datacentrealliance
The Data Centre Trade Association

EURECA now ready for launch……..

By Steve Hone, CEO, Data Centre Alliance

WE HAVE NOW REACHED the halfway point in the EURECA project and we are reaching an exciting point in the project. For those that missed the Industry Update Seminar held at Manchester University in early July, ahead of the DCA Data Centre Transformation Conference, I will explain what this is all about and why this is probably the most important data centre industry project yet.

First some background: it is the 2nd EU funded project the DCA have participated on, this time it is a Horizon2020 project of a duration of 30 months and it has a total budget of €1.5M.

There are eight project beneficiaries: University of East London, Leeds University, The DCA Trade Association, Green IT Amsterdam, Cerios Green NL, Carbon3IT, Maki Consulting GmbH, Norland IE and Telecity GmbH.

The project is our response to the EU Commission's funding call which laid down specific challenges to stimulate market transformation towards more sustainable energy products, buildings and services within the Public Sector. This included the European Energy Efficiency Directive which requires that central governments purchase only products, services and buildings with high energy-efficiency performance.

The call asked for proposals to overcome the operational barriers related to sustainable energy public spending, such as the lack of knowledge, practical training and tailored guidelines, and also the lack of willingness to change procurement habits or perceived legal uncertainties.

So in our world of data centres this is a real challenge, but also a great opportunity to influence a market many would say needs to be 'opened up'. On top of all that, consider all that is happening in our sector: data centre KPI's, best practices and maturity models coming at us from all directions, ISO/IEC, ITU, EU Code of Conduct, CENELEC EN50600 and ETSI to name some. This is becoming an increasingly complex landscape, even the excellent EU Code of Conduct now comprises of over 150 best practices.

So in very simple terms, the projects goal is to make all of this accessible to the non-expert both in language terms and to signpost the self-improvement roadmap of data centres, server rooms and computer rooms across the European Public Sector.

The methodology and the development of the tool, which will be used by public sector procurement officers and data centre operators, is now being beta tested and by the time you read this, it will be in the final test phase ready for launch.

Based on inputted data the output of the tool will provide maturity level scoring, along with recommended actions, relevant market navigation signposting and business case construction for improving the energy efficiency of their data centre. This project is designed to help build strong business cases which will in turn initiate tenders for the procurement of products and services from all our members - be it technology, consulting, design and build, data centre services or training.

We have now presented the project to the industry in multiple regions across Europe. The next two events being in Amsterdam which I understand will coincide with the next scheduled EU Code of Conduct meeting in September, followed by Paris in November. As this project gathers pace we would like all members to help with the project by referring public sector contacts who are both already energy efficiency 'champions' to help with knowledge sharing and also those who are looking to improve the energy efficiency of their data centre.

The next phase of the project is to select 'Pilot Users' from those public sector organisations who have registered their interest in being an early adopter of the service. Those selected to be pilot users will be fully supported through the entire process with consultative advice and training all free of charge as part of the project.

Updates and opportunities to comment or contribute to the development of EURECA together with upcoming events are freely available through either the DCA website or direct via a dedicated EURECA website.

If you are a supplier then I urge you to reach out to the DCA Trade Association and register your interest in being included in the recommended public sector supplier directory which will accompany this new Public Sector service.

datacentrealliance
The Data Centre Trade Association

OPERA

By Frank Verhagen, Director, Certios BV

OPERA- lOw Power heterogeneous architecture for nExt generation of smaRt infrastructure and platforms in industrial and societal Applications

Project Aims:
OPERA project aims at supporting these ambitious challenges with technological innovation on three main aspects:

- Design next generation Low Power (LP) and Ultra-Low Power (ULP) systems
- Improve energy efficiency in computing by means of heterogeneous architectures
- Provide smart and energy efficient solutions for the interaction between embedded smart systems and remote small form-factor data centers.

The objective is to implement an innovative computing infrastructure that covers different levels of the computing continuum by means of miniaturization and integration of existing cutting-edge technologies, such as Ultra-Low Power and Low Power architectures, next generation servers, 3D integrated circuits, optical interconnections.

The main results of the project will be represented by different platforms for the Ultra-Low Power computing applications accurately selected and defined in the project, and the deployment of a scalable small form-factor data center. The usage of these platforms foreseen in the defined scenarios will be applied in traditional field as well as in new areas of deployment.

