



The Immersion Cooling Authority

National Security Agency

Doing More with Less: Cooling Computers with Immersion Pays Off

Republished by GRC — 2018

Author:
David Prucnal, PE

Original Publication:
The Next Wave | Vol. 20 | No. 2 | 2013



DOING MORE WITH LESS: COOLING COMPUTERS WITH IMMERSION PAYS OFF

David Prucnal, PE

A consequence of doing useful work with computers is the production of heat. Every watt of energy that goes into a computer is converted to a watt of heat that needs to be removed, or else the computer will melt, burst into flames, or meet some other undesirable end. Most computer systems in data centers are cooled with air conditioning, while some high-performance systems use contained liquid cooling systems where cooling fluid is typically piped into a cold plate or some other heat exchanger.

Immersion cooling works by directly immersing IT equipment into a bath of cooling fluid. The National Security Agency's Laboratory for Physical Sciences (LPS) acquired and installed an oil-immersion cooling system in 2012 and has evaluated its pros and cons. Cooling computer equipment by using oil immersion can substantially reduce cooling costs; in fact, this method has the potential to cut in half the construction costs of future data centers.



Network servers are submerged into a tank of mineral oil. (Photo used with permission from Green Revolution Cooling: grcooling.com.)

THE FUNDAMENTAL PROBLEM

Before getting into the details of immersion cooling, let's talk about the production of heat by computers and the challenge of effectively moving that heat from a data center to the atmosphere or somewhere else where the heat can be reused.

In order for computers to do useful work, they require energy. The efficiency of the work that they do can be measured as the ratio of the number of operations that they perform to the amount of energy that they consume. There are quite a few metrics used to measure computer energy efficiency, but the most basic is operations per watt (OPS/W). Optimizing this metric has been the topic of many PhD theses and will continue to be the subject of future dissertations. Over the years, there has been progress against this metric, but that progress has slowed because much of the low-hanging fruit has been harvested and some of the key drivers, Moore's Law and Denard scaling, have approached the limits of their benefit. Improvements to the OPS/W metric can still be made, but they usually come at the expense of performance.

The problem is not unlike miles per gallon for cars. The internal combustion engine is well understood and has been optimized to the nth degree. For a given engine, car weight, and frontal area, the gas mileage is essentially fixed. The only way to improve the miles per gallon is to reduce the performance or exploit external benefits. In other words, drive slower, accelerate less, drift down hills, find a tailwind, etc. Even after doing all of these things, the improvement in gas mileage is only marginal. So it is, too, with computers. Processor clock frequencies and voltages can be reduced, sleep modes can be used, memory accesses and communications can be juggled to amortize their energy costs, but even with all of this, the improvement in OPS/W is limited.

A natural consequence of doing useful work with computers is the production of heat. Every watt of energy that goes into a computer is converted into a watt of heat that needs to be removed from the computer, or else it will melt, burst into flames, or meet some other undesirable end. Another metric, which until recently was less researched than OPS/W, is kilowatts per ton (kW/ton), which has nothing to do with the weight of the computer system that is using up the energy. Here, ton refers to an amount of air conditioning; hence, kW/ton has to do with the amount

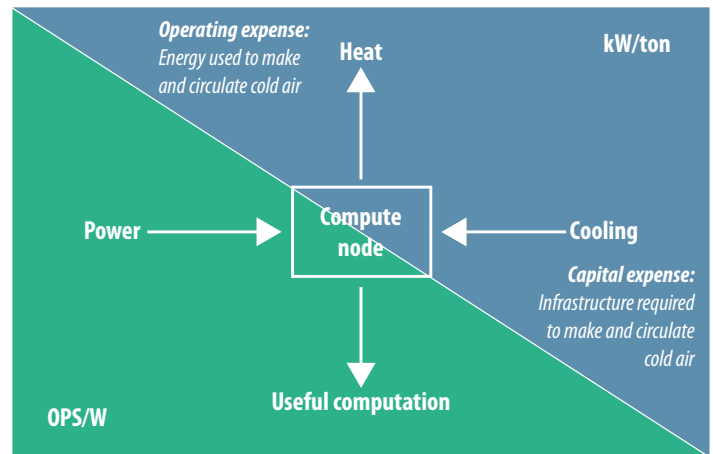


Fig. 1. There are two halves to the computer power efficiency problem: efficiency of the actual computation (green sector) and efficiency of the cooling infrastructure (blue sector).

of energy used to expel the heat that the computer generates by consuming energy (see figure 1).

