

WHAT'S COOLER THAN COOL? —

Bless the overclockers: In the data center world, liquid cooling is becoming king

"For high performance computing, you just can't do it with air."

ANDY PATRIZIO - 2/4/2019, 9:05 PM



Look, everyone loves John Slattery. But this is relevant, promise.

In *Iron Man 2*, there is a moment when Tony Stark is watching a decades-old film of his deceased father, who tells him "I'm limited by the technology of my time, but one day you'll figure this out. And when you do, you will change the world." It's a work of fiction but the notion expressed is legitimate. The visions and ideas of technologists are **frequently well ahead** of the technology of their times. *Star Trek* may have always had it, but it took the rest of us decades to get tablets and e-readers right.

The concept of liquid cooling sits squarely in this category as well. While the idea has been around since the 1960s, it remained a fringe concept when compared to the much cheaper and safer air cooling method. It took another 40 odd years before liquid cooling even started to take off in the 2000s, and then it was mostly confined to PC hobbyists who wanted to overclock their CPUs well beyond the **recommended limits set by Intel** and AMD.

Today, however, liquid cooling seems to be having a moment. You can buy a liquid cooling system for your PC for under \$100, and a whole cottage industry of enterprise and data center vendors (like CoolIT, Asetek, Green Revolution Computing, Ebullient, just to name four) are all promoting liquid cooling of data center equipment. Liquid cooling continues to be primarily used in areas of supercomputing, high performance computing (HPC), or other situations involving massive amounts of compute power where CPUs run at almost 100 percent utilization, but such a use case is becoming mainstream.

There are two common types of liquid cooling: direct to chip and immersion. In direct to chip, a heat sink is attached to the CPU just like with a standard fan, but instead of a fan there are two tubes connected. It has one tube where cold water comes in and cools the heat sink, which is absorbing heat from the CPU, and one pipe to take the hot water away. It is then cooled and returned to the CPU in a closed loop not unlike the human blood stream.

With immersion, the hardware is flooded with a liquid bath, which must obviously be non-conductive. Generally, this approach can best be compared to the pools used to cool nuclear reactor rods. Immersion is much more cutting edge and requires much more expensive coolant than direct to chip, which can still use plain-old water. On top of that, there is the risk of spillage with immersion. For that reason, direct to chip is much more popular for now.

For one major example, take Alphabet. When Google's parent company Alphabet introduced its TensorFlow 3.0 AI processors in May 2018, CEO Sundar Pichai **said** the chips are so powerful, "that for the first time we've had to introduce liquid cooling in our data centers." The switch was the price Alphabet paid for an eight-fold improvement in performance.

On the flip side, Skybox Datacenters recently announced a massive, **40,000 server supercomputer** for oil and gas exploration from DownUnder GeoSolutions (DUG). This initiative is expected to deliver 250 petaflops of computing power, more than any existing supercomputer—and the expected liquid cooling system would enclose the servers, using more than 720 enclosures in fully submerged tanks filled with dielectric fluid.

Either way, "Liquid cooling is the cooling of the future and always will be," said Craig Pennington, vice president of design engineering at data center operator Equinix. "It seems so obvious that's the right thing to do but no one has done it."

So how has liquid cooling gone from esoteric science on the fringe of computing to near-mainstream in modern data centers? Like all technologies, it was partially the result of evolution involving trial and error and a lot of engineering. But specifically for liquid cooling, today's massive data center operators should perhaps thank the early overclockers, who may really be the unsung heroes of liquid cooling.



The IBM System 360 data processing control panel. #VintageChic

What we talk about when we talk about liquid cooling

Liquid cooling really entered the popular imagination back in 1964 when IBM explored immersion cooling for the company's **System 360 mainframe**. (The System 360 was one of the company's first mainframe computers; while IBM had 700 and 7000 Series mainframes for more than a decade, the System/360 "ushered in an era of computer compatibility—for the first time, allowing machines across a product line to work with each other," according to **IBM**.) The concept was simple: chilled water would be run through a contraption that cooled the water down to below room temperature, and then the water would be supplied directly to the system. IBM's ultimate setup used what's now known as rear door cooling, where a radiator-like device was mounted on the back of a mainframe. This device drew in hot air from the mainframe via fans, and that air was then cooled by the water, much like how a radiator cools a car engine.

