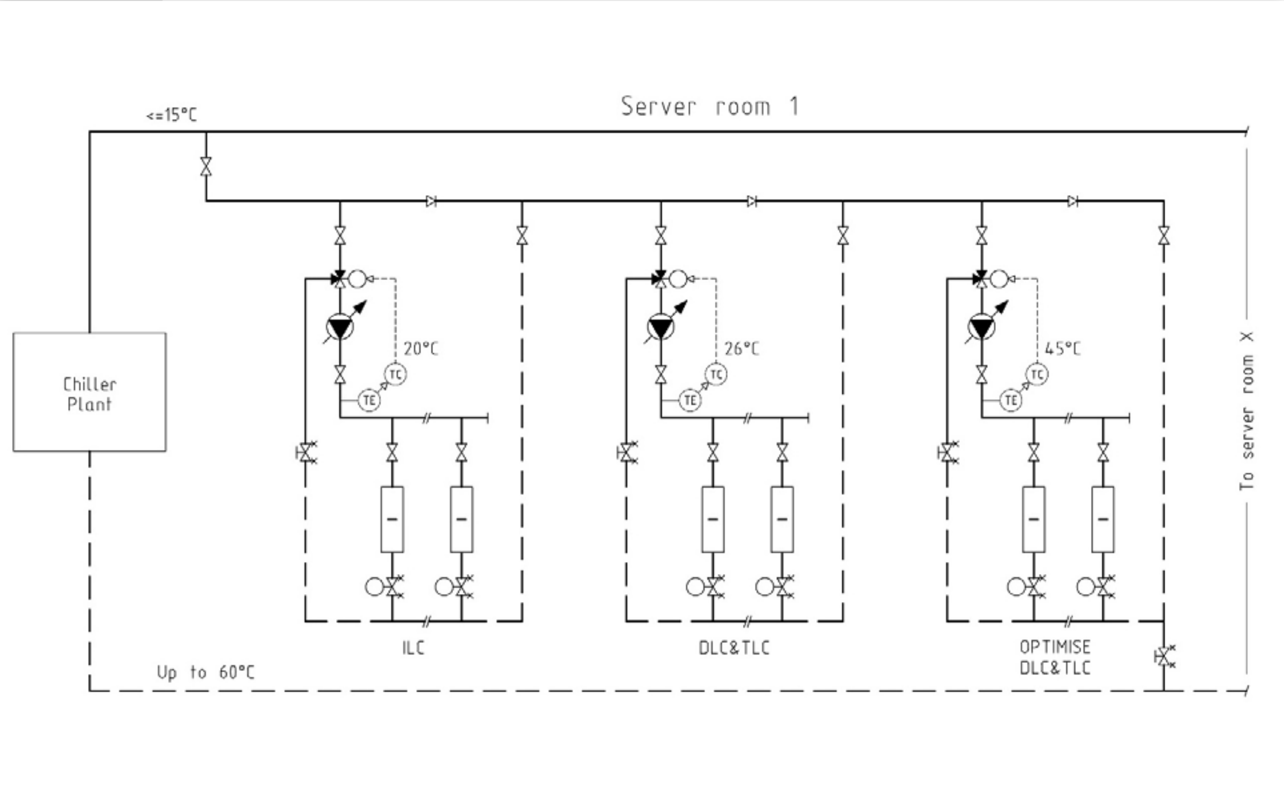

Image 11 Chiller based Temperature Chaining by Tebodin Bilfinger

4.3.4 Free cooling infrastructure

With water-base free cooling systems, the entire operation becomes even more economical. The prerequisite for this being no air-cooled systems, or that these systems are compatible with high environmental temperatures to account for high outside air temperature conditions. This means that the entire system should be able to be supplied with water at temperatures up to 36°C so that long periods of free cooling are possible (depending on geographical location).

