



PRODUCT MARKET ASSESSMENT & TECHNOLOGY LANDSCAPE

Immersive Cooling

Liquid cooling for high-density compute — market timing, economics, technical requirements and ecosystem.

Contents

01 Research Approach

How the assessment was built

02 Why Immersive Cooling Now?

Industry drivers and timing

03 Air Cooling vs. Immersion Cooling

Cost structure at 10kW/rack

04 ROI Analysis

5-year TCO for a 1MW deployment

05 Cooling Technology Comparison

Four technologies across 24 measures

06 Key Market Opportunities

The data center cooling crisis

07 Technical Requirements

Compliance, monitoring and integration

08 Market Size and Growth

Segments, regions and adoption timeline

09 Supply Chain: Ecosystem

Who builds what, and where cooling sits

10 Adjacent Opportunities

Where the same thermal problem recurs

Research Approach

Ezassi employed a qualitative approach, directly interviewing individuals with experience in data centers, specifically immersive cooling both from a technical and market opportunity perspective

Ezassi employed quantitative research to validate the key findings identified through the interviewing process. Utilizing:

- **Data Sources explored:** News feeds, web-monitoring, patents, publications, supplier/vendor databases, startup databases, grants, & trademarks
- **Geographic regions covered:** Global
- Benchmarks and other comparative mechanisms, where necessary



Technology to innovate faster.

Why Immersive Cooling Now?


Optimal timing for market entry

Industry drivers

- Traditional air cooling is reaching useful physical limits
- AI/ML workloads demanding higher energy densities
- ESG pressures for improved efficiency (sustainability)

Potential key benefits

- ~50% reduction in power consumption
- ~60% reduction in physical space
- Superior thermal performance
- Enhanced hardware lifespan



With CPUs reaching 200 to 400W and GPUs moving towards 1000 to 1200W, air cooling is becoming increasingly inadequate.

Air Cooling vs. Immersion Cooling Analysis

Baseline comparison (at 10kW/rack)

- Equal capital costs for both solutions
- 2MW data center analysis
- Traditional air-cooled vs. chassis-based immersion


Cost structure changes

Major savings

- Chiller/CRAH elimination (−\$0.91/W)
- Reduced UPS & switchgear (−\$0.14/W)
- Space & containment savings (−\$0.10/W)

Added costs

- Liquid cooling technology (+\$0.77/W)
- Dry coolers & CRACs (+\$0.31/W)
- Pumps & piping (+\$0.03/W)



ROI

In high-density data centers, liquid cooling improves the energy efficiency of both IT and facility systems compared with air cooling.

At the Leibniz Supercomputing Centre's SuperMUC-NG, warm-water cooling at 40 °C is reported to deliver around 30% energy savings — fanless servers, lower cooling power, energy-aware scheduling and less refrigeration — at a partial cooling PUE below 0.07.

5-YEAR TCO FOR A 1MW DEPLOYMENT

ROI Analysis

Total cost of ownership

- **Air cooling:** \$6.5–7.5M (\$1.1M CAPEX, \$5.4–6.4M OPEX)
- **Direct-to-chip:** \$5.0–6.0M (\$1.3M CAPEX, \$3.7–4.7M OPEX)
- **Immersion cooling:** \$4.5–5.5M (\$1.5M CAPEX, \$3.0–4.0M OPEX)

Payback period

- Direct-to-chip vs. air: 1.5–2.5 years
- Immersion vs. air: 2.0–3.0 years

Three-way technology comparison

Factor	Air Cooling	Direct-to-Chip	Immersion Cooling
Capital costs	\$0.90–1.10/W	\$1.10–1.30/W	\$1.20–1.50/W
Annual operating costs	\$0.08–0.10/W	\$0.06–0.08/W	\$0.04–0.06/W
Maintenance complexity	Low	Medium	High
Maintenance downtime	1–2 hours per server	2–4 hours per server	6–12 hours per system
Power reduction vs. air	Baseline	30–40%	40–50%
Water usage reduction	Baseline	60–70%	80–95%
Physical footprint	Baseline	20–30% reduction	40–60% reduction
Time to deploy	Fast (established)	Medium (emerging)	Slow (developing)

Cooling Technology Comparison Matrix (1 of 2)

Feature	Air Cooling	Direct-to-Chip Liquid	Single-Phase Immersion	Two-Phase Immersion
Technical Specifications				
Max power density	10–15 kW/rack	20–50 kW/rack	40–80 kW/rack	50–100+ kW/rack
Component temp	Variable ($\pm 10^{\circ}\text{C}$)	Stable ($\pm 5^{\circ}\text{C}$)	Very stable ($\pm 3^{\circ}\text{C}$)	Extremely stable ($\pm 1^{\circ}\text{C}$)
Cooling efficiency	Low	Medium	High	Very high
Max component power	~350W	~700W	~700W	1000W+
Implementation				
Installation complexity	Low	Medium	High	Very high
Implementation timeline	4–8 weeks	8–12 weeks	12–20 weeks	16–24 weeks
Facility requirements	Standard	Enhanced	Specialized	Highly specialized
Floor loading	Standard	Standard	Heavy (2000+ lbs/rack)	Very heavy (3000+ lbs/rack)
Financial Metrics				
CAPEX (\$/W)	\$0.90–1.10	\$1.10–1.30	\$1.20–1.40	\$1.40–1.60
OPEX (\$/W/year)	\$0.08–0.10	\$0.06–0.08	\$0.04–0.06	\$0.03–0.05
5-year TCO (per MW)	\$6.5–7.5M	\$5.0–6.0M	\$4.5–5.5M	\$4.0–5.0M
Payback period	Baseline	1.5–2.5 years	2.0–3.0 years	2.5–3.5 years