Over the time since, engineers improved upon this basic concept, and two dominant forms of liquid cooling ultimately emerged—immersion and direct contact. Immersion is exactly what it says: the electronics are sitting in a liquid bath that, for obvious reasons, cannot be water. The liquid must be non-conductive, or dielectric (companies like [3M](#) even engineer fluid specifically for this purpose).



FURTHER READING

50 years ago, IBM created mainframe that helped send men to the Moon

Immersion, though, has a lot of challenges and drawbacks. Because it sits in liquid, the server can only be accessed from the top. That's where external ports must be located. A 1U ([rack unit](#)) solution is impractical, so you can't stack the racks. The dielectric fluid, usually mineral oil, is very expensive and can be messy to clean up if it spills. You need special hard drives and will likely have to spend a lot to retrofit a data center. That's why in the case of the Houston supercomputer mentioned above, immersion is best done with a new data center rather than retrofitting an old one.

By contrast, direct contact liquid cooling is where a heat sink (or heat exchanger) sits on the chip, just like a regular heat sink. Instead of an attached fan blowing on the heat sink, however, this setup has two water pipe connections—one to bring cool water in to cool the plate, and another to take away the hot water created by contact with the heat plate. This has become the most common form of liquid cooling, adopted by major OEMs like HP Enterprise, Dell EMC, and IBM, as well as cabinet and enclosure makers like Chatsworth Systems and Schneider Electric.

Direct to chip uses water, though it's rather particular about the quality of water. You certainly can't use unfiltered municipal water. Just look at your faucet or shower head. Who wants a calcium buildup in their servers? At the very least, direct contact liquid cooling requires pure, distilled water, and sometimes water mixed with an antifreeze. This type of liquid coolant is a [very precise science](#) in and of itself.

The screenshot shows the Ars Technica website interface. On the left is a navigation menu with links like 'BUYER'S GUIDE', 'HOW-TOs & TWEAKS', 'PRODUCT REVIEWS', 'TECHNICAL BLACKPAPERS', 'CPU & CHIPSET GUIDE', 'THE FORUM', 'SEARCH ARS', 'WANKERDESK', 'ASK ARS!', 'DIARY OF A GEEK', 'GAME.ARS', 'WHO WE ARS', 'ADVERTISING', and 'LINKS'. The main content area features an article titled 'A brief history of clock' by Dr. Evil, with a sub-headline 'In which the history and mechanics of overclocking Intel CPUs is explored'. The article text discusses Intel's clocking and locking articles and mentions the 486 DX2 processor. To the right is a 'Wind it up' sidebar with a 'Recent:' section listing articles such as 'Game.Ars', 'Addtronics 7896 w/ modifications', 'Pentium 4 and G4e: an Comparison', 'Quantum Computing and PR', 'Updated System Guide', 'Ultimate Limits of...', and 'Cappuccino PC'. At the bottom of the sidebar is a link for 'A+ Certification Overview'.

[Enlarge](#) / Hmmm, it seems some of those early 2000s overclockers that helped spawned the liquid cooling renaissance may have [read A.R.S. Technical](#).

The Intel connection

How did we get from IBM's radiators to today's extravagant cooling solutions? Again, thank the [overclockers](#). Around the turn of the century, liquid cooling started catching on with PC overclockers and system builder hobbyists who wanted to run their computers at higher speeds than officially rated. Still, it was a very esoteric art with no standard design. Everyone did their own thing. It required user assembly MacGyvering that put Ikea products to shame. Most of the coolers didn't even fit well in the case.

In early 2004, things changed due to some internal politicking at Intel. An engineer from the company's Hillsboro, Oregon, design center—where most of the company's chip design work is done, despite Intel HQ being in Santa Clara, California—had been working on a custom cooling project for Intel for several years. The project had cost Intel more than \$1 million to develop at that time, and its aim was a liquid cooler for Intel CPUs. Unfortunately, Intel was about to kill the project.