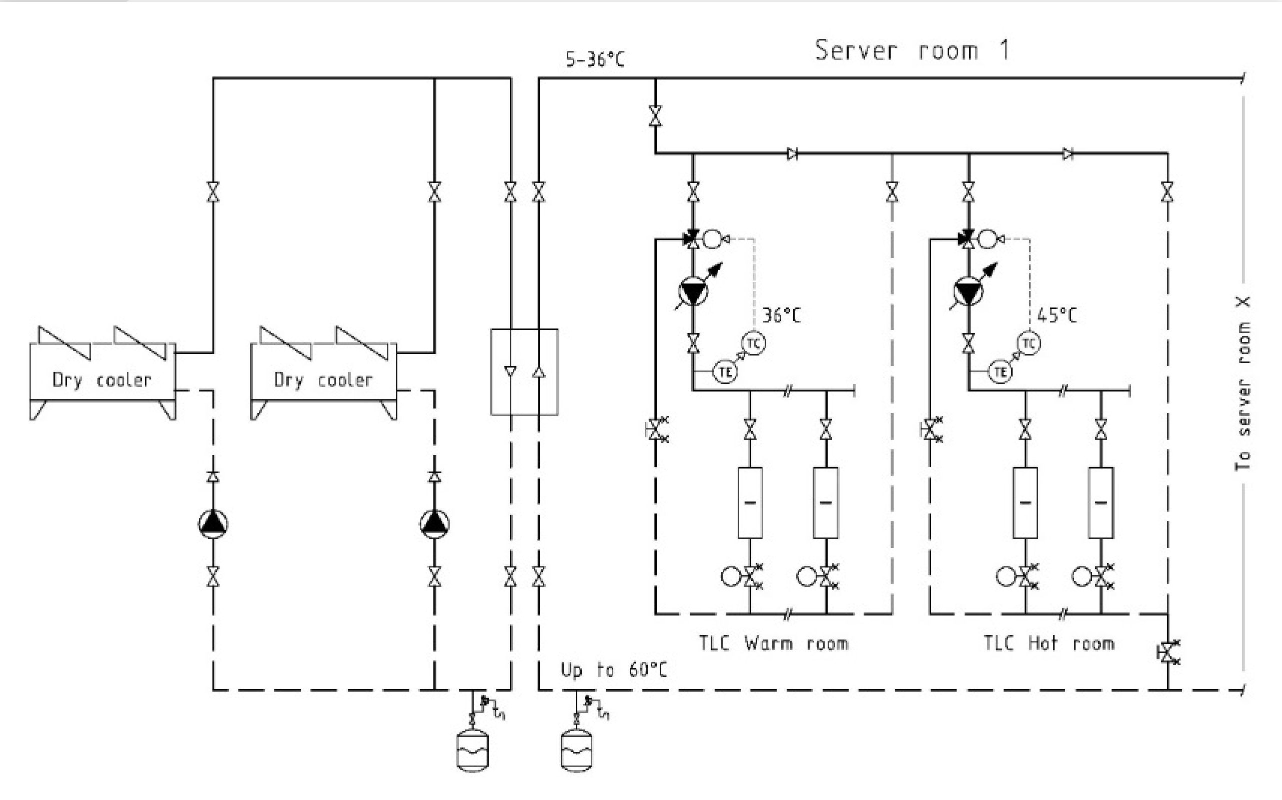

Image 12 Free Cooling based Temperature Chaining by Tebodin Bilfinger

5. The connected datacentre web

Immersed Computing® allows for easy deployable, scalable local edge solutions. These allow for rejecting heat to whatever reuse scenario is present, like thermal energy storage, domestic water, city heating etc. If no recipient of heat is available, only a dry cooler is sufficient.

Reducing or even eliminating the need for overhead installations like coolers for edge environments, allows for a different perspective on datacentres. Geographic locations become easier to qualify and high quantities of micro installations can be easily deployed with minimal requirements. These datacentres will be integrated in existing district buildings or multifunctional district centres. A convenient location for the datacentre is a place where heat energy will be utilised throughout the whole year. Datacentres can also be placed as a separate building in residential and industrial areas.

This creates the potential for a connected datacentre web consisting of mainly two types of datacentre environments.

- Large facilities (Core Datacentres) which are positioned on the edge of urban areas or even farther away;
- Micro facilities (Edge Nodes) which are focused on optimising the large network infrastructure and are all interconnected with each other and with all core datacentres.