The preliminary results in the project are already highlighting new possibilities, combining the reduction of power consumption of the components used, as well as a variety of methodologies for the power harvesting which are suitable for real-life scenarios and the optimization of communication through new technological concepts, merging the new wireless network of sensors with the existing infrastructure. Looking at the data center infrastructure, the project aims at showing not only the benefits of using FPGA accelerators to increase performance/Watts ratio, but also new ways for cloud applications to exploit such enormous computational power.

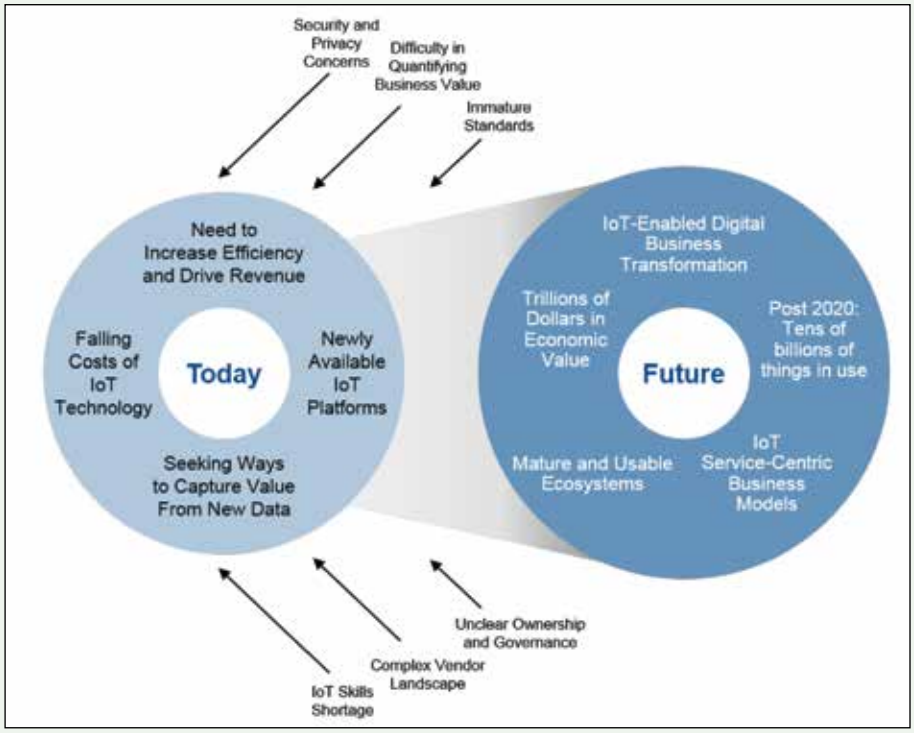

Figure 1 source: Gartner march 2016, IoT overview

The foreseen impact of the project is to drive the creation of a common working ground across key players in innovation across Europe, bridging for the first time extremely specialized skills such as Systems-on-Chip, highly performing FPGA, VHDL design, processing load optimization for multicore chips, multi-technology integration under low power constrains.

Participants:

- (Coordinator) STMicroelectronics SRL (ST) Italy
- IBM Israel - Science and Technology LTD (IBM) Israel
- HP (HP) France
- Nallatech Ltd (NALL) UK
- Istituto Superiore Mario Boella (ISMB) Italy
- Technion Israel Institute of Technology (TECH) Israel
- Consorzio per il Sistema Informativo Piemonte (CSI) Italy
- Neavia Technologies (NEAV) France
- Certios BV (CERT) NL
- Teseo SPA Tecnologie e sistemi elettronici ed ottici (TESEO) Italy
- Le Departement de l'Isère (LD38) France.

Innovation Potential of OPERA

What are the OPERA products and services actually are going to solve? Based on the usual reasons for success the identification of the markets opportunities are the following markets.

Internet of Things (IoT) analytics:
One of the markets that will benefit from Opera's "datacenter in a box" containing different ISA processors and an energy aware application management layer is the "Internet of things (IoT) analytics".
OPERA itself has a use case associated with the "internet of things" trend. A road management camera is such a device that generates a wealth of information and communicates directly with other automated systems with little to no human intervention.

