

Navigating the liquid cooling landscape for AI data centers in 2026 and beyond

A practical guide for European businesses

DELTA

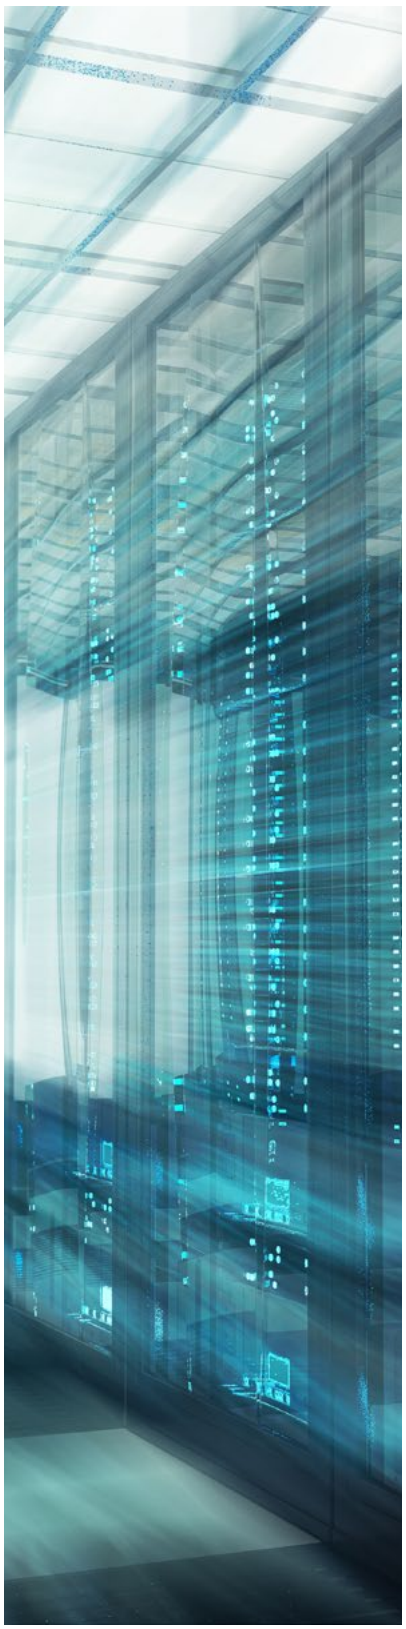

Introduction

The artificial intelligence revolution has already arrived for Europe's data centers, and it is rewriting the rules of infrastructure from the rack outward.

Three years ago, a standard data center rack drew three kilowatts of power. Today, AI-optimized racks exceed 650 kilowatts. By 2027, the industry is projecting one megawatt per rack, the thermal equivalent of managing a small power plant in a space the size of a wardrobe.

A single rack of this density houses 76 GPUs, generating heat that air cooling cannot physically handle.

The question facing European data center operators is how quickly they can adopt liquid cooling, which is better suited for these elevated cooling needs, and what's the best and most commercially viable way to roll it out.

Delta Electronics is that the heart of AI data centers' transition to liquid cooling. With 15,000 patents and 1.6 gigawatts of liquid cooling deployed globally in 2025 alone, Delta has established itself as the world's leading provider of advanced cooling infrastructure for AI data centers. That leadership is built on the principle that power systems, thermal management and intelligent monitoring must be engineered together, not assembled from separate vendors.

This white paper aims to guide European data center operators through the available cooling options and infrastructure paths to support intensive AI workloads.

The first is a time-to-market-led approach, retrofitting existing licensed facilities with proven liquid-to-air cooling technologies already deployed by hyperscalers such as Microsoft to generate revenue within 9 to 12 months.

The second is a future-proof approach: building new, purpose-designed facilities with liquid-to-liquid cooling infrastructure capable of supporting the next generation of AI workloads and the associated architectural shift to 800-volt DC power distribution, which promises the most significant step-change in data center energy efficiency in a decade.

The right choices depend on whether permits already exist, the scale of the intended facility, the operator's revenue timeline and their tolerance for technology obsolescence in a market where infrastructure decisions can be rendered outdated in less than 12 months. Delta works with operators across both paths, and in some cases across both at the same time.

For European operators, the urgency is increasing as EU data-center reporting requirements expand. Under the revised Energy Efficiency Directive and Delegated Regulation (EU) 2024/1364, data centers with an installed IT power demand of at least 500 kW must report annual energy-performance and sustainability indicators.

Standards such as ISO/IEC 30134 and EN 50600 provide the framework for defining and benchmarking key performance metrics. Liquid cooling is not required for compliance, but it can improve efficiency, support heat reuse and strengthen the operational case for high-density workloads.

This paper provides a practical guide to navigating that landscape: the technology options and how to choose between them; the design and implementation decisions that determine long-term performance; the operational discipline that liquid cooling demands; and the business case, including monetizing recovered heat through district heating and industrial partnerships, turning a cooling system into a revenue asset.

Delta's role in European data center infrastructure is as a strategic partner with the engineering depth, patent portfolio and deployment experience to help operators build infrastructure that performs today and adapts to future demands.



1. What is driving the need for liquid cooling?

For most of the past three decades, data center infrastructure evolved at a reasonably measured pace. Rack power densities increased gradually, cooling technology improved incrementally and operators could plan 5 or 10 years ahead with reasonable confidence. The rapid growth of AI has fundamentally altered that picture.



The graphics processing unit (GPU) is the engine of this transformation.

Existing NVIDIA GB200 and GB300 systems house 72 individual GPU chips on a 130kW rack, operating in close proximity to one another. This chip density is essential because GPUs must communicate with each other at near-zero latency to function collectively as a single AI processing unit.

Air cooling, the standard approach for traditional IT infrastructure, cannot adequately handle the intense thermal conditions this tight packing creates. Water is a much more effective thermal-transfer medium than air. At these high densities, that difference matters considerably, especially given that both AI demand and GPU technology are advancing at breakneck speed.

The next generation of NVIDIA GPUs (Vera Rubin NVL 144 and NVL 576) will squeeze 144 and 576 GPUs onto a rack, operating respectively from 144kW up to 190kW and 600kW per rack. With current GB200 and GB300 systems already demanding clustered configurations where a standard pod of eight cabinets consumes more than one megawatt (MW), we are likely to see substantially higher workloads when NVL144 and NVK 576-based systems launch.