In fact, many traditional data centers consume as much energy expelling heat as they do performing useful computation. This is reflected in a common data center metric called power usage effectiveness (PUE), which in its simplest form is the ratio of the power coming into a data center to the power used to run the computers inside. A data center with a PUE of 2.0 uses as much power to support cooling, lighting, and miscellaneous loads as it does powering the computers. Of these other loads, cooling is by far the dominant component. So, another way to improve data center efficiency is to improve cooling efficiency. The best case scenario would be to achieve a PUE of 1.0. One way to achieve this would be to build a data center in a location where the environmental conditions allow for free cooling. Some commercial companies have taken this approach and built data centers in northern latitudes with walls that can be opened to let in outside air to cool the computers when the outside temperature and humidity are within allowable limits. However, for those of us who are tied to the mid- Atlantic region where summers are typically hot and humid, year-round free cooling is not a viable option. How can data centers in this type of environment improve their kW/ton and PUE?

HOW COMPUTERS ARE COOLED

There are many different ways that computers are kept cool in data centers today; however, the most

common method is to circulate cool air through the chassis of the computer. Anyone who has ever turned on a computer knows that the computer makes noise when it is powered on. Central processing units (CPUs), memory, and any other solid-state components are completely silent, so what makes the noise? Spinning disk drives can make a little noise, but by far, the dominant noisemakers are the fans that are used to keep air moving across the solid-state devices that are all busily doing work and converting electrical power input into computation and heat. Even the power supply in a computer has a fan because the simple act of converting the incoming alternating current (ac) power to usable direct current (dc) power and stepping that power down to a voltage that is usable by the computer creates heat. All of the fans in a computer require power to run, and because they are not perfectly efficient, they too create a little heat when they run. The power used to run these fans is usually counted as computer load, so it ends up in the denominator of the PUE calculation, even though it does nothing toward actual computation.

But how do all of these fans actually cool the computer? Think of cooling as heat transfer. In other words, when an object is cooled, heat is transferred away from that object. What do people do when they burn their finger? They blow on it, and if they are near a sink, they run cold water on it. In both cases they are actually transferring heat away from their burnt finger. By blowing, they are using air to push heat away from their finger, and by running water, they are immersing their hot finger in a cool fluid that is absorbing and carrying the heat away. Anyone who has burned a finger knows that cold water brings much more relief than hot breath. But why? The answer depends on principles like thermal conductivity and heat capacity of fluids. It also helps to understand how heat moves.

HEAT ON THE GO: RADIATION, CONDUCTION, CONVECTION, AND ADVECTION

Imagine a campfire on a cool evening. The heat from the fire can be used to keep warm and to roast marshmallows, but how does the heat move from the fire? There are three modes of heat transfer at work around a campfire: radiation, conduction, and convection (see figure 2). As you sit around the fire, the heat that moves out laterally is primarily radiant heat. Now, assume you have a metal poker for stirring the coals and moving logs on the fire. If you hold the poker in the fire too long it will start to get hot in your




Fig. 2. There are three modes of heat transfer at work around a campfire: radiation, conduction, and convection.

hand. This is because the metal is conducting heat from the fire to your hand. To a much lesser extent the air around the fire is also conducting heat from the fire to you. If you place your hands over the fire, you will feel very warm air rising up from the fire. This heat transfer, which results from the heated air rising, is *convection*. Now, if an external source, such as a breeze, blows across the fire in your direction, in addition to getting smoke in your eyes, you will feel heat in the air blowing towards you. This is *advection*. In a computer, a CPU creates heat that is typically conducted through a heat spreader and then into the surrounding air. Convection causes the air to rise from the heat spreader, where it is then blown, or advected, away by the computer's cooling fan. Now that we know how heat moves, why is it that it feels so much better to dunk a burnt finger in water than to blow on it? This is where thermal conductivity and heat capacity of the cooling fluid come into play. First, a few definitions:

