



Immersion Cooling Performance for High-Density Data Centre Server Racks

Hala Anderson^{1*}, Aibek Radovic², Nurlan Subramaniam³

¹Department of Engineering, Loire Center for Engineering, Armenia

²Department of Engineering, Periyar School of Engineering, Romania

³Department of Engineering, Sharavati Center for Engineering, Bulgaria

Article Info

Volume 1, Issue 1, 2026

[doi: 10.62715/ijtet.2026.v1.i1.pp25-28](https://doi.org/10.62715/ijtet.2026.v1.i1.pp25-28)

Received: 29 March 2026

Accepted: 15 May 2026

Published: 15 July 2026

Abstract

Rack power densities in data centres have risen to levels at which air, even with careful containment, struggles to carry the heat away. Immersion cooling, in which servers are submerged in a dielectric liquid, is one of the responses now offered commercially. It comes in two forms: single-phase systems that circulate an oil or synthetic fluid, and two-phase systems in which a low-boiling fluid evaporates on the components and condenses on a coil above the bath. Claims made for both tend to rest on a single figure of merit, usually power usage effectiveness. This communication argues that the performance of an immersion-cooled rack is better understood as a chain of thermal resistances from chip to ambient, and that the weakest link differs between the two forms. It summarises what the published literature establishes about each, notes why power usage effectiveness flatters immersion systems when compared with air-cooled ones, and identifies questions about component-level heat flux limits, fluid stewardship and serviceability that remain open. The note is a commentary on published work and presents no new measurements. Its purpose is to help designers and operators ask more precise questions of the technology.

Keywords: Immersion cooling, Data centres, Dielectric fluids, Two-phase cooling, Power usage effectiveness, Thermal management

© 2026 Hala Anderson et al. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

1. Introduction

Data centres account for a visible share of world electricity use, and although efficiency gains have so far kept their consumption from growing in step with the computing they deliver (Masanet et al., 2020), the heat released per rack has climbed steeply with the spread of accelerator-based servers. Air cooling has been stretched to meet the trend through aisle containment, rear-door heat exchangers and higher permitted inlet temperatures (ASHRAE, 2021). The limits of air are physical. Its volumetric heat capacity is low, so large flows and fan powers are needed, and the heat sinks required for the hottest devices no longer fit the space available (Garimella et al., 2017).

Liquid cooling moves the heat into a medium that can carry far more of it per unit volume. Cold plates

* Corresponding author: Hala Anderson, Department of Engineering, Loire Center for Engineering, Armenia. E-mail: hala.anderson@loire.edu

do so for selected components. Immersion goes further and places the whole server in an electrically insulating liquid. Reviews of data centre cooling list it among the emerging options (Ebrahimi *et al.*, 2014; Capozzoli and Primiceri, 2015), and designs for dense racks have been described in some detail (Kheirabadi and Groulx, 2016). This communication comments on how the performance of such systems should be framed.

2. Two Forms of Immersion

In single-phase immersion the liquid, typically a mineral oil or a synthetic hydrocarbon, stays liquid throughout. It is pumped or allowed to circulate naturally past the boards and through a heat exchanger that passes the heat to facility water. Eiland *et al.* (2014) examined a server in mineral oil and showed how component temperatures respond to oil flow rate and inlet temperature. The fluids are inexpensive and have negligible vapour pressure. They are also viscous, with modest thermal conductivity, so the convective coefficient at a component surface is limited and the hottest devices may still need extended surfaces designed for liquid.


In two-phase immersion the bath is a fluorinated liquid chosen to boil at a temperature somewhat below the allowable component temperature. Vapour rises, condenses on a water-cooled coil in the headspace and falls back, with no pump on the dielectric side. Tuma (2010) set out the case for this open-bath arrangement. Boiling gives heat transfer coefficients well above those of single-phase convection in oil, and the bath temperature is pinned at saturation. The limit is the critical heat flux, which for dielectric liquids is low, so that high-flux devices require boiling-enhancement coatings or spreaders (Bar-Cohen *et al.*, 2006; Mudawar, 2001). Table 1 places both forms alongside air and cold-plate cooling.

Table 1: Rack cooling approaches compared by heat path, principal strength and principal limitation			
Approach	Heat path from component	Principal strength	Principal limitation
Air with containment	Heat sink to air, then to room cooling units	Familiar, serviceable, no liquid near electronics	Fan power and heat sink size at high density
Cold plates	Plate on selected chips to water or glycol loop	Very low resistance at the hottest devices	Remaining components still need air; many fittings
Single-phase immersion	All surfaces to circulating oil, then to water	Captures all heat; simple, low-cost fluid	Viscous fluid; limited coefficient at hot spots
Two-phase immersion	Boiling at surfaces; vapour condenses on coil	High coefficient; no pump on fluid side	Low critical heat flux; fluid cost and vapour loss
Hybrid immersion	Cold plates on processors inside an immersion tank	Combines targeted and bulk heat removal	Complexity; two liquid circuits to maintain

3. Performance as a Chain of Resistances

A rack is cooled adequately when every component stays below its temperature limit at the worst combination of load and ambient conditions. Heat generated in a chip must cross the series of stages shown in Figure 1, and the temperature difference between junction and ambient is shared among them. The useful question about any cooling method is which stage takes the largest share and how that changes with power.