The market for IoT analytics spans industrial, enterprise and consumer applications. Gartner emphasizes vendors innovative in the fields of industrial and consumer IoT analytics. In general, these vendors have novel ways to automate or make advanced analytics accessible to a broad range of

users by mining insights from streaming and other large and multi-structured data sources (relational, NoSQL, text, video, audio, social data and so on) in order to optimize industrial assets and processes and customer experiences. Their platforms typically feature connectors to, or partnerships with, device middleware vendors. They also offer industrial-domain- and asset-optimization-specific algorithms, customer-engagement-specific algorithms, analytical techniques and analytics applications.

Data and analytics leaders should consider IoT analytics solutions when they want to include an information component in an asset-based product in order to deliver more value to customers and engage with them better. They should also do so when they want to add sensors to their own assets and apply analytics to data derived from those sensors in combination with other sources in order to lower operating costs and extend asset life.

Most industrial companies are still at an early stage in their adoption of IoT analytics, and at various levels of maturity with it, as they strive to become digital companies. They tend to be conservative companies that have flat productivity and that are investing in a digital strategy and IoT analytics in an attempt to innovate again. They must embrace and use distributed data, analytics, processes, tools, skills and an expanded ecosystem to remain competitive as increasing numbers of them follow companies like GE into the digital future.

IoT analytics will be one of the industries that will benefit from the different instruction set architectures (ISA's) that are part of the OPERA solution. The expected increase in dataflow from sensors, including video, and the need to contain costs underline the targets set for Opera, doing things faster and more efficient then can currently be done.

IoT Communication

Gartner in its document "Internet of Things Primer for 2016" (Gartner, 2016), predicts that we are quickly approaching the era where tens of billions of network enabled devices are in use.

In order to sustain that many, mostly wireless, devices, both communication and the device itself must operate with extreme energy efficiency and Ultra low power (ULP).
One of the many types of IoT devices is a camera, specifically OPERA will work on a traffic management camera usable in area's that lack both electrical as well as wired communication grids.

The technology developed for OPERA includes a reconfigurable directional antenna for wireless communication and an ULP image processing addition to the camera that will limit the amount of data sent by the camera and support autonomous operation of the device for traffic safety purposes.

Environmental management

Interestingly, environmental management is a market that itself is creating an enormous amount of data collected by sensors that need to be analysed for decision support. Gartner states in an article "Cool Vendors in Sustainability, 2016" (2) that by 2020, 38% of the 6.8 billion installed IoT devices will be operating with off-grid energy sources. Experience with just such sensors, specifically with energy harvesting, energy efficient communication and data collection, is collected within OPERA through a use case where an off-grid road management camera is being developed. The technology involved, pre-processing video data and a reconfigurable antenna capable of communicating with different networks and tight radio beams will be of great use in curbing the energy demand from these billions of sensors. In this way, OPERA will contribute to a greater usability and proliferation of these sensors, supporting a better decision support system in environmental management.

Hyper converged infrastructure

Hyper converged infrastructure has, according to Gartner, a growing enterprise adoption. The same article however states: "The increasingly vibrant and congested HCIS/SDS market creates challenges for start-ups to truly differentiate themselves versus rivals"

The datacenter in a box that OPERA is creating could be at the heart of such a differentiating Hyper Converged Infrastructure Solution, offering better energy efficiency and greater capacity for application hosting then current solutions in the market.

Cloud

Cloud is a booming industry where, according to Gartner, containers are an emerging technology for delivering cloud services. Through the VDI use case, OPERA adds energy aware resource management to the mix. This knowledge will enhance and strengthen the container marketspace and through energy efficiency resulting from the resource management across different ISA's contribute to lowering the environmental footprint of clouds across Europe.

Website: **www.operaproject.eu**

The standardisation of cloud computing SLA and contractual terms: The SLALOM project

By Breda Beyer, Director of Membership, Cloud Industry Forum

SLALOM has now published a standardised set of contract terms and technical specifications for Cloud Service Agreements. The project's aim to deliver legally-worded, open and ready to use SLA contract templates and technical models that provide consistent specifications and clarify the general responsibilities and liabilities of end users and Cloud Service Providers (CSPs).

Cloud contracts remain a major stumbling block for service users, particularly smaller organisations with limited resources. The current lack of consistency makes it difficult for businesses to compare suppliers or know precisely what they are getting into when entering into Cloud service agreements and

where liabilities and responsibility for Cloud service delivery lie. This is particularly the case for SMEs who have much to gain from moving to the cloud.

In providing a standardised set of contract terms and technical specification, SLALOM can reduce that uncertainty, putting end users on firmer ground when moving to the Cloud. Such templates provide a trusted verifiable starting point for provider and business user to negotiate SLAs for doing business in the cloud in a simple, fair and transparent way.