This speed is unprecedented in data center design, and its impact goes beyond single racks: it affects how entire systems are built and organized. A conventional two-megawatt data center built for traditional IT workloads and an AI-ready facility of the same power rating are fundamentally different buildings: in layout, cooling philosophy, power distribution and in the proportion of space dedicated to compute versus infrastructure. Scaled to 50 megawatts, those differences become even more pronounced.

Power availability is now a constraint as significant as technology itself. Securing permits for a large AI facility, typically in the 50 to 100 megawatt range, can take up to five years in Europe. Infrastructure must therefore be designed not just for today's GPU generation but for the one after next, even as those specifications remain in flux. This creates a genuine tension for every operator: move fast to be competitive, or build carefully to remain viable over the long term.

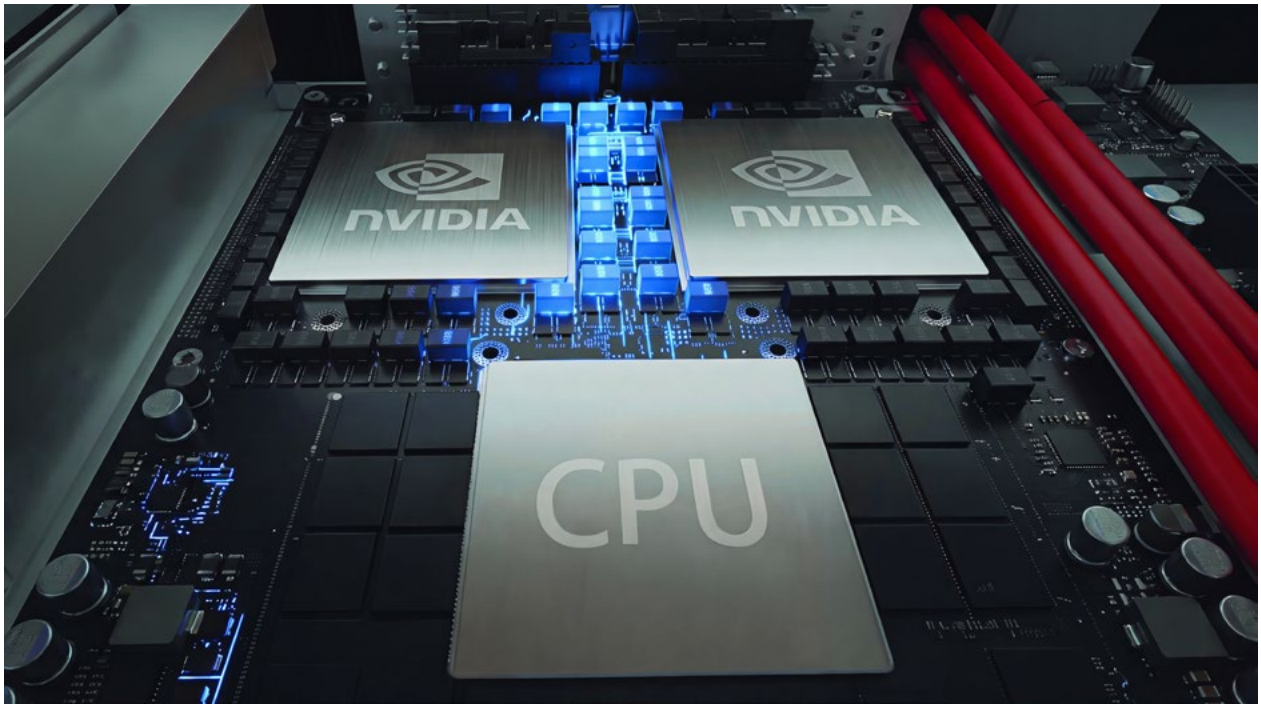

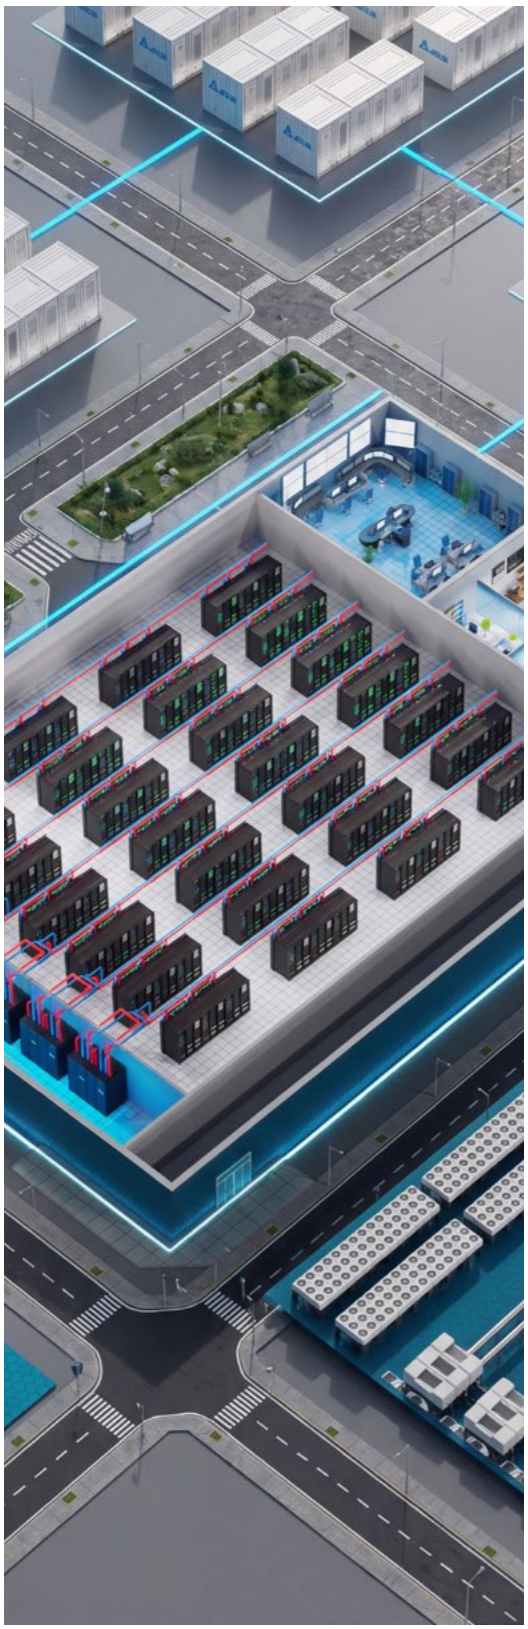

There is no single correct answer, which is why infrastructure strategy has become as important as technology selection.