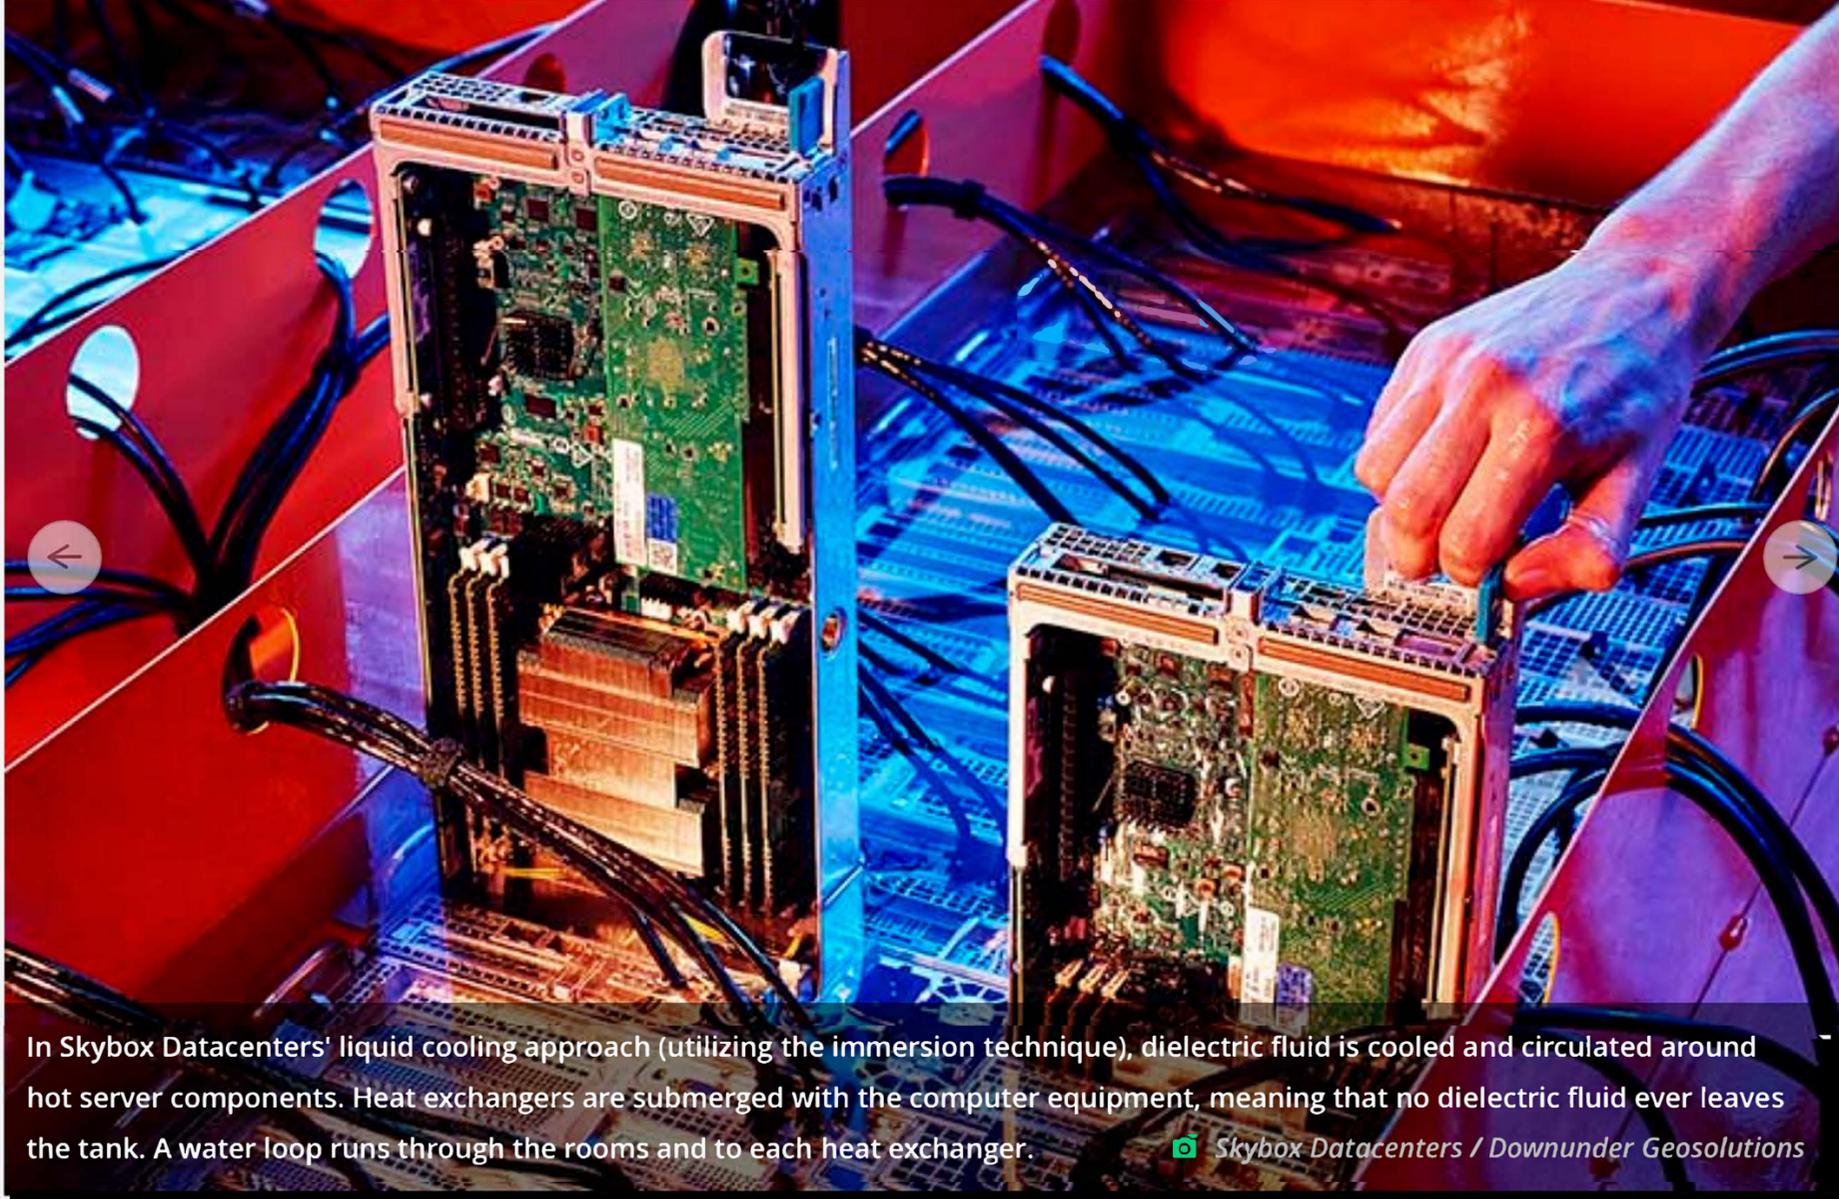
The engineer hoped for a different result. To save his project, he brought the idea to **Falcon Northwest**, a builder of top-of-the-line systems for gamers in Portland. “The reasoning was as a company, they saw liquid cooling as a tacit endorsement of overclocking, and overclocking back then was verboten,” said Kelt Reeves, president of Falcon Northwest. Intel had logical reasons for this stance. At the time, unscrupulous retailers in Asia were selling overclocked PCs at higher-than-rated clock speeds with poor cooling, and this somehow became Intel’s problem in the public discourse. Thus, the company opposed overclocking.