The two main SLALOM deliverables are (1) cloud service agreement model terms and conditions; and (2) service level metric definitions based on a technical model which is aligned to ISO and allows the unambiguous definition of metrics.

The consortium working on SLALOM brought together expertise from industry and research environments: global service provider ATOS, the international legal firm Bird&Bird, the National Technical University of Athens, the Cloud Industry Forum and the University of Piraeus.

Ready to use cloud service agreement – the legal track

The SLALOM legal model is effectively a complete contract for use between a cloud user (adopter) and their provider. (Placeholders are included for situation-specific content such as the description of services, and charges.) The model is set out as a cloud service agreement (CSA) document, following a common industry approach, with more detailed content (e.g. specific service levels, privacy, and consideration) as annexes. The contract is ready to use, and can provide a starting point for creating new contracts, or it can be used as a benchmark for comparing with existing ones. It has been designed to be extendible or modifiable because it is recognized that in general variations will be needed.

The legal track, led by Bird & Bird drafted the main legal and the contractual provisions of the proposed CSA model taking into consideration the laws and regulations of Italy, Germany, France, Greece and the UK. A report on jurisprudence and case law discussing cloud computing contracts in these jurisdictions was produced as part of this exercise.

An extensive survey was conducted on the clauses commonly in use in the market examining standard CSAs used by leading

cloud service providers. it was not possible to comprehensively review the legislation and regulations that are specific to certain industry sectors (e.g. financial services) in the above-mentioned countries. However, an effort was made to focus on the main terms and conditions usually governing the relationship between the “Providers” and the “Adopters” of the services regardless of the specific market sector.

The CSA is freely available under Creative Commons license and can be downloaded from the SLALOM website.

Using metrics to improve cloud SLAs – the technical track

The SLALOM technical model has been developed by National Technical University of Athens (NTUA), incorporating revisions to reflect work which SLALOM has contributed to the ISO committee responsible for cloud standards. We established a formal liaison relationship with ISO early in the project to influence and align with ISO’s SLA work. This should mean that anyone implementing SLALOM’s recommendations will be ISO compliant.

The SLALOM reference model aims to standardise the definition of SLAs in a manner that serves the whole lifecycle of SLAs for cloud services and overcomes the shortcomings of existing approaches. It aims to do so by eliminating ambiguities in the definition and calculation of metrics and facilitating the measurement, monitoring and enforcement of SLAs to achieve non-repudiability, so that these measurements cannot be contested. Another objective has been to abstract the metric definitions as much as possible so as to facilitate the direct comparability of SLAs among providers.

The SLALOM reference model is ISO-compliant, utilising the classes and parameters of the draft ISO/IEC 19086-2 metric model, but further allows for the use of a ‘sampling’ class for concretely defining the sampling process of the metric. Furthermore, all metrics defined via the SLALOM model should be machine understandable.

In order to prove the applicability of the SLALOM model in a wide range of real world scenarios, a number of metrics directly stemming from commercial cloud providers were successfully mapped to the SLALOM model.

Through a survey conducted by the SLALOM project, feedback was collected from various stakeholders, which was used in order to prioritize cloud related SLA metrics according to their importance. Based on this prioritization, the SLALOM example SLA metrics provide usable definitions for

- availability [accessibility and functionality]
- response time [transactional and incident]
- incident resolution time
- performance of virtual cores.

The SLALOM technical model and its proposed metrics should give reasonable assurance that there is no bias towards either the cloud service provider or the cloud service adopter (i.e. customer). This is not only because of the way the metrics are specified, but also because of the metrics which are proposed. It can be downloaded from the SLALOM website.

Market reception – adoption / interest

Led by Cloud Industry Forum and the University of Pireaus, SLALOM consulted throughout the project with a wide range of

stakeholders from both research and industry to guide us in producing results which can have practical application. We presented at a number of conferences across Europe, held webinars, meetings, workshops and conducted detailed questionnaires to elicit feedback which has been published as part of the project deliverables.

Early indications are that SME providers will be the earliest users of the legal model because they have the most flexibility with respect to their legal contracting. SME providers can benefit from an independently created template contract which provides a fair and balanced baseline from which to negotiate with their customers. There is already at least one SME provider which is using the SLALOM legal model.