The shift to AI has also redefined what “efficiency” means in a data center context. In an AI facility, the primary measure of economic performance is cost per token, the unit of output that large language models and AI systems produce. Reducing that cost requires maximizing compute density, minimizing energy loss at every point in the power chain and removing the thermal constraints that would otherwise limit GPU performance. This, alongside regulatory drivers, is accelerating the adoption of liquid cooling as the standard approach for AI workloads.

The question European data center operators face is no longer whether to transition to liquid cooling, but how quickly, at what scale, and in which configuration. At the same time, deciding on the most suitable liquid cooling technology to support this rapid expansion will be vital.

2. The AI/HPC thermal inflection point in Europe

2.1 Workload and density trends

The data center industry has lived through several waves of technological change, but none has compressed so much disruption into such a short timeframe as the current AI buildout.

What began as a gradual shift toward high-performance computing (HPC) has accelerated sharply, driven by large language models, GPU-based inference and the proliferation of AI workloads across both hyperscale and enterprise environments.

Rack densities that once averaged 10 to 20 kilowatts are now routinely reaching 50 to 100 kW in AI deployments, with the most demanding GPU clusters reaching 130kW today. But 600kW and 1MW are not far off with the rapid technological advances processors are undergoing. These are not theoretical projections: they reflect hardware already in production. The thermal consequences for existing air-cooled facilities are significant: hot spots, stranded capacity and cooling systems operating beyond their design envelopes. Many data halls built to current standards are already approaching the limits of what air cooling can reliably manage.

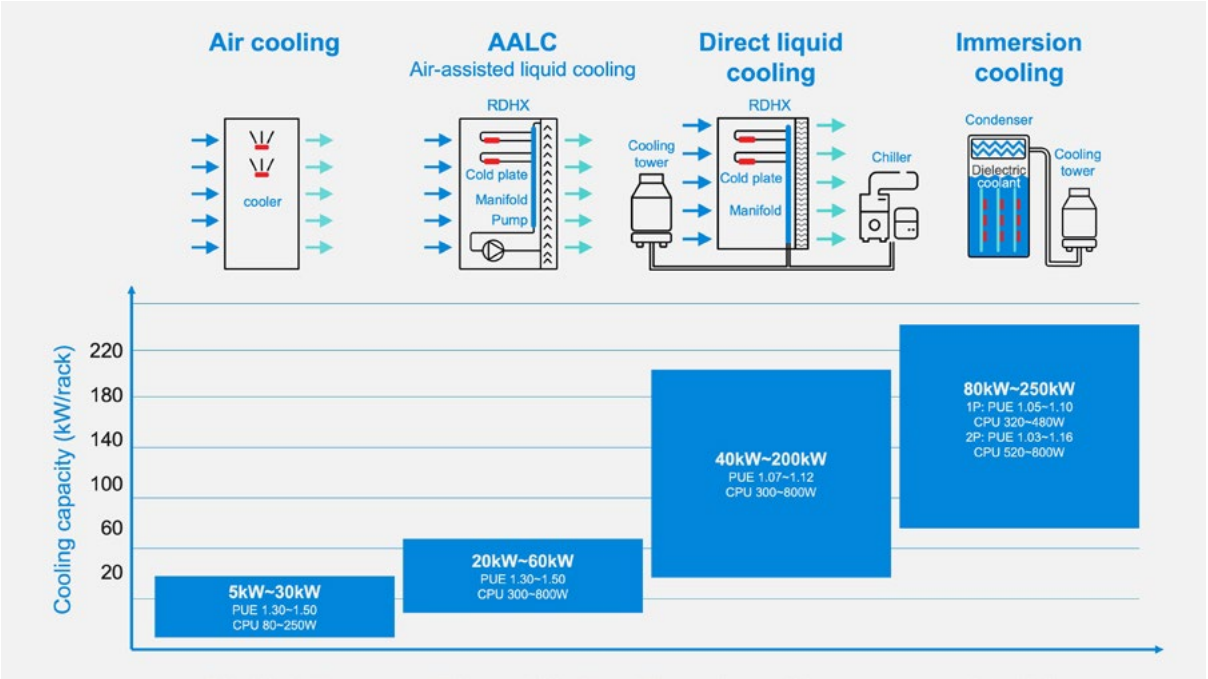

2.2 The emerging regulatory and ESG landscape

In Europe, this technical challenge arrives alongside an accelerating regulatory agenda. The EU Energy Efficiency Directive (EED) places mandatory obligations on data center operators with a total rated power of one megawatt or more, requiring the monitoring and public reporting of key performance indicators (KPIs), including power usage effectiveness (PUE), water usage effectiveness (WUE) and energy reuse factor (ERF).

The EN 50600 standard and ISO/IEC 30134 metrics provide the measurement framework, and operators that can demonstrate leadership on these KPIs are better positioned for both regulatory compliance and environmental, social and governance (ESG) credibility.

Beyond reporting, national and local governments across Europe are increasingly looking to data centers as potential contributors to district heating networks and urban energy systems. Liquid cooling, which recovers heat at temperatures useful for downstream reuse, is a practical enabler of these circular energy models in a way that air cooling is not.

2.3 Why liquid and hybrid cooling are becoming the default path

Air cooling remains viable for lower-density, mixed enterprise workloads. For AI, however, it is increasingly unfavorable. The thermal resistance of air is high relative to water, and the energy required to move sufficient air volume through a dense GPU rack becomes disproportionate at scale. Liquid cooling addresses this directly, transferring heat more efficiently and enabling higher compute density within the same footprint.

The market is responding accordingly. Liquid and hybrid cooling are no longer niche or emerging, they are becoming the standard approach for new AI deployments and a growing proportion of retrofits. For European operators, the question is how to sequence and scale the transition to liquid cooling, given their specific site constraints, workload profile, and regulatory obligations.