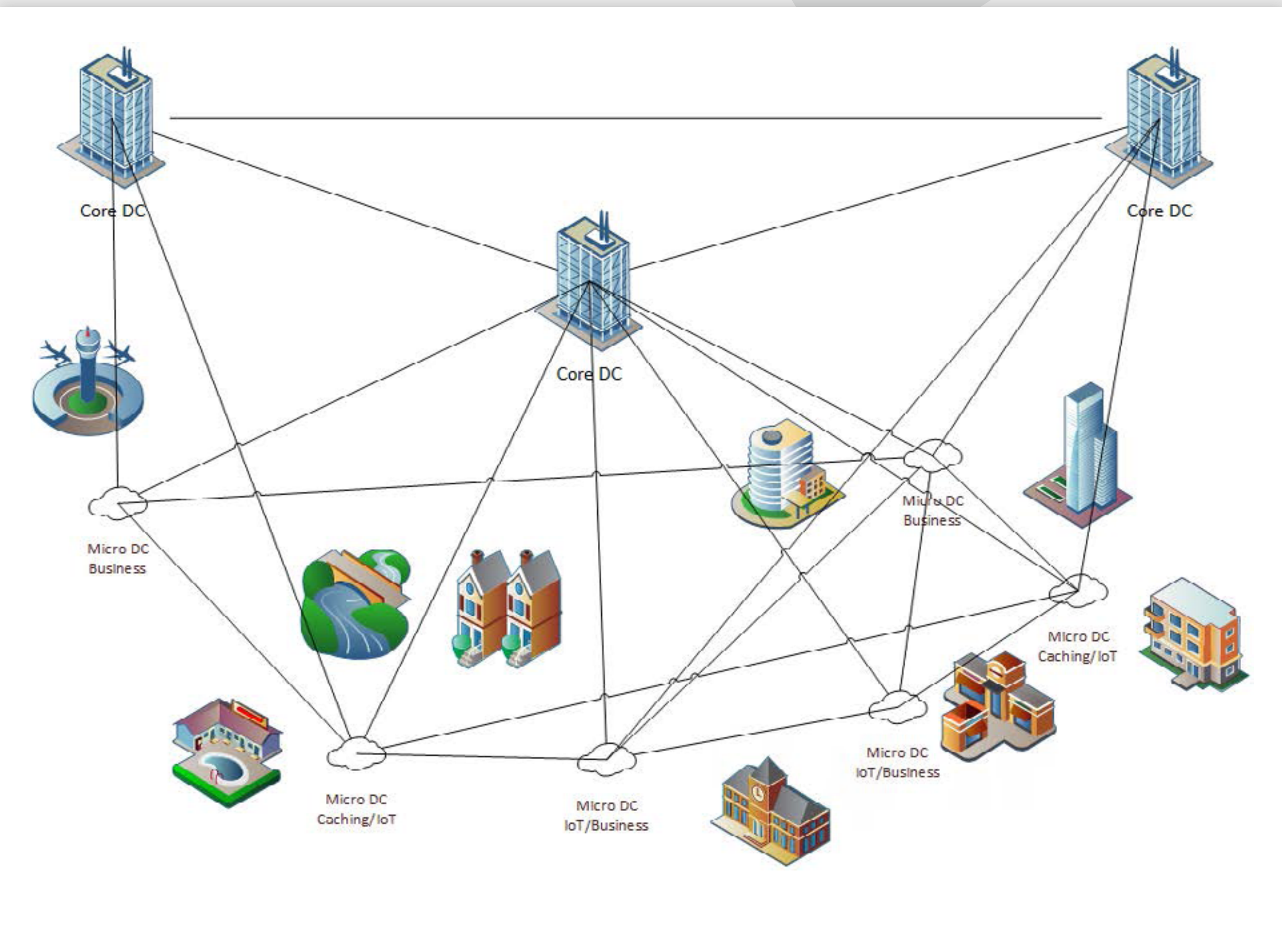

Image 13 Datacentre environment consisting of core datacentres and edge nodes

The main purpose of the edge nodes is to reduce the overall network load and act as an outpost for IoT applications, content caching and high bandwidth cloud applications.

The main function of the core datacentres is to ensure continuity and availability of data by acting as data hubs and high capacity environments.