- ▶ **Thermal conductivity** is the ability of a material to conduct heat; it is measured in watts per meter degree Celsius, or $W/(m \cdot ^\circ C)$.
- ▶ **Heat capacity** is the amount of heat required to change a substance's temperature by a given amount or the amount of heat that a substance can absorb for a given temperature increase; it is measured in joules per degree Celsius ($J/^\circ C$).
- ▶ **Specific heat capacity** is the heat capacity per unit mass or volume; it is typically given per unit mass and simply called specific heat (C_p); it is measured in joules per gram degree Celsius, or $J/(g \cdot ^\circ C)$.

The answer to why it feels so much better to dunk a burnt finger into water than to blow on it can be found in table 1. First, water is a much better conductor of heat than air, by a factor of 24. Think of it as having 24 times more bandwidth for moving heat. Second, water can hold far more heat than air. In fact, 3,200 times more. So, water provides 24 times more heat transfer bandwidth and 3,200 times more heat storage than air. No wonder the finger feels so much better in the water.

One more thing to consider about heat transfer; heat naturally flows from hot to cold, and the rate of heat transfer is proportional to the temperature difference. This is why the colder the water, the better that burnt finger is going to feel.

COOLING COMPUTERS

By now it should be apparent that the fans in a computer are there to advect (i.e., move) a cooling fluid (e.g., air) across the heat producing parts (e.g., CPUs, memories, and peripheral component interconnect cards) so that the cooling fluid can absorb heat through conduction and carry it away. This can be described by the following mass flow heat transfer equation: $Q = mc_p \Delta T$

In this equation, Q is the rate of heat transfer in watts, m is the mass flow rate of the cooling fluid in grams per second, c_p is the specific heat of the cooling fluid, and ΔT is the change in temperature of the cooling fluid. What it says is that the cooling depends on production of the amount of coolant flowing over the heat source, the ability of the coolant to hold heat, and the temperature rise in the coolant as it flows across the heat source.

How much air does it take to keep a computer cool? There is a rule of thumb used in the data center design

world that 400 cubic feet per minute (CFM) of air is required to provide 1 ton of refrigeration. One ton of refrigeration is defined as 12,000 British thermal units per hour (Btu/h). Given that 1 kilowatthour is equivalent to 3,412 British thermal units, it can be seen that a ton of refrigeration will cool a load of 3,517 W, or approximately 3.5 kW. The mass flow heat transfer equation can be used to confirm the rule of thumb. Air is supplied from a computer room air conditioning (CRAC) unit in a typical data center at about 18°C (64°F). Now, 400 CFM of air at 18°C is equivalent to 228 grams per second, and the specific heat of air is equivalent to 1 J/(g·°C). Solving the mass flow heat transfer equation above with this information yields a change in temperature of 15°C. What all this confirms (in Fahrenheit) is that when 64°F cooling air is supplied at a rate of 400 CFM per 3.5 kW of computer load, the exhaust air from the computers is 91°F. Anyone who has stood in the “hot aisle” directly behind a rack of servers will know that this rule of thumb is confirmed.

It is not unusual for a server rack to consume over 10 kW. Using the rule of thumb above, a 10.5 kW server rack requires 1,200 cubic feet of cooling air— enough air to fill a 150 square foot office space with an 8 foot ceiling—per minute. That’s a whole lot of air! Simply moving all of that air requires a significant amount of energy. In fact, for racks of typical one-unit (1U) servers, the energy required to move cooling air from the CRAC units and through the servers is on the order of 15% of the total energy consumed by the computers. Remember—this is just the energy to move the cooling air, it does not include the energy required to make the cold air. If there was a way to cool computers without moving exorbitant quantities of air, it could reduce energy consumption by up to 15%. This may not seem like much, but consider that a 15% improvement in OPS/W is almost unheard of,