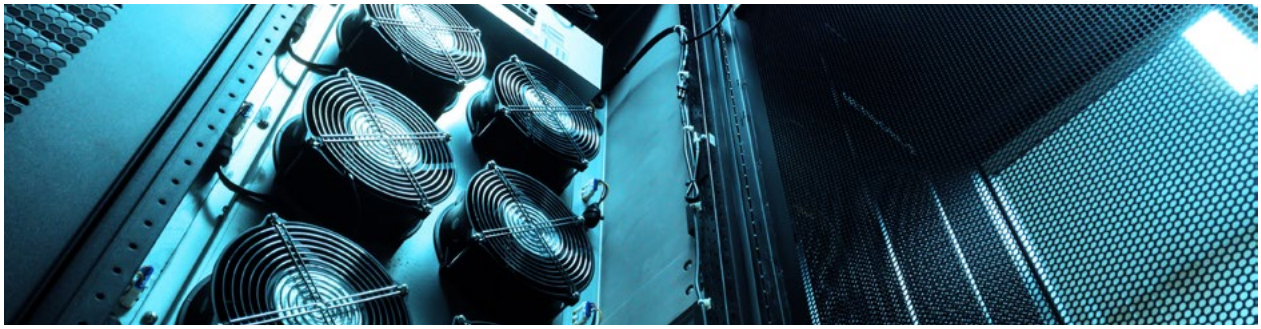

3. Assessing readiness

3.1 Time-to-market vs future-proof: how to choose your path

Once you've decided to add AI capability to your data center, the next question is usually about time. How quickly does the facility need to be operational, and how long does it need to remain viable?

The answer to those two questions, more than budget or geography, is likely to determine the right approach.

Two broad paths are available. The first prioritizes speed: retrofitting an existing facility that already has power permits in place, deploying liquid-to-air cooling to bring AI racks online within 9 to 12 months, generating a return on investment quickly.

The facility may not run on the latest hardware generation and may have a shorter operational lifespan than a purpose-built alternative, but it can be profitable well before a new-build would break ground. In some cases, operators choose to sell or repurpose a retrofitted site once a cleaner, larger facility is ready, treating the retrofit as a bridging asset rather than a long-term one.

The second path prioritizes longevity: building a new facility from the ground up means it can be designed around liquid-to-liquid cooling, 800V DC power architecture, and the density requirements of future GPU generations. This approach takes longer, permitting alone can run to five years in parts of Europe, but it produces a facility with a stronger total cost of ownership (TCO) position over a 10-year horizon and avoids the constraints that come with adapting existing infrastructure to a workload it was not originally designed to support.

Regional patterns offer useful context: in Europe, retrofitting is the more common choice, particularly among operators looking to move quickly and avoid the complexities of new-build planning. In the Middle East, where large greenfield projects are more prevalent, purpose-built AI facilities are the norm. Hyperscalers often pursue both simultaneously, building at scale for the long term while renting or retrofitting existing capacity to meet near-term demand.

What makes the decision difficult is that the two approaches are not easily compatible.

The project timelines, design standards and cooling architectures diverge early, and a facility started under one model cannot simply be converted to the other midway through. Choosing between them requires a clear view of the business plan: the expected workload trajectory, the revenue timeline and the operator's tolerance for infrastructure obsolescence.

Delta works with customers at this decision point: assessing the site, the business objectives and the regulatory environment to recommend an approach that is realistic for the timeline and viable. In most cases, the right answer is somewhere between the two, with a phased strategy that delivers capacity quickly while preserving the option to scale.

3.2 Technical baseline: power, cooling, space and water

A liquid cooling transition begins with an assessment of what the existing facility can support. For most operators, four variables define the starting point.

Power is typically the most immediate constraint. AI workloads draw significantly more power per rack than traditional IT, and many existing facilities lack the capacity headroom, or the grid connection, to support a meaningful GPU deployment without upgrades. Medium-voltage readiness, redundancy levels and uninterruptible power supply (UPS) capacity all need to be reviewed against projected AI load.

Cooling infrastructure varies considerably between facilities. The relevant questions include: what is the current chilled water architecture, what is the PUE baseline, and are there existing constraints, such as chiller capacity, free cooling availability, and pipe routes that would limit the introduction of liquid cooling hardware?

Facilities with chilled water already circulating through the data hall are generally well positioned for liquid cooling integration. Those relying on computer room air handlers (CRAHs) may require more significant preparatory work.

Space and structure affect what is physically deployable. Liquid cooling introduces pipework, manifolds and coolant distribution units (CDUs) that require floor space, ceiling clearance and, in some cases, structural reinforcement. Rack layout and aisle containment configuration will also influence which cooling architecture is most practical.

Water is a consideration that is often underweighted. Liquid-to-liquid cooling requires a reliable supply of facility water. In regions with water stress, something that's increasingly common across southern and central Europe, operators need to understand their WUE baseline and the implications of different cooling approaches for water use.

3.3 Regulatory and ESG readiness

Alongside the technical baseline, operators should assess their position relative to current and forthcoming regulatory requirements. The EU EED obligations that took effect for qualifying facilities include mandatory data collection on PUE, WUE, ERF, and carbon usage effectiveness (CUE). While many operators are collecting some of this data already, fewer have the systems in place to report it consistently and publicly.

Understanding the gap between current data collection and reporting obligations is a practical first step. It also helps identify where liquid cooling can deliver compliance value; for example, by reducing PUE through more efficient heat removal, or by enabling waste heat recovery that improves ERF performance.

Internal ESG commitments are a related consideration. Many organizations have net-zero targets or Scope 2 emissions reduction goals that are increasingly difficult to meet as AI workloads scale on air-cooled infrastructure. Liquid cooling can form part of a credible pathway toward those targets, particularly when combined with renewable energy sourcing and heat reuse.

3.4 Delta's assessment and advisory services

Delta supports operators through this assessment phase with site audits, computational fluid dynamics (CFD) modeling, and scenario analysis.

CFD modeling maps airflow patterns, identifies hot spots and tests the impact of density increases before any physical changes are made. TCO analysis compares air-only, hybrid, and full liquid cooling architectures over a 10-year horizon, incorporating energy costs, maintenance and anticipated regulatory exposure.

The output is a structured readiness assessment: a clear view of where the facility stands, what the options are, and what a phased transition roadmap could look like given the operator's timeline and business objectives.

3.5 The options: a structured overview

Several liquid cooling architectures are available, each suited to a different combination of density requirements, facility constraints and operational preferences.

Rear-door heat exchangers (RDHx) attach to the back of existing racks and use circulating water to capture heat as it leaves the servers. They require no modification to the IT equipment itself and can coexist with existing air cooling infrastructure, making them a practical first step for operators moving beyond traditional air cooling in an existing data hall.