Asperitas
IMMERSED COMPUTING®

5.1 Core datacentres

The core datacentres are large datacentres which have most of the characteristics of datacentres today. The facilities are optimised for large amounts of IT with all the redundancy and continuity aspects of a high tiered datacentre. The main difference is that the entire facility is equipped with liquid. It utilises an open circuit cooling approach with Temperature Chaining. The main purpose of this datacentre environment is to provide availability and data protection and to function as a main network hub.

The core datacentres are geo-redundant, which means that the facility is effectively replicated with other core datacentres to provide optimal availability of data.

If the core datacentre is supporting multiple tenants, the infrastructure is set up in such a way, that any customer of the datacentre will benefit from the overall networked datacentre infrastructure and architecture.

5.2 Distributed micro edge nodes

The micro edge nodes (10-100 kW) function as forward locations of the core datacentres. The edge nodes provide services like data processing for IoT systems, data caching for digital content (YouTube, Netflix, etc.) and fast access to cloud services. The edge nodes are continuously replicated with the core datacentres and several strategic other edge nodes. This provides constant availability through geo-redundancy.

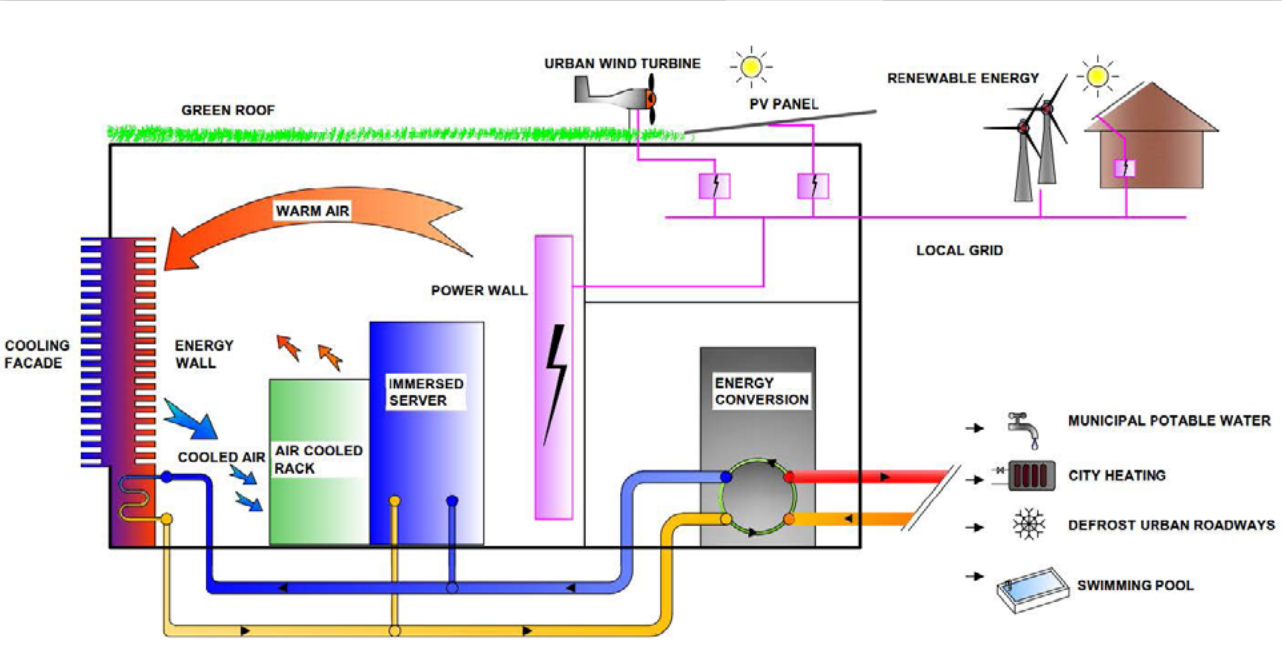

Image 14 Concept for a purpose built micro edge datacentre by Tebodin Bilfinger

By making information available in multiple locations at the same time, it becomes easy to exchange between different physical structures when interacting with the information. The capacity of overhead installations can be minimised to allow only for normal operation and a shutdown phase in case of emergency while active data processes are moved to a different facility.