SLALOM legal terms have also been adopted by the Athens Bar Association (ABA) Scientific Committee and made available on subscription to more than 22,000 users (Athens lawyers and other institutional and individual users).

Based on the consultations between PICSE (Procurement Innovation for Cloud Services in Europe), Helix Nebula and SLALOM, HNSciCloud has made use of the results of SLALOM technical track in the preparation of their 5.3M€cloud service tender for the establishment of a European hybrid cloud platform, launched 21st of July (closure September 19, 2016). The purpose of the platform is to support the deployment of high-performance computing and big-data capabilities for scientific research

A number of organisations are considering using both SLALOM models for benchmarking against their current practices, which is likely to influence those practices, and we anticipate further commercial uptake as more people and organizations learn about the SLALOM legal and technical deliverables.

SLALOM is a project funded by the European Union's Horizon 2020 research and innovation programme under grant agreement no. 644270.

Energy efficient PUE: If water is the solution, how do we best solve water

By Dave Harper, Mechanical Engineer, Dunwoody LLP

WATER is often touted as the almost magical solution to data centre PUE considerations as the addition of adiabatic cooling can radically drop almost any PUE. Mains water however comes with its own problems having in itself, in many locations, significant embedded carbon and in some cases insufficient capacity to deal with extended peak temperature periods for large data centres. In order to provide a partial solution to both of those limitations we can look to rainwater capture.

The major element of rainwater is that in addition to requiring treatment and storage, it is not reliable. Worst of all there will tend to be a close correlation between the periods of highest consumption and the periods of lowest supply. In order to mitigate this, we need to look at a full year solution to water harvesting allowing winter storage to make up for summer consumption and drought.

When aiming to use water to significantly lower PUE it is likely that, for at least some years, the optimal water use will exceed the available rainwater capture. This leads to two important questions. Firstly, how do we go about identifying when the rainwater shortfall will occur? Secondly, what do we do about it when it does?

Historical data for a site lets us build up a profile of yearly temperatures and rainfalls. From that temperature data we can build up a demand curve and identify a worst case scenario for water requirement. After adding an additional margin for a future worst case year this gives us a flightpath for storage levels, any time the stored water is above the flightpath level even the worst of droughts cannot produce a problem as storage is sufficient for the remainder of the year.

When the water level drops below the flightpath then action needs to be taken. Water shortfall can be approached in two ways. In discussion with cooling equipment manufacturers a reduced water consumption regime which reduces water usage at the cost of increased energy usage can be established for for lower water levels or by means of mains supply top up.

For sites with relatively plentiful rainfall discussion with equipment manufacturers can allow the default setup to be seen essentially as a boost setting on top of the efficient reduced consumption regime. That regime will then have a secondary flightpath to identify when rainwater is in fact becoming critical. In this case or for sites where rainwater is in fact more sparse mains top up is required.

On sites where the mains supply is more than sufficient, a simple buffer can be put in place equal to the longest likely mains interruption, equalling or exceeding Uptime's twelve-hour design day storage requirement. For sites where the mains cannot be relied upon to such a degree the flight path comes into its own, a portion of the tank can be set aside for guaranteed rainwater capture with the remainder being topped up as needed to try to meet the flightpath.

Examination of past years can show how much of the tank can afford to be set aside for guaranteed capture without ever compromising having sufficient storage in the peak demand season.

Once usage and top up criteria have been established it is then possible to plot how such a regime will operate over the course of the past years. The simplest manner by which to do this is an hour by hour simulation of the tank water level changes according to the described criteria. Daily rainfall can be approximated as spread throughout the day since very few places keep records of hourly rainfall. Rain can then be considered as added to the storage tanks proportional to capture area multiplied by a capture co-efficient, since no capture system is 100% effective.

In addition to the question of water management there is also the issue of legionella prevention in the storage of large amounts of water. In the case of underground storage, the ground will act as a moderating influence. Unfortunately for many locations it is likely that the required storage volume is larger than can be cost effectively dug out. In this case where above ground storage is needed it is useful to calculate temperatures in order to identify how thorough the insulation must be both to avoid overheating in summer and to avoid freezing in winter.

In order to calculate this there are a number of other important inputs needed. Surface area and facing of the tanks, mean insolation of the site, shading of the tanks, thermal mass of the tanks, U value of the tank insulation, U value of the tank surface to air bridge, solar reflectivity and emissivity of the tanks, and temperature of mains water. Using the simulated track of tank volumes it is then possible to simulate temperature for each hourly step.