Liquid-to-air cooling routes coolant directly to cold plates mounted on CPUs and GPUs, removing heat at the source rather than after it has already entered the air. This approach delivers significantly better thermal performance than RDHx for high-density AI racks and is well-suited to GPU clusters, accelerator pods and dedicated AI workloads. Delta's GoCool-225 liquid-to-air CDU, for example, supports up to 225 kilowatts per rack in this configuration.

Liquid-to-liquid cooling uses a secondary circuit of chilled facility water to cool the primary coolant loop coming from the chip. Delta's GoCool-3000 CDU is designed for this architecture, offering precise control of flow, pressure and temperature with stainless-steel plumbing and integrated filtration. This approach is standard for large-scale new builds where chilled water infrastructure can be designed in from the outset.

Immersion cooling submerges IT hardware directly in a dielectric fluid, which absorbs heat with high efficiency. It enables the highest rack densities of any current approach and is well-suited to specialized AI or HPC environments where hardware standardization and long operational cycles are feasible. The trade-off is a more constrained hardware ecosystem and greater operational complexity.

Two-phase direct-to-chip cooling is technically promising for very high heat densities, using refrigerant that changes state to absorb heat. Products and standards are still developing, and most operators are best served by tracking this technology for inclusion in medium-term infrastructure roadmaps rather than deploying it at scale today.

Hybrid architectures combine liquid cooling for AI and GPU workloads with continued air cooling for lower-density legacy and enterprise loads. This is the most common configuration for facilities in transition, allowing operators to introduce liquid cooling progressively without disrupting existing infrastructure.

4. Liquid cooling technology landscape and how to make your decision

4.1 Decision framework: matching technology to use case

The right cooling architecture depends on a combination of factors that interact in ways specific to each facility. Three questions provide a useful starting structure.

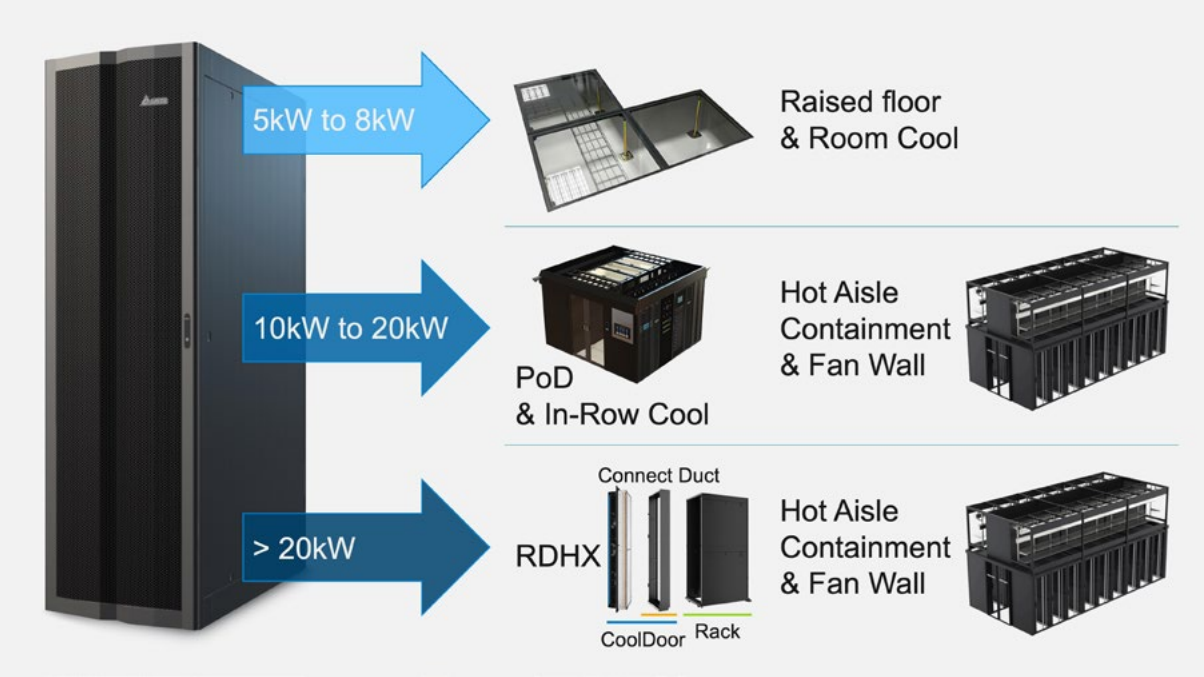

What is the expected heat load per rack? Racks drawing up to around 30 kilowatts can generally be managed with air cooling or RDHx. Loads between 30 and 160 kilowatts are well served by air-assisted liquid cooling. Direct-to-chip becomes the practical option above that threshold, with immersion relevant for the most demanding configurations above 80 kilowatts where hardware standardization allows it.

Is this a new build or a retrofit? New builds allow cooling architecture to be designed in from the ground up: liquid-to-liquid systems with chilled water distribution, high-voltage DC power and pre-planned manifold routes. Retrofits require working within existing constraints: available floor space, existing chilled water or CRAH systems, and the need to maintain continuity of operations during transition. Liquid-to-air solutions, which do not require facility-side chilled water, are often the more practical choice in retrofit scenarios.

What are the primary constraints and priorities? Operators with limited water availability may prefer liquid-to-air over liquid-to-liquid. Those with strong ESG commitments around waste heat reuse should consider how different architectures affect output water temperatures and heat recovery potential. Facilities with limited in-house fluid management expertise may prioritize solutions with simpler operational requirements and strong vendor support.