The micro edge nodes are small locations with minimised overhead installations. They will have simplified configurations which consists of a small data floor, switchboard and energy delivery. Often without redundancy in power or cooling infrastructure (there is a significant thermal buffer with Immersed Computing®), but with sufficient sustainable Li-ion battery power (Tesla Power-pack) to allow for replication and shutdown. The facilities are based on Immersed Computing® and additional liquid technologies when required. This allows these facilities to become enclosed air environments which prevents environmental impact like noise or exterior installations. The liquid infrastructure is cooled with any present external cooling strategy which is available.

Image 15 Tesla Powerpack

Due to the minimised overhead, these nodes can be deployed in large quantities near or within network hubs for urban or office areas or even as part of a non-datacentre facility which can directly benefit from the reusable heat. This allows for fast network access and simple energy reuse.

Energy reuse by water mains

An example for heat reuse from micro datacentres in urban areas

A heat exchanger onto the water mains is utilised as the cooling strategy, changing the water mains temperature with a minimal amount (0,5-2°C, depending on actual volume and flow). The energy saving is calculated only for the hot water creation by boiler in each household in the district which it facilitates. All cold-water use is part of the inefficiency.

To estimate the impact, we must first determine the impact of hot water use per household (based on gas boilers and statistics from The Netherlands):

Average gas use per household for hot water	**300 nm³/year**
Boiler efficiency	90%
Calorific value gas	31.65 MJ/nm³
Energy for hot tap water	8546 MJ/year
ΔT heating hot tap water	55°C
Avg. hot tap water usage	36994 L/year
Hot tap water usage	**101 L/day**
Total tap water usage	350 L/day
Hot tap water share of total	29%

We can now determine the impact of a micro edge datacentre on a small district.

Heat production small micro edge	16 kW
Total transfer to water mains	140160 kWh/year
Energy reduced for hot water	**40587 kWh/year**
Gas equivalent	**5129 nm³/year**

Note: The unit 'nm³' stands for "Normal Cubic Meter". It is the amount of gas in cubic metres at the pressure of 101.325 kPa and a temperature of 15°C.

5.3 Management platform

The management of the distributed datacentre model, will be possible through the emergence of software platforms providing ubiquitous management of data, network and computation capacities. These kind of platforms already exists for traditional centralised infrastructure, but new challenges emerge from this hybrid and distributed architecture. Closer to the end users, edge nodes in urban areas have new constraints in terms of energy consumption and heat production. Containerisation, through technologies like Docker or Singularity, opens great opportunities to make applications more scalable, flexible and less dependent on the infrastructure. Many frameworks appeared recently (Swarm, Kubernetes) to manage decentralised clusters. Some of them also integrate energy and heat management by design like "**Q.ware**" developed by **Qarnot computing**. This positive dynamic in the software industry is an essential pillar to enabling core datacentres and edge nodes with an integrated architecture.

See www.qarnot.com for more information on the management of distributed environments.

Visit qarnot.com

5.4 Network optimisation

The use of core datacentres and edge nodes allows for network optimisation by preventing long distance transport of raw (large) data and allowing the processing of data close to the source. By bringing data which is in high demand closer to the end user (caching), high volume data transmission across long distance backbones is greatly reduced, as well as latency which is a critical factor for delivering good end user experience.

Cooling strategies in the edge

There are numerous edge cooling strategies which are optimal for the scale of micro edge nodes. All strategies require 24/7 thermal rejection, thus completely eliminate the need for cooling installations.

Here are a few commonly available cooling strategies in urban areas:

- Spas and swimming facilities with multiple pools have a constant demand for heating due to constant convection (100% reuse).
- Hospitals and hotels equipped with warm water loops which require constant 24/7 thermal input (100% reuse).
- Urban fish and vegetable farms using aquaponics (100% reuse).
- Aquifers for energy storage, these can normally be supplied with thermal energy 24/7 (75% reuse).
- Water mains can provide distributed energy savings (29% reuse).
- Canals, lakes and sewage water can be used for heat rejection when reuse is not possible (0% reuse).

5.5 Energy grid balancing

One of the limitations for datacentre growth today is the capacity of the existing power grid. In most areas in the world, the power grid was designed and implemented long before datacentres even existed. There are numerous areas where the power grid will reach its maximum capacity within the next 3-5 years. The traditional datacentre approach causes high loads on very specific parts of the grid. By applying the distributed datacentre model, the power grid is more balanced and the impact of expansion greatly reduced.