	Thermal Conductivity, W/(m·°C) at 25°C	Specific Heat (C_p), J/(g·°C)	Volumetric Heat Capacity (C_v), J/(cm ³ ·°C)
Air	0.024	1	0.001297
Water	0.58	4.20	4.20
Mineral Oil	0.138	1.67	1.34
Aluminum	205	0.91	2.42
Copper	401	0.39	3.45

Table 1. Thermal conductivity and heat capacity of common substances

and for a moderate 10 megawatt (MW) data center, a 15% reduction in energy consumption translates into a savings of \$1.5 million per year.

PUMPING OIL VERSUS BLOWING AIR

Unfortunately, we cannot dunk a computer in water like a burnt finger since electricity and water do not play well together. Mineral oil, on the other hand, has been used by electric utilities to cool electrical power distribution equipment, such as transformers and circuit breakers, for over 100 years. Mineral oil only has about 40% of the heat holding capacity and about one quarter the thermal conductivity of water, but it has one huge advantage over water—it is an electrical insulator. This means that electrical devices can operate while submerged in oil without shorting out.

While mineral oil does not have the heat capacity of water, it still holds over 1,000 times more heat than air. This means that the server rack discussed earlier that needed 1,200 CFM of air to keep from burning up could be kept cool with just about 1 CFM of oil. The energy required to pump 1 CFM of oil is dramatically less than the energy required to blow 1,200 CFM of air. In a perfectly designed data center, where the amount of air blown or oil pumped is matched exactly to the heat load, the energy required to blow air is five times that required to pump oil for the same amount of heat removed. In reality, the amount of air moved through a data center is far more than that required to satisfy the load. This is due to the fact that not all of the air blown into a data center passes through a computer before it returns to the CRAC unit. Since the air is not ducted directly to the computers' air intakes, it is free to find its own path back to the CRAC unit, which is frequently over, around, or otherwise not through a server rack. As we will soon see, it is much easier to direct the path of oil and to pump just the right amount of oil to satisfy a given computer heat load. Thus, the energy required to circulate oil can be more than 10 times less than the energy required to circulate air.

IMMERSION COOLING SYSTEM

Now that we have established that mineral oil would be a far more efficient fluid to use for removing heat from computers, let's look at how a system could be built to take advantage of this fact.

Imagine a rack of servers. Now imagine that the rack

is tipped over onto its face. Now convert the rack into a tub full of servers. Now fill the tub with mineral oil.

Figures 3 and 4 show the system that LPS acquired and is using in its Research Park facility. The system is comprised of a tank filled with mineral oil that holds the servers and a pump module that contains an oil-to-water heat exchanger and oil circulation pump. In this installation, the heat exchanger is tied to the facility's chilled water loop; however, this is not a necessity, as will be discussed later. The oil is circulated between the tank and the heat exchanger by a small pump. The pump speed is modulated to maintain a constant temperature in the tank. This matches the cooling fluid supply directly to the load. The design of the tank interior is such that the cool oil coming from the heat exchanger is directed so that most of it must pass through the servers before returning to the heat exchanger. The combination of pump speed modulation and oil ducting means that the cooling fluid is used very efficiently. The system only pumps the amount needed to satisfy the load, and almost all of what is pumped passes through the load.

There are three interesting side benefits to immersion cooling in addition to its efficiency. The first is due to the fact that the system is designed to maintain a constant temperature inside the tank. Because the pump is modulated to maintain a set point temperature regardless of changes in server workload, the servers live in an isothermal environment. One of the causes of circuit board failures is due to the mismatch in the coefficients of thermal expansion, or CTEs. The CTEs for the silicon, metal, solder, plastic, and fiberglass used in a circuit board are all different, which means that these materials expand and contract at different rates in response to temperature changes. In an environment where the temperature is changing frequently due to load changes, this difference in CTEs can eventually lead to mechanical failures on the circuit board. Oil immersion reduces this problem by creating a temperature-stable environment.