Two temperatures should be simulated, tank skin temperature and bulk water temperature. To simplify the number of calculation steps temperatures for each from the previous hour can be used rather than hunting an iterative equilibrium. Captured rain can be approximated as equal to ambient temperature and along with mains water top up and the previous tank temperature provide a baseline mixed temperature for the hour by means of weighted average. Onto that temperature the differential to the previous hour's skin temperature is combined with the contact area of the stored water to the tank skin and the insulation level to give an energy transfer to the water which then gives, along with the heat capacity of water, the temperature change over that hour. Once the water temperature is calculated that thermal transfer is then included in the skin temperature calculation along with insolation, radiation, and convective transfer to the air to give the new skin temperature. By setting up the inputs as variables acceptable values of such can be identified by means of a simple pattern search with target tank temperatures to be maintained.

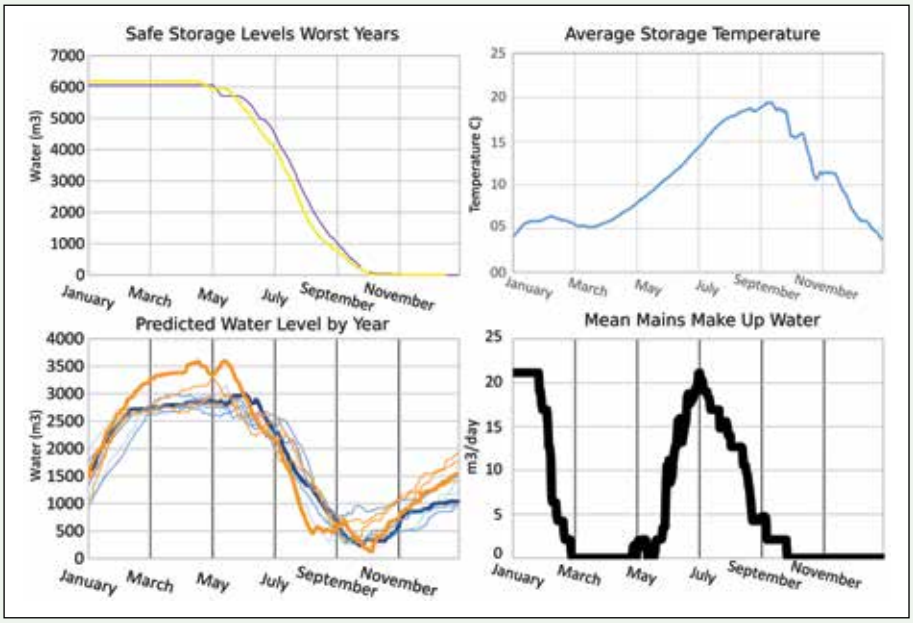

Simple spreadsheet calculation allows the pre-emptive optimisation of rainwater capture storage and usage for data centre cooling. This allows the correct strategies and equipment to be selected in advance without compromising resilience.

Liquid cooling research and development

By Dr Jon Summers, Institute of Thermofluids, School of Mechanical Engineering, University of Leeds

When you hear the words liquid cooling, what does it conjure up in your mind? We are talking about thermal management of any device or process that generates unwanted heat that if left unregulated would cause the device or process to malfunction.

Transforming high voltages generates unwanted heat and for many years liquids have been used to cool such power electronics for many reasons: one interesting reason is that the transformer is kept at a regulated temperature passively so that there is no need for moving parts as would be required if fans or p umps were employed. The heat is naturally convected away from the device and absorbed by the environment.

In the information and communication technology (ICT) sector, heat is well known to be a problem that has to be managed, whether it is from your smartphone or from the data centre that your smartphone interacts with. Liquid cooling of ICT equipment can be achieved in a myriad of ways.

A chilled water loop to the Computer Room Air Handling (CRAH) unit can be considered liquid cooling, as too would be an in-row CRAH or a rack level rear door heat exchanger on a chilled water loop, or a liquid loop to chips (microprocessors) where cooling is achieved using a very small heat sink with microchannels for the liquid coolant to collect the heat. The main difference between all of these liquid cooling solutions is their proximity to the heat source. The latter solution of on-chip liquid cooling is a situation where the facility invades the ICT space and this overlap of the two siloed domains is perhaps one of the biggest challenges to the large scale adoption of proximity liquid cooling.

