

# Addressing The Power Issue With Mears Advanced Technology Group



By [Matt Kimball](#), Contributor. © Senior Analyst, Servers  
for [Moor Insights and Strategy](#).

Follow Author

Published Feb 23, 2023, 10:42am EST



This article is more than 3 years old.

Sustainability. If you are in IT—either as an IT executive or as a solutions provider selling into IT—this is a term you are learning to embrace. ESG has become a major focus for virtually every enterprise, meaning that major contributors to ESG efforts (like, ahem, IT) are being measured on how well they are aiding a company’s sustainability efforts—in this case, the reduction of organization-wide power consumption. If your organization hasn’t yet jumped on the ESG bandwagon, get ready. Whether through free will or regulation, an ESG initiative is coming your way.

In this post, I’ll discuss the challenges I’ve been hearing about from IT executives across many industries, and spotlight an interesting solution I recently had the chance to preview.

## Do more with less power

For many businesses, IT is a unit that contributes significantly to overall power consumption—and one that’s being tasked with reducing that power footprint. IBM recently conducted a study that showed that datacenters consume about 1% of the world's power—about as much as the entire power consumption of Australia. The country. In the not-too-distant future, that number is projected to climb to 3%.

Seemingly at odds with sustainability initiatives are the digital transformation projects and edge computing expansions that IT executives are also tasked with implementing. They’re told to employ a cloud operating model and deploy micro

datacenters everywhere along the edge so more data can be gathered—data that will be fed into machine learning models so the business can be more innovative and agile. But they must also do so while reducing power consumption by a significant amount. Makes perfect sense, no?

This scenario of contradicting initiatives is what IT executives face every day, leaving many of them scratching their heads and wondering if they are being set up for failure. For those not in IT, it's important to understand that while power has always been a consideration, it has not always been treated as such a precious commodity to conserve.

### **It's also about density**

While sustainability is a big deal, there is another challenge that IT executives face, especially as they look to the edge: compute density. 174ZB of data—that's *zettabytes*—will be generated in 2025, with well over half of the average company's data being generated at the edge. This mass of data must be aggregated, transformed and analyzed in real time. And for this scenario to work, edge environments will require more than a couple of tower servers with some hard drives and a gateway. The edge requires rich compute platforms that use many high-end CPUs, GPUs and other accelerators. More than that, these servers are deployed in some of the harshest conditions—on oil rigs, at the base of cellular towers, on the battlefield and so on.

---

MORE FROM FORBES ADVISOR


### **Best Travel Insurance Companies**

By Amy Danise Editor



### **Best Covid-19 Travel Insurance Plans**

By Amy Danise Editor



---

Edge computing is the next frontier for IT; Its power, space and ruggedization barriers will continue to stall many IT projects. If only there were a way to cram a lot of compute into a small package that can also provide the necessary power and cooling. If only.

## The answer is two-phase immersion cooling

While those in the datacenter space have been experimenting with various cooling techniques over the years, I think there may be a solution to all the challenges of edge computing described above—power, cooling, space, compute density, ruggedization: two-phase immersion cooling.

Two-phase immersion cooling houses servers in tanks filled with dielectric fluid. That's right, servers run while they are immersed in fluid. And, yes, it defies conventional thinking. It works because the dielectric fluid removes heat from the server and its components, while not interfering with the server's functionality. The result is that the fluid turns into a vapor and rises. As the vapor rises, cooling coils turn it back into a fluid, which then returns to a reservoir so the cycle can continue over and over again.




MATG two-phase cooling is targeting the edge  
MOOR INSIGHTS & STRATEGY

The power savings associated with two-phase cooling seems almost too good to be true. Estimates from Mears Advanced Technology Group ([MATG](#)) put two-phase power savings in the 60% range, with overall TCO savings of up to 50%. These are staggering numbers.



MATG's modular design can reduce HVAC cooling and datacenter floorspace by up to 90%  
MOOR INSIGHTS & STRATEGY

What I find most impressive about two-phase immersion cooling is the compute density that can be achieved—which only adds to the improvement in TCO. Think about this: a single two-phase immersion bay can house 72 servers, replacing 10 standard datacenter racks.



Multiple servers are immersed in - and cooled by a dielectric fluid while running workloads  
MOOR INSIGHTS & STRATEGY

Edge computing is a market where a product like MATG's two-phase solution should find immediate success. The product is a secure container with power and networking built in that can be dropped into even a harsh field environment and simply plugged in. It can hold all of the compute capacity the edge environment will need—for today and tomorrow. On top of all the other benefits, this reduces the ever-important time-to-value metric that every IT organization cares about in today's digitally transformed setting.

I have been looking into two-phase immersion cooling, and several interesting players have entered the market. But MATG is uniquely positioned for a couple of reasons. The first is the product: an industrial design that offers incredible compute density and can scale from the edge to an HPC lab to a hyperscale datacenter.

The second is the company itself, and there are a couple of important aspects to this. For starters, MATG is part of [Quanta Services](#)—number 285 on the Fortune

500 list. That makes MATG essentially a well-funded startup that can weather any economic storm.

The other aspect is the complementary nature of MATG to Quanta's core business: industrial construction. Quanta Services works with telcos and other large customers that require these containers for ruggedized edge environments. MATG enables these customers to use a single vendor to lay the fiber, pour the concrete and deliver the micro datacenter.