The second side benefit is server cleanliness. Aircooled servers are essentially data center air cleaners. While data centers are relatively clean environments, there is still some dust and dirt present. Remember, a typical server rack is drawing in a large office space full of air every minute. Any dust or dirt in that air tends to accumulate in the chassis of the servers. Pumping oil versus blowing air tends to accumulate in the chassis of the servers.



Fig. 3. The immersion cooling system at the Laboratory for Physical Sciences, like the one pictured above, uses mineral oil to cool IT equipment. (Photo used with permission from Green Revolution Cooling: grcooling.com.)

The final side benefit of immersion cooling is silence. Immersion cooling systems make virtually no noise. This is not an insignificant benefit, as many modern air-cooled data centers operate near or above the Occupational Safety and Health Administration's allowable limits for hearing protection.

In addition to efficient use of cooling fluid and the side benefits mentioned above, there is another advantage to immersion cooling—server density. As mentioned earlier, a typical air-cooled server rack consumes about 10 kW. In some carefully engineered HPC racks, 15-20 kW of load can be cooled with air. In comparison, the standard off-the-shelf immersion cooling system shown in figure 3 is rated to hold 30 kW of server load with no special engineering or operating considerations.

DOING MORE WITH LESS

Let's take a look at how immersion cooling can enable more computation using less energy and infrastructure. Air cooling infrastructure.



Fig. 4. Network servers are submerged into a tank of mineral oil and hooked up to a pump that circulates the oil. (Photo used with permission from Green Revolution Cooling: grcooling.com.)

Air cooling infrastructure

Cooling air is typically supplied in a computer room with CRAC units. CRAC units sit on the computer room raised floor and blow cold air into the under-floor plenum. This cold air then enters the computer room through perforated floor tiles that are placed in front of racks of computers. Warm exhaust air from the computers then travels back to the top of the CRAC units where it is drawn in, cooled, and blown back under the floor. In order to cool the air, CRAC units typically use a chilled-water coil, which means that the computer room needs a source of chilled water. The chilled water (usually 45-55°F) is supplied by the data center chiller plant. Finally, the computer room heat is exhausted to the atmosphere outside usually via evaporative cooling towers.

Oil-immersion systems also need to expel heat, and one way is through the use of an oil-to-water heat exchanger; this means that oil-immersion systems, like CRAC units, need a source of cooling water. The big difference however is that CRAC units need 45-55°F water; whereas, oil-immersion systems can operate with cooling water as warm as 85°F. Cooling towers alone, even in August in the mid-Atlantic area, can supply 85°F water without using power-hungry chillers. Because oil-immersion systems can function with warm cooling water, they can take advantage of various passive heat sinks, including radiators,

geothermal wells, or nearby bodies of water.

The takeaway here is that there is a significant amount of expensive, energy-hungry infrastructure required to make and distribute cold air to keep computers in a data center cool. Much of this infrastructure is not required for immersion cooling.

Fan power

One of the primary benefits of immersion cooling is the removal of cooling fans from the data center. Not only are the energy savings that result from the removal of cooling fans significant, they are compounded by potentially removing the necessity for CRAC units and chillers.

Cooling fans in a typical 1U rack-mounted server consume roughly 10% of the power used by the server. Servers that are cooled in an oil-immersion system do not require cooling fans. This fact alone means that immersion cooling requires approximately 10% less energy than air cooling. Internal server fans, however, are not the only fans required for air-cooled computers. CRAC unit fans are also necessary in order to distribute cold air throughout the data center and present it to the inlet side of the server racks.

This CRAC unit fan power must be considered when determining the actual fan-power savings that can be realized by immersion cooling systems. Table 2 compares the power required to move 1 W of exhaust heat into a data center's chilled water loop for fanblown air cooling versus pump-driven oil-immersion cooling. The third column shows this power as a percentage of IT technical load. It shows that the power required to run all fans in an air-cooled system is equal to 13% of the technical load that is being cooled. This is contrasted with the power required to run pumps in an oil-immersion cooling system, which is equal to 2.5% of the technical load that is being cooled. The difference,

10.5%, represents the net fan-power savings achieved by switching from an air-cooled to immersion-cooled data center. This analysis assumes that in both the air-cooled and immersion-cooled cases, the cooling infrastructure is matched exactly to the load. The last column in table 2 uses a similar analysis but assumes that the cooling infrastructure capacity is provisioned at twice the load. It shows that overprovisioned fan power grows faster than overprovisioned pump power. This is further illustrated in figure 5.