On the 12th July 2016 at Datacentre Transformation (DT) Manchester (www.dtmanchester.com), one of the workshops chaired by myself was devoted to proximity liquid cooling, that being liquid cooling of ICT technologies where the liquid invades the ICT patch. Even within the realms of proximity

datacentrealliance
The Data Centre Trade Association

liquid cooling there are variations, but in fact there are really only two approaches, namely Direct Liquid Cooling (DLC) and Total Liquid Cooling (TLC). The former is where liquids are used to collect the majority of the heat from a number of discreet components, such as the Processing Units (Central and/or Graphical), the memory and perhaps the power supply, but there is still a requirement for the remainder of the heat to be collected via the air – so DLC covers those solutions where a major percentage is cooled by liquid and the remainder by air. The latter case of TLC on the other hand uses liquids to capture all of the heat, i.e. the extreme case of DLC.

The proximity liquid cooling workshop at DT Manchester was supported by four different suppliers at various stages of development. From the DLC side we were joined by Canadian based CoolIT Systems (www.coolitsystems.com), which provides a solution that uses a centrally pumped water-glycol liquid loop to a rack based manifold to collect heat mainly from the processing units using copper based cold blocks with microchannels.

Also categorised as a DLC approach is the solution from the new US based Ebullient (www.ebullientcooling.com), which also uses cold plates and rack manifolds, but with a low global warming potential (GWP), chemically engineered dielectric that has a low boiling point and therefore employs the latent heat of vaporisation of the liquid as it is pumped from one hot component to the next, in a sense trading heat for vapour pressure so that each component in the loop sees a similar temperature, with the useful attribute of being able to extract heat from more components in series than other DLC approaches that employ liquid coolants in a single phase.

The workshop was also joined by two TLC solution providers, namely UK based Iceotope (www.iceotope.com) and a new Netherlands based company called Asperitas (www.asperitas.com). Both liquid cooled solutions are based on immersion of the microelectronics into a dielectric liquid, hence being able to collect all of the heat. The approach of Iceotope is to immerse the server in an inert, low viscosity chemically engineered dielectric liquid, a similar liquid to that used by the Ebullient solution, and each server is enclosed and orientated vertically.

This approach does not require moving parts in the enclosed server as heat is transferred by natural convection to a liquid cooled cold plate on one side of the enclosure, which has a centrally pumped liquid coolant supplied through a manifold for simultaneous distribution to multiple servers. The solution by Asperitas on the other hand immerses all of the servers that would be in a single rack into a bath of high dielectric strength oil, which has cosmetic applications.

The heat from the immersed servers is harvested by the immersion of two heat exchangers that are connected to a standard chilled water loop either side of the servers with a near to full bath depth insulated separating wall that promotes natural convection circuits. As with the Iceotope solution, there are no moving parts in the dielectric liquid. After a series of ten minute presentations from the four suppliers of proximity liquid cooled ICT systems, the workshop went into an audience participated panel session. The questions from the audience teased out some interesting aspects of the different approaches. On the one hand the DLC solutions can be retrofitted to standard rack mountable ICT systems, but requires infrastructure employed for standard air cooled ICT systems. The TLC solutions however, do not breathe air, so they can be deployed outside of a standard data centre and, if well insulated, they are sufficiently quiet to deploy into human occupied workspace. All solutions, however, will require facility water to the rack, cabinet or bath, but again many data centres do have chilled water at their periphery and so perhaps we will see data centres adopt a hybrid approach of a mixture of air and liquid cooled ICT.

As an end user of these types of proximity liquid cooled solutions, one measurement is universal to all, namely that the difference between the ICT system temperature and the temperature at the point of final heat rejection is much smaller than encountered in air cooled ICT systems.

Proximity liquid cooled solutions therefore offer a greater geographical range for free cooling methods purely as a result of higher rejected temperatures even when operating the liquid cooled ICT systems at temperatures similar to air. But also, through consideration of the second law of thermodynamics, proximity liquid cooling offers a far greater opportunity for the reuse of heat by considering the so called exergy of the integrated system. No – it is not a spelling mistake; it is a quantity that is a maximum at the heat source and ideally would not diminish as it is transferred to the point of use.

Minimal drop in exergy is best achieved using proximity liquid cooled systems that are located close to where the heat could be used. We are now getting into the realms of urban based micro-data centres offering combined electrified heat and digital services with the possibility of meeting the potential future demands of the Internet of Everything and Smart Cities.