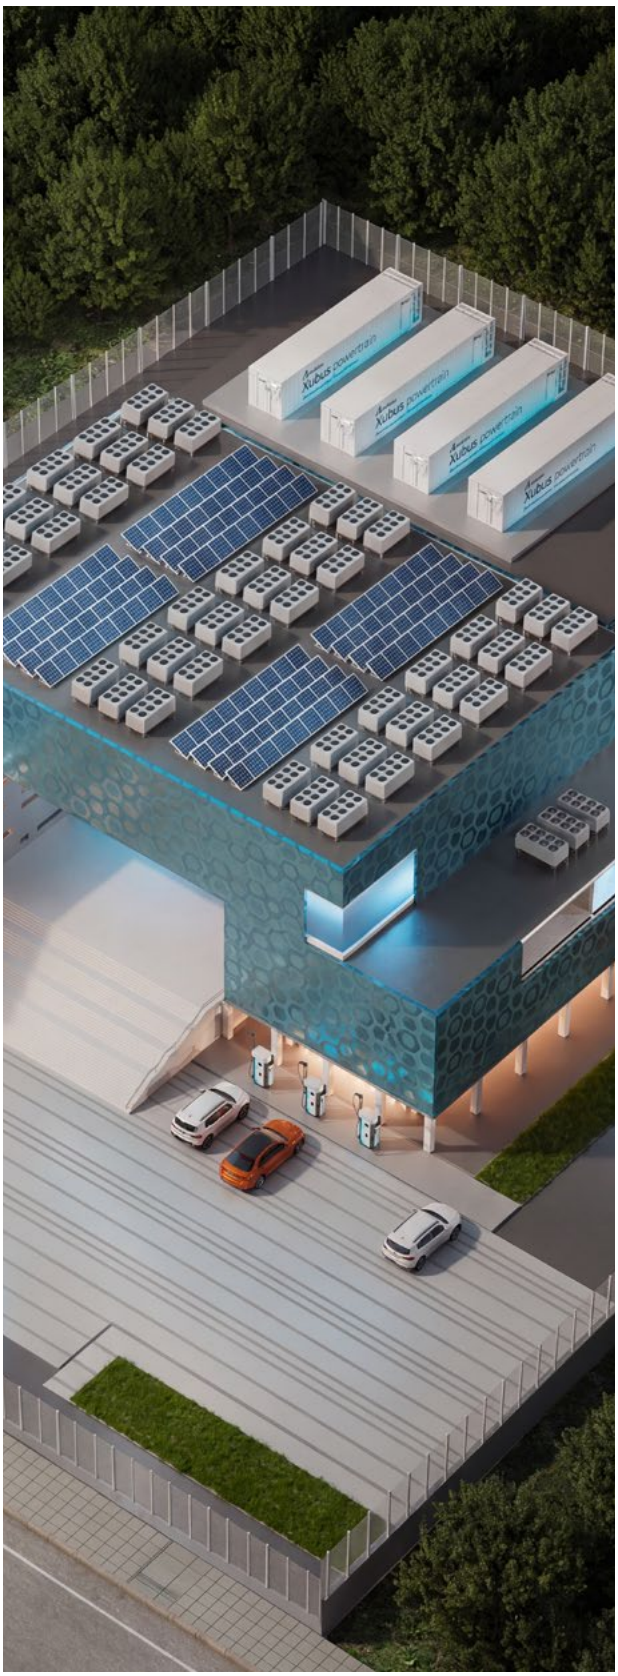

4.2 Where each technology fits in practice

Rear-door heat exchangers are a sensible starting point for facilities with moderately dense racks, up to around 50 kilowatts, where the operator wants to improve cooling performance without redesigning the data hall or modifying IT hardware. They work well alongside existing air infrastructure and can be deployed incrementally.

Direct-to-chip liquid cooling should be the default consideration for dedicated AI or GPU pods. It delivers the thermal performance that high-density accelerator workloads require and is compatible with current GPU generations from NVIDIA and others. For operators retrofitting a section of an existing facility to support AI (say, a 10-megawatt data center adding a single AI), rack liquid-to-air direct-to-chip is typically the practical choice, as it avoids the need for facility-side chilled water plumbing.

Immersion cooling is worth evaluating when rack densities are at the upper end of what direct-to-chip can handle, or when a new-build environment allows for the hardware standardization it requires. It is not well suited to mixed or variable-workload environments.

Two-phase cooling is best considered in the context of future planning. Operators designing facilities now should size power and space envelopes with the possibility of two-phase adoption in mind, without committing to it before the product and standards landscape matures further.



Delta's GoCool-225 L2A CDU supports up to 225 kW per rack with no raised floor or facility piping required

4.3 Delta's liquid cooling portfolio

Delta's liquid cooling portfolio spans the range of architectures described above. The GoCool-225 liquid-to-air CDU supports direct-to-chip cooling at up to 225 kilowatts per rack, with a closed-loop design that requires no raised floors or facility-side piping, well-suited to retrofit deployments. Stainless steel plumbing and integrated 50-micron filtration maintain fluid quality throughout the system lifecycle. This is crucial for AI servers, as even a speck of dust can clog the microchannels on the cool pad inside a server and cause malfunctions. Customers can also choose a finer 25µm filter.

The GoCool-3000 liquid-to-liquid CDU addresses larger-scale requirements, with dual power feed, redundant components and a 10-inch touchscreen interface for real-time monitoring. It supports racks above 100kW that produce output temperatures of between 40–60°C. It is also suitable for waste heat recovery down the line, for example for district heating. The GoCool-3000 has a very high discharge head to achieve the required pumping pressure for both NVIDIA's AI server racks, requiring pressures of 30-35PSI, and Google's, which need between 80-90PSI.

Both products integrate with Delta's broader data center infrastructure: racks, busways, PDUs, UPS systems and monitoring platforms, allowing cooling and power to be managed as a unified system rather than separate disciplines.

This integration is particularly relevant for operators planning AI deployments at scale, where the interaction between power delivery, cooling performance and compute efficiency has a material impact on both reliability and operating cost.



	LZL CDU GoCool-1500+	L2L CDU GoCool 3000
Cooling Capacity	1,500kW @ 1.0kW/LPM 1,000kW @ 1.5KW/LPM	3,000kW @ 1.0kW/LPM 2,000kW @ 1.5KW/LPM
Coolant Type	Primary: Water Secondary: 25%PG	Primary: Water Secondary: 25%PG
ATD	5 °C	4 °C
External Pressure Drop	289 kPa (41.9psi)	334 kPa (48.4psi)
Dimension (WxDxH)	900 x 1200 x 2300 mm	1200 x 1500 x 2300 mm

Comparison of some Delta solutions

5. From concept to design

5.1 Power and electrical integration

Designing for AI-density racks requires a different approach to electrical infrastructure than most existing data centers were built to accommodate. At 50 to 100 kilowatts per rack and above, circuiting, redundancy and cable management all need to be revisited.

A rack drawing 100 kilowatts requires roughly five times the power feed of a conventional 20-kilowatt rack, and the switchgear, busways, and PDUs serving it need to be sized accordingly.

