

INDUSTRY BRIEF

PROTECTING HIGH-FREQUENCY
TRADING PERFORMANCE WITH

WATERLESS DIRECT-TO-CHIP COOLING

legrand

IR150

legrand

MR250

THE CHALLENGE: COOLING RISK IS BUSINESS RISK

For a high-frequency trading (HFT) firm, compute performance, latency consistency, and uptime all depend on cooling performance. Trading platforms and the modeling workloads behind them depend on processors operating within tight thermal limits, especially when market activity and compute demand spike. If temperatures rise, processors can throttle, and performance can become less predictable. If a conventional water-based cooling system leaks, the result may be damaged equipment, emergency remediation, and unplanned downtime.


HFT organizations already understand what every minute of availability is worth to their business. The question is how to increase cooling capacity without introducing additional risk of failure or adding maintenance burdens at sites that may not have dedicated liquid-cooling personnel.

The thermal demands of HFT infrastructure are also changing. Trade execution may remain concentrated on latency-optimized CPUs and FPGAs, but the research, simulation, and predictive-modeling environments supporting modern strategies increasingly incorporate power-intensive CPUs and GPUs. These workloads do not make every trading facility an ultra-high-density environment. They do, however, create concentrated thermal loads that can strain air cooling and reduce the thermal headroom available during periods of peak activity.

SOLUTION AT A GLANCE: WATERLESS COOLING WITHOUT ADDED COMPLEXITY

Two-phase direct-to-chip cooling provides a different approach. Using a dielectric refrigerant instead of water enables trading firms to:

1. **Increase cooling capacity**
2. **Eliminate the risk of water exposure inside the IT environment**
3. **Support next-generation CPUs and GPUs**
4. **Simplify operations and reduce maintenance**
5. **Deploy incrementally within existing facilities**



A DIFFERENT APPROACH TO DIRECT-TO-CHIP COOLING

Water-based direct-to-chip cooling can address rising thermal loads, but it introduces operational requirements of its own. Water quality must be maintained, connections monitored, and response procedures established for leaks near valuable electronics. These responsibilities can be especially burdensome for organizations operating across multiple private facilities or colocation sites without specialized liquid-cooling personnel on-site.

Two-phase direct-to-chip cooling changes this risk equation. Instead of circulating water through cold plates at the server, it uses a dielectric, nonconductive refrigerant that absorbs heat directly from the chip. As the refrigerant changes from liquid to vapor, it removes latent heat, not just sensible heat, from the processor before condensing and circulating through the closed loop again.

No water circulates through the servers or the IT cooling loop. If refrigerant escapes, it does not present the same electronic damage risk as water. The closed-loop design also avoids the water treatment and fluid-quality maintenance required by water-based systems. It reduces dependence on rack-level water-leak monitoring and mitigation, as well as the need for specialized on-site cooling support.

DEPLOY ADVANCED COOLING WITHOUT REBUILDING THE DATA CENTER

The Accelsius IR150 delivers waterless two-phase cooling in a pre-integrated rack-level platform, combining a 150 kW two-phase coolant distribution unit, a 42U IT rack, and integrated liquid and vapor manifolds in a single enclosure. It helps organizations deploy advanced cooling without extensive facility modifications. Redundant pumps and controls help support continuous operation.

Unlike facility-wide cooling designs, the IR150 can be deployed incrementally where cooling capacity is needed most. The integrated platform can be rolled into place and connected to facility power and heat-rejection infrastructure, enabling deployment in existing facilities without extensive white-space reconstruction. Firms can introduce two-phase cooling where it is needed, one deployment at a time, rather than committing an entire facility to a new cooling architecture. This incremental approach limits disruption, preserves flexibility, and keeps the cooling failure domain at the individual-rack level. It also allows infrastructure teams to align cooling investments with actual compute growth rather than making facility-wide changes in anticipation of future demand.



MAINTAIN PREDICTABLE PERFORMANCE DURING PEAK DEMAND

By removing heat directly at the processor, two-phase cooling helps maintain consistent thermal conditions during periods of increased market activity. This enables CPUs and GPUs to operate within optimal thermal ranges, helping sustain performance during high-demand periods, and reducing the risk that temperature fluctuations or thermal throttling will affect latency-sensitive workloads.

That consistency is valuable during market events, when trading activity, analytics, and supporting compute workloads can rise rapidly. Infrastructure teams gain a more predictable thermal environment without relying on colder air temperatures or overprovisioned mechanical cooling.

EXPAND COMPUTE CAPACITY WITHIN EXISTING FACILITIES

Many trading firms scale across multiple private data centers and colocation environments where space is constrained. A rack-integrated, retrofit-ready solution provides a practical way to add current and next-generation CPU and GPU capacity without expanding facilities or redesigning entire data centers.

For HFT environments, the goal is not extreme density for its own sake. It is gaining the compute capacity needed for evolving strategies while retaining control over availability, operational risk and infrastructure costs.

MODERNIZE COOLING WITHOUT ADDING WATER RISK

As thermal demands continue to increase across trading, analytics and AI environments, cooling infrastructure must evolve to support performance, availability and operational simplicity. Two-phase direct-to-chip cooling delivers the benefits of liquid cooling while addressing one of the primary concerns many HFT firms have about its adoption: water near high-value electronics. With no water in the IT cooling loop, lower fluid-maintenance demands and a pre-integrated deployment model, the Accelsius IR150 helps trading infrastructure teams protect availability, maintain consistent compute performance, and expand on their own terms. The result is a cooling strategy that helps trading firms increase compute capacity while maintaining control over operational risk.

Ready to learn more? Contact Legrand to discuss your cooling strategy today.

To learn more visit

legrand.us/cooling

Legrand.us

877.295.3472 | productspecialist.DAT@legrand.us

©2026, Legrand. All rights reserved. The industry-leading brands of Approved Networks, Kratos Industries, Ortronics, Raritan, Server Technology, Starline, and ZPE Systems empower Legrand's Power & Thermal Management division to produce innovative power, cooling, connectivity and management solutions for data center white space and gray space, building networks, and facility infrastructures. Our division designs, manufactures, and markets world-class products for a more efficient and sustainable future. The exceptional reliability of our technologies results from decades of proven performance and a dedication to research and development. Lit #V2446