5.6 Becoming an energy producer

By focusing on the reuse of energy, each edge node rejects its thermal energy directly into a reusable heat infrastructure (district heating/ heat storage), building heating (hospitals/ industry), water heating (hospitals/zoos) or other heat users. The core datacentres become large suppliers of district heating networks or will be connected to 24/7 industries which require constant heating within a large scale industrial process.

6. Sidenote: need for a different KPI

PUE has been embraced as a major KPI for datacentres for many years. It has done a great job for the greening of the datacentre industry because of its simplicity and effectiveness in differentiating between core business and overhead. Simply put:

Datacentres exist to facilitate IT. Therefore, IT power consumption is essential and everything else is overhead. Overhead must be minimised. Cooling is the greatest factor for overhead, so this is where you save energy.

However, reality is a bit more complicated. PUE was never intended to be an “efficiency” indicator among datacentres. PUE works well for power thirsty legacy installations and datacentres. Reality is much different today. We are already building to the highest PUE standards so the smaller incremental improvements are rarely worth the additional investment.

The enormous growth, driven by IoT and the unprecedented move to cloud based computing, demands that we now move inside the IT space to reduce energy usage. Direct and Total liquid cooling strategies do exactly this by reducing or eliminating the fan energy within IT which saves 6-45% of the IT energy footprint.

There's only one problem…

Saving IT energy, ruins our efficiency model since we're saving in the denominator which is on the wrong side of the equation.

Therefore, as soon as any datacentre hits a PUE of 1.2 or when the IT environment includes liquid cooling scenarios, a different metric must be used.

There are many metrics out there, but the challenge with each one of them is the fact that they are not as straightforward as PUE. Either because the required information is not easily available or because the formulas involved are too complex.

Action and involvement from governments, IT manufacturers, datacentre operators and datacentre suppliers is required to allow for more efficient datacentre design and a better metric to stimulate these.

7. When is this future?

Sustainability has been a theme within the IT industry since the introduction of the energy star label for hardware. Since 2007 sustainable digital infrastructure became a topic internationally shaped by a variety regional organisations running green IT programs to reduce the footprint of layers within digital infrastructure including datacentres. Reducing the overall footprint of the datacentre industry is an impossible challenge as the rate growth of the adoption of digital services, and therefore the need for infrastructure, has outpaced the rate at which the energy footprint can be reduced. Industry orchestration will be needed if this challenge is to be met and new standards for datacentre sustainability are to be set.

Years have passed and excess datacentre heat reuse is still not delivering its promise, often because technology, organisational and operational elements cannot be matched. Recently stakeholders have been aiming to change that by stimulating a new role for datacentres, which is the transformation we still need to see on a large scale: the transformation from energy consumers to flexible energy prosumers.

Governments are important stakeholders within the development of sustainable datacentres, from an EU level to local city halls. Today, policies have a limited reach and focus on elements which should be part of a holistic vision. Few regions acknowledge the fact that the datacentre industry is a key player to reach their climate or circular goals, let alone include specific ambitions for this industry as part of their overall sustainability plans and policies.

This document describes the datacentre of the future.

However, all the technologies which are described in this document are available today. There is nothing which still needs to be developed, no technology which cannot be implemented and the impact of datacentres on the world is real today.

Datacentre of the future?

The future is now

This statement is fully supported by the “Liquid Alliance”

8. Asperitas and Tebodin Bilfinger

Asperitas is a clean-tech company focused on greening the datacentre industry by introducing Immersed Computing®.

Since 2014 Asperitas has worked on validating and developing Immersed Computing® as a unique approach to the datacentre industry. Building on existing liquid immersion cooling technologies by adding integration of power and network components, improving cooling physics with a strong focus on design and engineering for usability, Asperitas has come up with a complete and integrated solution which can be effectively utilised in most, if not all situations.

The Asperitas development partners include Intel, SuperMicro, Boston, Shell, Schleifenbauer Aqualectra and Brink Industrial. Asperitas is leading the industry through OCP and ASHRAE. Asperitas is furthermore recognised and supported by the Netherlands Enterprise Agency as a Cleantech company.

Tebodin is part of the international engineering and consulting firm **Bilfinger** and offers an integral range of consulting & engineering solutions for datacentres, often combining knowledge and insights from other industry sectors. For example, the securing of critical systems in the Oil & Gas industry.