For operators planning AI pods or islands within a larger facility, a busway-based power distribution approach offers practical advantages: it allows power to be routed overhead to rack positions without floor-level disruption, and it can be reconfigured as rack positions change. Delta's busway and PDU systems are designed to support this kind of modular, high-density layout, integrating with UPS and monitoring infrastructure to give operators visibility across the full power chain.

Looking ahead, the transition toward 800-volt high-voltage DC (HVDC) power architecture is a relevant consideration for new builds. Current data centers convert power between AC and DC multiple times, each conversion introducing around 1% energy loss.

Moving to a single AC-to-DC conversion and distributing at 800V DC through the facility reduces those losses and improves overall energy efficiency. NVIDIA has published its HVDC architecture standard, and this approach is expected to become increasingly common in large-scale AI facilities over the coming years.

5.2 Thermal and hydraulic design

CDU selection and sizing should be determined by the selected rack density profile and the cooling architecture. For a direct-to-chip deployment, the CDU needs to match the thermal load of the racks it serves to maintain chip temperatures within operating range under peak load. Row-level CDUs serving four to eight racks require careful flow balancing across the manifold to ensure each rack receives consistent coolant supply.

Piping architecture influences both performance and maintainability. Redundant loop designs, where each CDU is connected via two independent supply and return paths, allow for isolation and maintenance without taking the cooling system offline. Thermal insulation on primary facility network pipework is important where coolant temperatures approach ambient dew point, to prevent condensation forming on pipe surfaces.

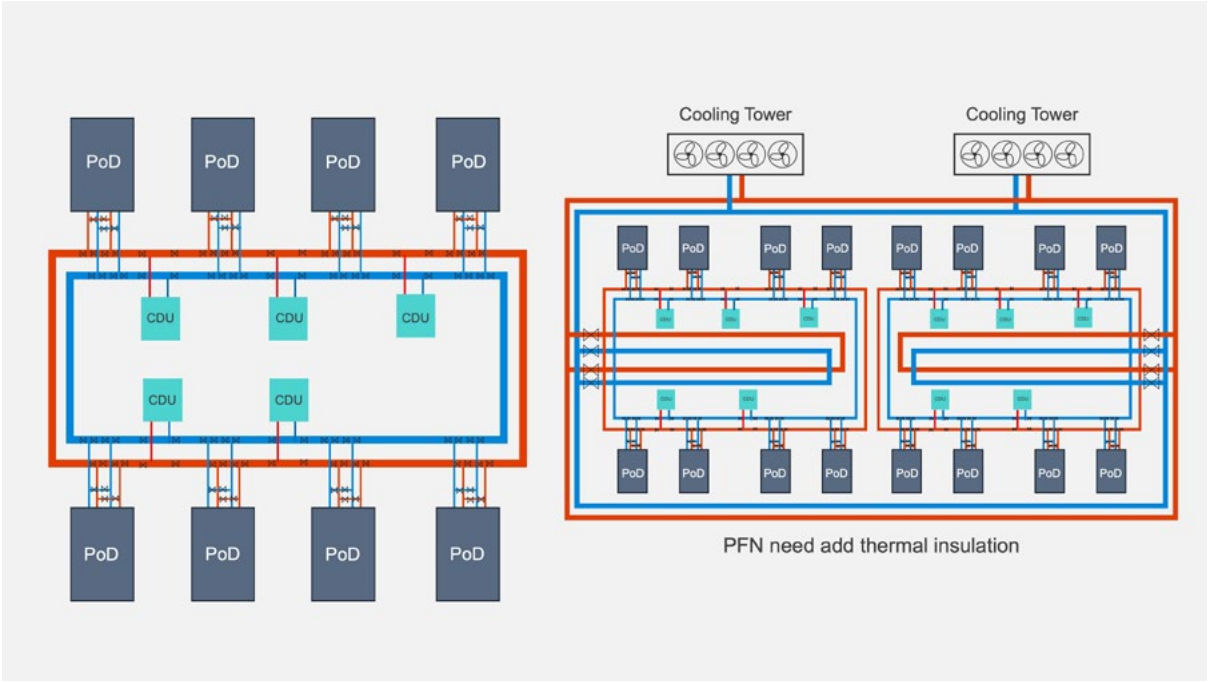

Control matters as much as hardware selection. Delta's auto-tuning algorithm for its air-assisted liquid cooling (AALC) systems automatically optimizes valve and pump parameters to maintain target temperatures with less than 5% dynamic response variation. This means that when one CDU unit is taken offline for maintenance, the remaining units quickly compensate for the change, minimizing any impact on IT racks.

5.3 Materials, fluids and compatibility

The choice of fluids and the materials it comes into contact with are decisions that affect both performance and long-term maintenance costs. Water-glycol mixtures are the most common choice for direct-to-chip systems. Glycol provides freeze protection and inhibits corrosion, and the mixture is straightforward to manage. Deionized water is used in some high-performance applications but requires more careful quality monitoring. Engineered dielectric fluids are used in immersion systems and some two-phase applications.

Material compatibility across the coolant loop, metals, elastomers, plastics and any galvanic interfaces should be reviewed at design stage. Copper and aluminum components in the same loop without appropriate inhibitor chemistry can cause corrosion over time.

Delta's CDUs use stainless steel plumbing and integrated 50-micron filtration to maintain fluid quality throughout the system lifecycle.

5.4 Space planning and constructability

Liquid cooling introduces physical infrastructure that needs to be accommodated within the data hall layout: CDUs, manifolds (the pipework headers that distribute coolant to individual racks), flexible hoses and pipe routes. Safety clearances around CDUs need to be incorporated from the outset rather than resolved after installation to avoid piping making direct contact with racks and simplify maintenance access. Similarly, safety equipment such as drip trays, lead detection cables and zoned valve control for automatic shutdown must be factored in at the design stage.

In retrofit scenarios, working around existing raised floors, cable trays and aisle containment structures poses additional challenges. Floor loading may also be a consideration as CDUs and associated pipework add weight that may need to be assessed against structural limits.

For both new-builds and retrofits, data center operators often prefer overhead to underfloor piping. Not only does it not interfere with existing structures, but it is also more secure, as any leakage will be easier to detect and address than with underfloor pipes. Furthermore, with the growth of high-density applications, underfloor space might be needed for additional electrical cabling and double-layer piping specified for safety. With limited room underfloor, overhead piping is increasingly becoming the standard, especially for the global hyperscalers.