But this Oregon engineer thought if he could get a customer and market case for the cooler, Intel would ultimately relent. (What Intel had built also happened to be a far better solution than what was on the market elsewhere, Reeves told Ars.) So after a little internal advocating and negotiating between companies, Intel allowed Falcon to sell the cooling systems—partially since Intel had already produced thousands of them. The only catch? Falcon couldn’t acknowledge Intel’s involvement. Falcon agreed, and soon it became **the first PC maker** to ship a fully sealed, all-in-one-liquid cooler.

That pioneering modern liquid cooling solution wasn’t exactly consumer friendly, Reeves noted. Falcon had to modify cases to make the radiator fit and had to invent a cold plate to cool the water. But over time, CPU cooler makers like ThermalTake and Corsair studied what Intel did and aimed for incremental improvements. From there, server products and vendors like CoolIT and Asetek sprung up to bring liquid cooling specifically to the data center. Some of what they do—like the tubing that will not break, crack, or leak for up to seven years—was eventually licensed to consumer CPU cooler vendors, and this sharing of advancements back and forth has become the norm.

As this market started growing in multiple ways, even Intel eventually changed its tune. It now markets the overclocking capabilities of the K and X series of its CPUs, for instance, and the company doesn’t even bother to provide a stock fan with its top-of-the-line CPUs for gamers.

“[Liquid cooling] is tried and true—everyone on the consumer side is doing it,” Reeves said. “Intel stopped shipping stock fans with high-end CPUs because they require liquid cooling; it’s been proven by and even blessed by Intel at this point. I don’t think you will find anyone who would say an all-in-one is not reliable enough.”



In Skybox Datacenters' liquid cooling approach (utilizing the immersion technique), dielectric fluid is cooled and circulated around hot server components. Heat exchangers are submerged with the computer equipment, meaning that no dielectric fluid ever leaves the tank. A water loop runs through the rooms and to each heat exchanger. [Skybox Datacenters / Downunder Geosolutions](#)



The case for liquid cooling—practicality

For a long time, traditional data center design involved a raised floor with tiny holes in it so cold air could come up and be sucked in to the mainframe or the server cabinet. This was known as a CRAC unit, for computer room air conditioner. The problem is blowing cold air through pinholes in the floor is simply not enough anymore.

Truly, the primary reason for the recent boom in the liquid cooling industry is necessity. Today's CPUs are too hot and servers are packed too densely for air to effectively cool them anymore, as Google itself noted. Water has approximately 3,300 times the heat capacity than air, and a water cooled system can move 300 liters of water per minute, as opposed to 20 cubic feet of air per minute.

To state it plainly, water can cool much more efficiently and in a much smaller space. So after years of trying their best to keep power consumption down, CPU vendors can now throw power to the wind and crank up the voltage to get maximum performance—they know liquid cooling can handle it.

"The chip level power we are being asked to cool is accelerating past 500 watts," said Geoff Lyon, CEO of CoolIT. "Some CPUs not on the market yet will be hitting 300 watts. This is all being driven by **AI and machine learning**. It can't grow fast enough."

Lyon said CoolIT is looking at extending cooling to chipsets, power conditioning, networking chips, and memory. "It's not that big a lift to go the extra distance and do RAM, too," he added. "There's flavors of RAM with advanced packaging that are 18 watts per DIMM. The typical DIMM is four to six watts. In a high memory system, we see servers with 16, 24 DIMMs, which is a lot of heat."



FURTHER READING
The AI revolution has spawned a new chips arms race

Vendors all over are seeing this increasingly hot demands. Equinix has seen its average densities move from five kilowatts to seven or eight kilowatts and now to 15-16 watts, with some equipment in the 40 kilowatt density. "So the total volume of air to be moved through there is too vast. It isn't happening immediately, but the next two years will be the fundamental adoption of liquid cooling," said Equinix's Pennington.

The case for liquid cooling—energy efficiency

Energy consumption has been a concern for the data center industry for a while now (the [EPA](#) has been monitoring this for at least a decade). Today's data centers are massive facilities that consume an estimated [two percent of the global electrical power](#) worldwide and emit as much CO2 as the airline industry. There remains an overwhelming energy concern regarding this. Luckily, liquid cooling helps reduce the power bill.

The first savings comes from turning off the air conditioners in the data center. Second is the elimination of fans. Every server rack from 1U up has multiple fans in them to move air, but that can be reduced to little or none, depending on the density.

And with a technique known as dry cooling, which omits refrigeration, even greater savings can be attained. Originally, direct to chip cooling circulated the water through a refrigerator which cooled the water way down, to about 15-25 degrees Celsius. But it was eventually observed that liquid coolers just let the water snake through a long set of piping and fans cooled the pipes as they heated up from the hot water, natural thermal diffusion would also cool the water down enough to be effective.

“Because it’s so efficient, you don’t have to worry about water temperature being as cold as they could be,” said Equinix’s Pennington. “Warm water cooling still captures all the heat from the servers rather effectively. You don’t need compression cycle, you can just use dry coolers.”