Cooling Technology Comparison Matrix (2 of 2)

Feature	Air Cooling	Direct-to-Chip Liquid	Single-Phase Immersion	Two-Phase Immersion
Operational Factors				
Maintenance complexity	Low	Medium	High	Very high
Maintenance downtime	1–2 hours/server	2–4 hours/server	6–12 hours/system	8–16 hours/system
Personnel requirements	Standard DC techs	Specialized training	Extensive training	Expert specialists
Component compatibility	Universal	Limited	More limited	Most limited
Future Readiness				
AI workload support	Limited	Good	Excellent	Superior
Density scaling	Poor	Good	Excellent	Superior
Heat reuse potential	Low (25–35%)	Medium (40–60%)	High (70–85%)	Very high (85–95%)
Energy efficiency	Low	Medium	High	Very high
Implementation Timeline				
Current adoption	Mainstream	Early mainstream	Early adopters	Innovators
Widespread adoption	Declining	2024–2025	2025–2027	2027–2029
Best use case	Standard compute	Mixed workloads	AI clusters, HPC	Extreme density, specialized

Key Market Opportunities


The data center cooling crisis

Current challenge

- AI/GPU power consumption increasing from 500W to 2000W per processor
- Traditional data centers struggling with 10–15kW per rack
- New AI servers demanding up to 14.3kW each
- Cooling costs consuming up to 40% of data center energy bills
- Standards within the liquid cooling industry are being defined through communities such as the Open Compute Project (OCP) and ASHRAE

Market leaders already moving

- Microsoft partnering with Dell/Submer
- Meta exploring immersion solutions
- Google investing in innovative cooling technologies
- Amazon's 1000 Watt next-gen Trainium AI chip lays the groundwork for liquid-cooled data centers to house new AI chips



Technical Requirements

Essential compliance & integration

Safety & certification

- Regional certifications (UL, CE)
- Minimum dielectric strength ≥ 6 kV/mm
- Flash point $\geq 150^{\circ}\text{C}$
- Auto ignition point $\geq 300^{\circ}\text{C}$
- Safety protocols

Monitoring & control

- Temperature monitoring (minimum 2 locations)
- Leak detection systems
- Pump status monitoring
- Redfish-compliant management interface

Facility integration

- Water system pressure compatibility (50–1000 kPa)
- Standard interface support (MODBUS, SNMP, BACNet)
- Material compatibility with cooling systems

Market Size and Growth

Global data center liquid cooling market

- \$2.8B (2023), projected to reach \$12.5B by 2030 (CAGR 24.5%)
- **Regional growth rates:** APAC (29.8%), North America (22.7%), Europe (21.4%)
- **Market segments:** Hyperscale (42%), Enterprise (28%), Colocation (19%), Edge (11%)

Key industry developments

- **NVIDIA's market influence**
 - Releasing immersion-specific GPUs optimized for liquid cooling environments in Q3 2025
 - Grace Hopper Superchip (GH200) supporting liquid cooling configurations
 - Next-gen accelerators (expected 2025–2026) requiring liquid cooling for full performance

Technology adoption timeline

- **Current early adopters (2024):** Microsoft, Google, Meta, select hyperscalers
- **Near-term (2025–2026):** Colocation providers, financial services, HPC
- **Mid-term (2027–2028):** Enterprise data centers, telecommunications
- **Mainstream (2029+):** Standard deployment for high-density computing


Market penetration projections

- **2025:** 15% of new high-density deployments (>30kW/rack)
- **2027:** 35% of new high-density deployments
- **2030:** 60% of new high-density deployments

Supply Chain: Ecosystem

Ecosystem is made up of all the players that make up a data center and the infrastructure.

This includes OEMs as well as semiconductor manufacturers — all of whom play a role in the movement to liquid cooling.



Adjacent Industry & Application Opportunities

Electric Vehicles & Charging

- Fast-charging stations
- High-power converters
- Power distribution units
- Battery thermal management
- Power electronics cooling

Diagnostic Equipment

- MRI cooling systems
- CT scanner electronics
- X-ray equipment

Marine & Military Applications

- Ship propulsion systems
- Navigation equipment
- Offshore power systems

Industrial Testing Equipment

- Battery testing facilities
- Power cycling equipment
- Thermal stress testing

Industrial Motors & Drives

- High-power motor controllers
- Variable frequency drives
- Industrial automation systems

Welding Equipment

- Industrial welding machines
- Plasma cutters
- High-power laser systems