### **How does two-phase cooling grow, and what does MATG need to be successful?**

In an ideal world, the datacenter infrastructure market would embrace two-phase cooling to transform itself immediately, with endless server racks being replaced by tanks. No more need for earplugs and sweatshirts in noisy, over-air-conditioned server farms.

Not so fast, though. Despite the breakneck speed at which the technology housed inside evolves, the datacenter market as a whole moves very slowly. These datacenters have a 20- to 30-year life span, and their operators will milk every bit of that to maximize the return on their considerable capital investment.

Instead, two-phase cooling will find its initial footing in edge environments and in support of specific workloads like high-performance computing that are always hungry for more compute and more workload acceleration. Like any new technology, two-phase cooling must find its early adopters to cross the proverbial chasm—and I believe the two types of deployments just mentioned will be this market's "killer apps." Next will come the hyperscale datacenters that will ultimately realize incredible TCO savings as they adopt two-phase cooling. Eventually, mainstream datacenters will act as market laggards through new datacenter builds.

How does MATG find success in this market? I believe the company's biggest challenge is the same one faced by other two-phase players: the entrenched hardware ecosystem. The trick to achieving high compute density with two-phase cooling is to use a type of server blade motherboard different from what is being offered today by big server players such as [Dell](#), [HPE](#), [Lenovo](#), and [Supermicro](#). And while there are some good designs by [Gigabyte](#), [2CRSi](#), and other server makers, the big players must jump into this market with both feet for it to thrive.

If MATG can partner with these players, the rest of the ecosystem will fall into line, from chips to components and on up the stack to software. While such enablement of the ecosystem represents a heavy lift, MATG and others are sure to find willing partners along the way. Why? Because we all know that the status quo in datacenter power and cooling is unsustainable, and that change in this market is both good and necessary. The only question is, who moves first? I have my own thoughts on that, but that's a topic for another day.

*Note: Moor Insights & Strategy writers and editors may have contributed to this article.*

Moor Insights & Strategy provides or has provided paid services to technology companies like all research and tech industry analyst firms. These services include research, analysis, advising, consulting, benchmarking, acquisition matchmaking, and video and speaking sponsorships. The company has had or currently has paid business relationships with 8x8, Accenture, A10 Networks, Advanced Micro Devices, Amazon, Amazon Web Services, Ambient Scientific, Ampere Computing, Anuta Networks, Applied Brain Research, Applied Micro, Apstra, Arm, Aruba Networks (now HPE), Atom Computing, AT&T, Aura, Automation Anywhere, AWS, A-10 Strategies, Bitfusion, Blaize, Box, Broadcom, C3.AI, Calix, Cadence Systems, Campfire, Cisco Systems, Clear Software, Cloudera, Clumio, Cohesity, Cognitive Systems, CompuCom, Cradlepoint, CyberArk, Dell, Dell EMC, Dell Technologies, Diablo Technologies, Dialogue Group, Digital Optics, Dreamium Labs, D-Wave, Echelon, Ericsson, Extreme Networks, Five9, Flex, Foundries.io, Foxconn, Frame (now VMware), Fujitsu, Gen Z Consortium, Glue Networks, GlobalFoundries, Revolve (now Google), Google Cloud, Graphcore, Groq, Hiregenics, Hotwire Global, HP Inc., Hewlett Packard Enterprise, Honeywell, Huawei Technologies, HYCU, IBM, Infinidat, Infoblox, Infosys, Inseego, IonQ, IonVR, Inseego, Infosys, Infiot, Intel, Interdigital, Jabil Circuit, Juniper Networks, Keysight, Konica Minolta, Lattice Semiconductor, Lenovo, Linux Foundation, Lightbits Labs, LogicMonitor, LoRa Alliance, Luminar, MapBox, Marvell Technology, Mavenir, Marseille Inc, Mayfair Equity, Meraki (Cisco), Merck KGaA, Mesosphere, Micron Technology, Microsoft, MiTEL, Mojo Networks, MongoDB, Multefire Alliance, National Instruments, Neat, NetApp, Nightwatch, NOKIA, Nortek, Novumind, NVIDIA, Nutanix, Nuvia (now Qualcomm), NXP, onsemi, ONUG, OpenStack Foundation, Oracle, Palo Alto Networks, Panasas, Peraso, Pexip, Pixelworks, Plume Design, PlusAI, Poly (formerly Plantronics), Portworx, Pure Storage, Qualcomm, Quantinuum, Rackspace, Rambus, Rayvolt E-Bikes, Red Hat, Renesas, Residio, Samsung Electronics, Samsung Semi, SAP, SAS, Scale Computing, Schneider Electric, SiFive, Silver Peak (now Aruba-HPE), SkyWorks, SONY Optical Storage, Splunk, Springpath (now Cisco), Spirent, Splunk, Sprint (now T-Mobile), Stratus Technologies, Symantec, Synaptics, Syniverse, Synopsys, Tanium, Telesign, TE Connectivity, TensTorrent, Tobii Technology, Teradata, T-Mobile, Treasure Data, Twitter, Unity Technologies, UiPath, Verizon Communications, VAST Data, Ventana Micro Systems, Vidyo, VMware, Wave Computing, Wellsmith, Xilinx, Zayo, Zebra, Zededa, Zendesk, Zoho, Zoom, and Zscaler. Moor Insights & Strategy founder, CEO, and Chief Analyst Patrick Moorhead is an investor in dMY Technology Group Inc. VI, Fivestone Partners, Frore Systems, Groq, MemryX, Movandi, and Ventana Micro.

Editorial Standards

Reprints & Permissions