Dry coolers also save water. A large data center can consume millions of gallons of water per year if it uses refrigeration, but a data center using dry coolers doesn’t use any water. That saves power *and* water, which can be very helpful if the data center is in a municipal area.

“We’re not taking in lots of water,” said Pennington. “If you design it carefully, it’s a closed loop system. There is no water flows in and out, just make up water to keep the tanks topped off about once a year. You don’t top off the water in your car all the time, you wouldn’t do it here either.”

Where performance lives, adoption follows

As just one real-world example: after adopting liquid cooling techniques, Dell observed up to a 56 percent improvement in PUE, or [Power Usage Effectiveness](#), according to Brian Payne, vice president of product management and marketing for PowerEdge at Dell EMC. The PUE is the ratio of power needed to cool the systems versus the power needed to run them. A PUE of 3 means it takes twice as much power to cool the systems as to run them, while a PUE of 2 means an equal amount of power is needed to both run and cool the systems. A PUE of 1 is impossible because *some* kind of cooling is needed, but getting the PUE down to as close to 1.0 is a near obsession among data center operators.

In addition to the PUE, Dell sees its liquid-cooling adoptive customers getting up to 23 percent more compute power without having to overhaul the cooling system. “Based on the investment required to deliver it, we forecast a one-year ROI,” Payne says. “I would liken it to buying a higher efficiency air conditioner for your home. You invest a little bit but see the payback in the power bill over time.”

A QUICK WORD ON IMMERSION

Green Revolution Cooling is firm focusing on liquid immersion rather than direct to chip, and CEO Peter Poulin says there are two reasons immersion in particular better addresses energy efficiency. First, you remove all the fans from servers when go the immersion route. By simply removing fans from servers, you reduce power servers draw on average by 15 percent. One Green Revolution customer reduced their power draw by 30 percent.

Tangentially, there is another benefit to removing the fans: quiet. Despite often using very small fans, servers are unbelievably loud and being in a data center can be unpleasant from both a heat and noise perspective. Liquid cooling makes them a much more pleasant place to work.

The next big benefit is how little power is required to run the immersion cooling system. It has just three moving parts: a pump to circulate the coolant, a pump to move coolant to coolant tower, and fan for cooling tower. When replacing air, this can mean the electric load drops to just five percent of what it was on air conditioning. “So you get this massive reduction in energy consumption, which enables a whole lot of other things,” said Poulin. “Depending on the customer, they can be more energy friendly or reduce the carbon footprint associated with building more data centers.”

To consider a completely different new liquid-cooling believer, take the Ohio Supercomputer Center (OSC). It runs 1,800 nodes total in a cluster. But after adopting liquid cooling, Doug Johnson, chief systems architect and HPC systems group manager, said the center has hit a PUE of 1.5. OSC uses an outside loop so the water is routed out of the building to be cooled by the natural ambient temperature, which averages 30C in the summer and a lot less in winter. And with the chips hitting 70C, even if the water hits 40C, that's still a lot cooler than the chip and serves its purpose.

Like so many early adopters, OSC was completely new to this. It went from not doing any liquid cooling five years ago to 25 percent cooling today. It hopes to have 75 percent liquid cooling in three years and be fully liquid cooled a few years after that. Even with today's reduced setup, Johnson said it took four times as much power to cool the systems before moving to liquid cooling, and he claims the solution has reduced the overall power consumption by two-thirds. "I think that percentage will only increase in time as we get GPUs integrated into the cooling system."

From any customer's standpoint, it takes time and energy to evaluate any newer technology—it's why big companies like Dell have even partnered with CoolIT when it comes to making the liquid cooling sales pitch. Not surprisingly, fear of leaks remains the first thing customers tend to bring up. But hesitations aside, it turns out there isn't much of a choice as present if you want the best performance.

"That leak fear has always existed," said CoolIT's Lyon. "What's changed is you don't have other options. For high performance computing, you just can't do it with air."

Andy Patrizio is a freelance technology journalist based in Orange County, California, not entirely by choice. He prefers building PCs to buying them, has played too many hours of Angry Birds on his iPhone, and collects old coins when he has some to spare.

*Note: This article originally referred to the IBM System 360 as the company's first mainframe computer, but the text has been corrected to reflect predecessors like the **700 and 7000 Series mainframes**.*