Tebodin and Asperitas have established a partnership to bring Immersed Computing® to data-centres and to serve the enterprise market. Tebodin is adopting Immersed Computing® in sustainable and innovative datacentre concepts, offering customers significant cost and energy footprint reduction, enhanced stability and modular expansion as and when required.

The first concept of an energy efficient **datacentre edge node** was presented by Tebodin at the international launch event of Asperitas at Cloud Expo Europe in London.

See www.tebodin.bilfinger.com for more information about Tebodin Bilfinger.

Visit tebodin.bilfinger.com

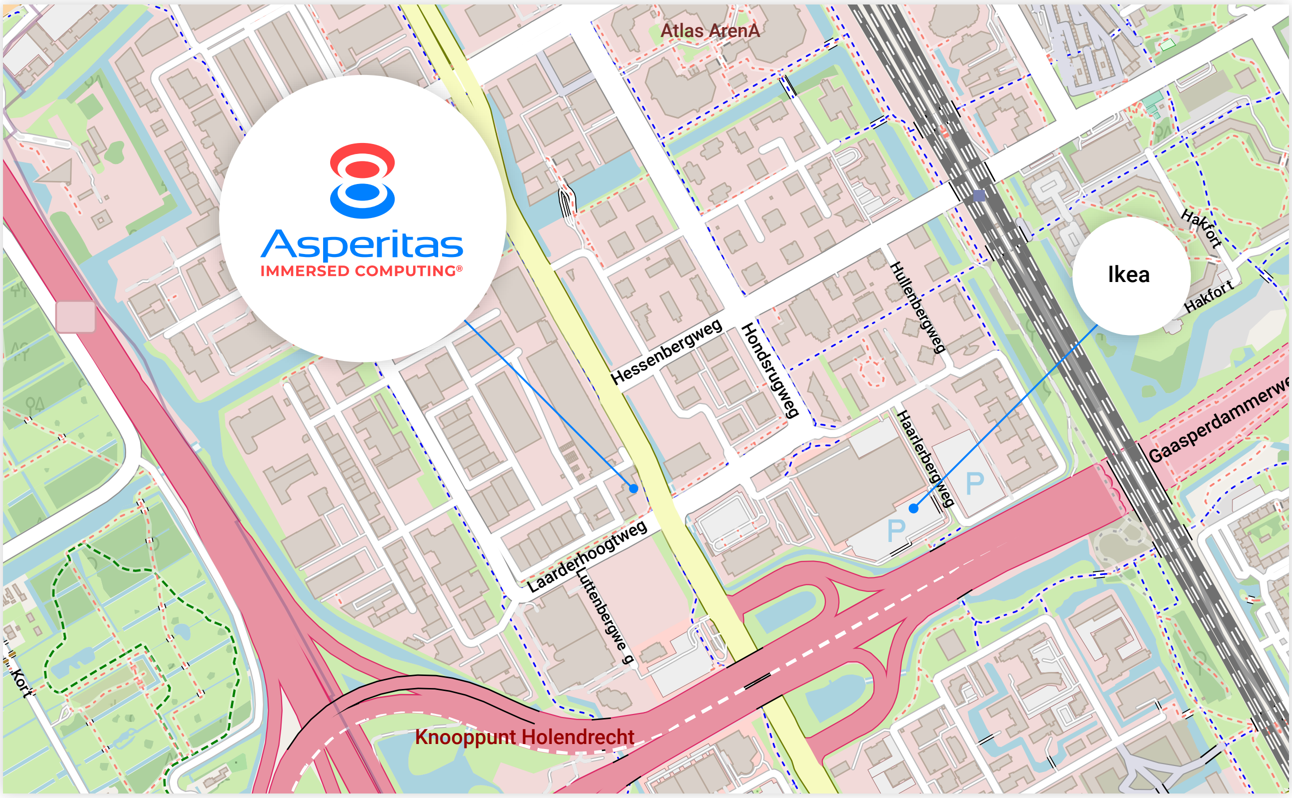

Laarderhoogtweg 18

1101 EA Amsterdam

The Netherlands

+31 88 96 000 00

info@asperitas.com

Visit www.asperitas.com