For single-phase immersion the dominant stage at high power is usually the transfer from the package to the liquid. The remedy is flow management inside the tank: ducting the liquid through the hottest regions instead of letting it bypass them, and fitting heat sinks with fin spacing suited to a viscous medium. For two-phase immersion the surface coefficient is high until the heat flux at the lid approaches its critical value, at which point the device temperature runs away. The governing stage is then the spreading of heat within the package, and at rack level the capacity of the condenser, since vapour that is not condensed raises tank pressure or escapes.



The stages beyond the tank matter as much. Immersion systems can accept warmer facility water than chilled-air systems, which permits heat rejection by dry coolers for much of the year and makes the rejected heat more useful for reuse (Ebrahimi et al., 2014). That benefit is realised only if the temperature differences inside the tank are small enough to leave room for it.

4. What the Usual Metric Hides

Power usage effectiveness, defined in the international standard as the ratio of total facility energy to the energy delivered to the information technology equipment (ISO/IEC, 2016), is the figure most often quoted for immersion installations. It needs careful reading. In an air-cooled server the internal fans are powered by the server and are counted as information technology load. Remove them, as immersion does, and the denominator shrinks along with the numerator. Two facilities with the same ratio may therefore differ in total energy, and a move to immersion can save energy that the ratio does not show. Comparisons between cooling methods are sounder when based on total energy for a defined computing workload, and when they state the component temperatures maintained, since a method that runs chips hotter has not delivered the same service.

5. Open Questions

Three matters are not settled by the published literature. The first is the heat flux that each form can sustain at the component, as opposed to the rack. Rack-level power is an average, and processors continue to concentrate power in small areas; published boiling data for dielectric liquids on enhanced surfaces come largely from small test heaters and not from production packages. The second is the fluid. Fluorinated liquids are costly, some have high global warming potential, and the regulatory outlook for this family of substances is uncertain. Hydrocarbon oils avoid those concerns but are combustible and can affect elastomers, cable jackets and labels over time.

The third is operation. Servers must be lifted out to be serviced, fluid is carried out with them, and standard components are qualified for air. Long-term reliability data for immersed hardware in the open literature are sparse compared with the record for air cooling, and operators weighing the technology would be helped by published field experience that reports failures as well as energy figures.

6. Conclusion

Immersion removes the air-side bottleneck that constrains dense racks, and both of its forms are technically sound. Their performance is not captured by a single efficiency ratio. It is better described by

where the temperature difference from chip to ambient is spent: at the component surface in single-phase systems, and at the approach to critical heat flux and in the condenser in two-phase ones. Framing the comparison in those terms, and reporting total energy and component temperatures together, would make published claims easier to test.

Conflict of interest

To be declared by the author(s) before publication.

Funding

To be declared by the author(s) before publication.

References

- ASHRAE (2021). [Thermal Guidelines for Data Processing Environments](#), 5th edition. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta.
- Bar-Cohen, A., Arik, M. and Ohadi, M. (2006). [Direct liquid cooling of high flux micro and nano electronic components](#). *Proceedings of the IEEE*, 94(8): 1549-1570.
- Capozzoli, A. and Primiceri, G. (2015). [Cooling systems in data centers: State of art and emerging technologies](#). *Energy Procedia*, 83: 484-493.
- Ebrahimi, K., Jones, G.F. and Fleischer, A.S. (2014). [A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities](#). *Renewable and Sustainable Energy Reviews*, 31: 622-638.
- Eiland, R., Fernandes, J., Vallejo, M., Agonafer, D. and Mulay, V. (2014). [Flow rate and inlet temperature considerations for direct immersion of a single server in mineral oil](#). *Proceedings of the Fourteenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 706-714.
- Garimella, S.V., Persoons, T., Weibel, J.A. and Gektin, V. (2017). [Electronics thermal management in information and communications technologies: Challenges and future directions](#). *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 7(8): 1191-1205.
- ISO/IEC (2016). [ISO/IEC 30134-2:2016, Information technology: Data centres, Key performance indicators, Part 2: Power usage effectiveness \(PUE\)](#). International Organization for Standardization, Geneva.
- Kheirabadi, A.C. and Groulx, D. (2016). [Cooling of server electronics: A design review of existing technology](#). *Applied Thermal Engineering*, 105: 622-638.
- Masanet, E., Shehabi, A., Lei, N., Smith, S. and Koomey, J. (2020). [Recalibrating global data center energy-use estimates](#). *Science*, 367(6481): 984-986. <https://doi.org/10.1126/science.aba3758>
- Mudawar, I. (2001). [Assessment of high-heat-flux thermal management schemes](#). *IEEE Transactions on Components and Packaging Technologies*, 24(2): 122-141.
- Tuma, P.E. (2010). [The merits of open bath immersion cooling of datacom equipment](#). *Proceedings of the 26th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM)*, 123-131.

Cite this article as: Anderson H., Radovic A. and Subramaniam N. (2026). [Immersion Cooling Performance for High-Density Data Centre Server Racks](#). *International Journal of Thermal Engineering and Technology*. 1(1), 25-28. doi: 10.62715/ijtet.2026.v1.i1.pp25-28