Lower operating expenses

Table 3 compares the fan power versus pump power required to serve a 1 MW technical load, assuming the cooling infrastructure is sized to serve 150% of the load. It shows that the fan power to circulate cold air exceeds the pump power to circulate oil by 158 kW per megawatt of technical load. At one million dollars per megawatt-year, this equates to \$158,000 a year in additional cooling energy operating expense. This represents the savings due solely to circulating cooling fluid. When the cost of making cold air is considered, the energy savings of immersion cooling becomes much more significant.

Table 4 summarizes the energy required for aircooling that is not needed for immersion cooling. The values in Table 4 are typical for reasonably efficient data centers. One ton of refrigeration will cool approximately 3,500 W of technical load; therefore, 1 MW of technical load requires a minimum of 285 tons of refrigeration. At 2.1 kW/ton, the air-cooled data center cooling infrastructure consumes about 600 kW to cool 1 MW worth of technical load. This equates to \$600,000 per year per megawatt of technical load. Almost all of this energy cost can be eliminated by immersion cooling since chillers, CRAC units, and server fans are not required.

Method of Cooling	Power Required to Move 1W of Waste Heat into Chilled Water Loop (W)	Percentage of Technical Load to Power Fan or Pump (at 100%)	Percentage of Technical Load to Power Fan or Pump (at 200%)
Fan-Powered Air	0.13 W	13%	26%
Pump-Powered Oil Immersion	0.025 W	2.5%	5%
<i>Net savings due to fan removal</i>		10.5%	21%

Table 2. Power usage for air-cooled versus immersion-cooled data centers




Figure 5. The power required to run the fans in an air-cooled data center (purple line) accounts for about 13% of the center’s technical load (26% if run at twice the technical load); whereas, the power required to run the pumps in an immersion-cooled data center (green line) accounts for about 2.5% of the center’s technical load (5% if run at twice the technical load). As is illustrated, overprovisioned fan power grows faster than overprovisioned pump power.

Lower capital expenses

Immersion cooling requires far less infrastructure than air cooling; therefore, building data centers dedicated to immersion cooling is substantially less expensive.

Cooling infrastructure accounts for a major portion of data center construction costs. In high reliability/availability data centers, it is not uncommon for the cooling infrastructure to account for half of the overall construction cost. According to the American Power Conversion Data Center Capital Cost Calculator, cooling infrastructure accounts for **at least 43%** of data center construction cost.

For large data centers, where the technical load is in the neighborhood of 60 MW, construction costs can approach one billion dollars. This means that about 500 million dollars is being spent on cooling infrastructure per data center. Since immersion-cooled systems do not require chillers, CRAC units, raised flooring, and

Cooling Equipment	Power Usage (kW/ton)
Chillers	0.7 kW/ton
CRAC Units	1.1 kW/ton
Server Fans	0.2 kW/ton
<i>Total</i>	<i>2.1 kW/ton</i>

Table 4. Power usage of cooling equipment in air-cooled data centers

temperature and humidity controls, etc., they offer a substantial reduction in capital expenditures over air-cooled systems.

IMMERSION COOLING FAQs

Several recurring questions have emerged over the many tours and demonstrations of the LPS immersion cooling system. Here are answers to these frequently asked questions.

Q. What server modifications are required for immersion?

There are three modifications that are typically required including:

1. Removing the cooling fans. Since some power supplies will shut down upon loss of cooling, a small emulator is installed to trick the power supply into thinking the fan is still there.
2. Sealing the hard drives. This step is not required for solid-state drives or for newer sealed helium-filled drives.
3. Replacing the thermal interface paste between chips and heat spreaders with indium foil. Some server vendors are already looking at providing immersion-ready servers which will be shipped with these modifications already made.

Method of Cooling	Fan or Pump Power as a Percentage of Technical Load (at 150% capacity)	Total Power (at 150% capacity)
Fan-Powered Air	19.5%	1.195 MW
Pump-Powered Oil Immersion	3.75%	1.0375 MW
<i>Delta</i>		<i>158 kW</i>

Table 3. Power usage for air-cooled versus immersion-cooled data centers with 1 MW of technical load

Q. Are there hazards associated with the oil? (e.g., fire, health, spillage)

With regard to flammability, the mineral oil is a Class IIIB liquid with a flammability rating of 1 on a scale of 4. Accordingly, immersion cooling does not require any supplemental fire suppression systems beyond what is normally used in a data center. The health effects are negligible. The oil is essentially the same as baby oil.

Spills and leaks are considered a low probability; however, for large installations, some form of spill containment is recommended. Spill decks, berms, curbs, or some other form of perimeter containment is sufficient.

Q: How much does the system weigh?

A 42U tank fully loaded with servers and oil weighs about 3,300 pounds, of which the oil accounts for about 1,700 pounds. This weight is spread over a footprint of approximately 13 square feet for a floor loading of approximately 250 pounds per square foot. A comparably loaded air-cooled server rack weighs about 1,600 pounds with a footprint of 6 square feet, which also translates to a floor loading of about 250 pounds per square foot.

Q: How is the equipment serviced or repaired?

Basic services such as device and board-level replacements are not significantly different than for air-cooled equipment. Hot-swaps can be done in the oil. For services requiring internal access, the server can be lifted out of the tank and placed on drainage rails above the surface of the oil. After the oil drains, component replacement is carried out the same way as for air-cooled servers.

For rework at the circuit board level that requires removal of the oil, there are simple methods available to ultrasonically remove oil from circuit boards and components.

Q: Are there other types of immersion-cooling systems besides oil immersion?

Yes. What this article has covered is called singlephase immersion. That is, the oil remains in the liquid phase throughout the cooling cycle. There are some people looking into two-phase immersion-cooling systems. In a two-phase cooling process, the cooling liquid is boiled off. The resulting vapor is captured and condensed before being recirculated. The phase change from

liquid to gas allows for higher heat removal but adds to the complexity of the system. Also, the liquid used in two-phase systems is extremely expensive compared to mineral oil. At this time, there are no two-phase immersion-cooling systems commercially available.

CONCLUSION

Computers consume energy and produce computation and heat. In many data centers, the energy required to remove the heat produced by the computers can be nearly the same as the energy consumed performing useful computation. Energy efficiency in the data center can therefore be improved either by making computation more energy efficient or by making heat removal more efficient.

Immersion cooling is one way to dramatically improve the energy efficiency of the heat removal process. The operating energy required for immersion cooling can be over 15% less than that of air cooling. Immersion cooling can eliminate the need for infrastructure that can account for half of the construction cost of a data center. In addition, immersion cooling can reduce server failures and is cleaner and quieter than air cooling.

Immersion cooling can enable more computation using less energy and infrastructure, and in these times of fiscal uncertainty, the path to success is all about finding ways to do more with less.

ABOUT THE AUTHOR

David Prucnal has been active as a Professional Engineer in the field of power engineering for over 25 years. Prior to joining NSA, he was involved with designing, building, and optimizing high-reliability data centers. He joined the Agency as a power systems engineer 10 years ago and was one of the first to recognize the power, space, and cooling problem in high-performance computing (HPC). He moved from facilities engineering to research to pursue solutions to the HPC power problem from the demand side versus the infrastructure supply side. Prucnal leads the energy efficiency thrust within the Agency's Advanced Computing Research team at the Laboratory for Physical Sciences. His current work includes power-aware data center operation and immersion cooling. He also oversees projects investigating single/few electron transistors, three-dimensional chip packaging, lowpower electrical and optical interconnects, and power efficiency through enhanced data locality.