5.5 Designing for heat reuse from day one

Liquid cooling recovers heat at temperatures that vary by architecture and that determine what downstream reuse is practical. For direct-to-chip systems, output water temperatures typically range from 40 to 60 degrees Celsius, warm enough for supply to district heating networks operating at low-temperature standards, space heating, or domestic hot water preheat.

Liquid-to-liquid systems operating at lower supply temperatures may require a heat pump to upgrade output temperatures to levels suitable for reuse. Where heat reuse is an objective, the architecture choice is therefore part of the heat recovery design, not a separate consideration.

The regulatory case for planning this from day one is also worth noting. The EU Energy Efficiency Directive requires qualifying facilities to report on energy reuse factor, the proportion of energy that is recovered and reused rather than wasted. Facilities that have been designed for heat recovery from the outset will have a materially stronger ERF position than those attempting to retrofit heat exchange infrastructure into an existing system. As with most aspects of liquid cooling design, the cost and complexity of getting this right at the start is considerably lower than addressing it later.

Where heat reuse is an objective, the design implications are best addressed early: heat exchanger sizing, control schemes for variable heat export and interface requirements with external energy partners all affect the facility design in ways that are more practical to plan for than to retrofit later.

Start with a solid foundation

- Computational Fluid Dynamics (CFD) Modeling and Analysis
- Aisle Containment
- Blanking Panels
- Brush Grommets
- Humidification / Dehumidification
- Proper air filtration
- Perforated Floor Tile flow rate audits with anemometer

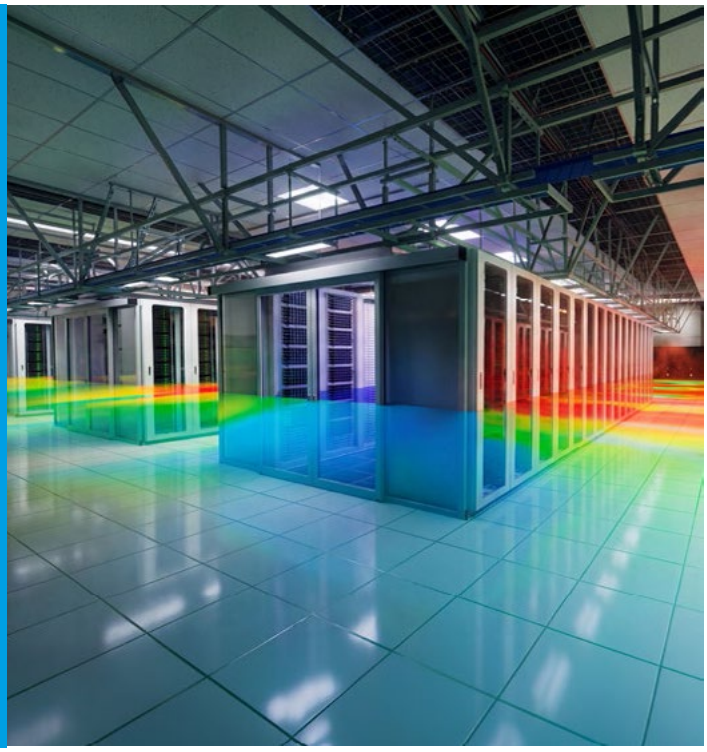

6. Implementation, commissioning and phased deployment

6.1 Phased deployment strategies

For most operators, a phased approach to liquid cooling is more practical than a facility-wide transition.

Starting with a pilot rack or pilot pod of typically four to eight racks allows the team to develop operational familiarity, identify any site-specific integration challenges, and validate performance assumptions before committing to a larger rollout.

A well-scoped pilot defines success criteria upfront: target rack temperatures under load, coolant flow rates, delta-T across the CDU, and any PUE or WUE improvement targets. Lessons from the pilot, on installation sequence, commissioning time, and operational procedures, can then be applied to subsequent phases with greater confidence and efficiency.

Scaling from a pilot pod to a full AI zone or facility-wide rollout introduces coordination requirements that go beyond the technical. IT teams and facilities teams typically operate on different planning cycles and with different risk tolerances. Aligning them around a shared deployment schedule is as important a project management task as the physical installation work.





6.2 Commissioning best practices

Commissioning a liquid cooling system follows a defined sequence that should not be compressed under time pressure. Pressure testing and flushing the pipework before connecting to IT equipment is essential, as contaminants introduced at this stage can cause long-term reliability issues. Leak testing at operating pressure, with adequate dwell time before sign-off, is a non-negotiable step.

Control loop tuning should be carried out under representative IT load conditions, and this includes setting pump speeds, valve positions and temperature setpoints.

A system tuned at low load may behave differently when racks are fully populated. Staged load tests, including transient tests that simulate rapid load changes, verify that the control system responds stably and that chip temperatures remain within specification.

Documentation and handover should include as-built drawings of the pipework and manifold layout, commissioning test records, fluid analysis results and operating procedures for routine maintenance and emergency response. This documentation supports both ongoing operations and any future modifications to the system.

6.3 Minimizing risk and downtime during transition

In live facilities, the transition from air to liquid cooling needs to be managed carefully to avoid impact on running workloads.

Where possible, liquid cooling infrastructure should be installed and pre-commissioned before any air cooling systems are decommissioned, operating both in parallel through the cutover period. Workload migration can then be sequenced to transfer compute to the liquid-cooled racks before the air-cooled hardware is taken offline.

Modular and prefabricated components, including pre-assembled manifold sections, factory-tested CDU skids and pre-terminated hose kits, reduce installation time in the data hall and the associated exposure window. Delta's AALC manifold and hose kit systems are designed with this in mind, supporting faster deployment without compromising system integrity.



6.4 Training and preparing for change

Liquid cooling introduces new operational disciplines that most data center teams have not previously needed. Fluid handling, leak response procedures, coolant quality monitoring and the interaction between cooling system behavior and IT workload are all areas where training is needed.

Safety procedures deserve specific attention. A pressurized coolant circuit carrying water-glycol mixture at operating temperature presents hazards that are different in character from the electrical risks that data center teams are more familiar with. Spill response protocols, including how to isolate the affected zone, contain fluid safely and manage the impact on running IT equipment, should be documented, rehearsed and visible to all staff working in the data hall before the system is commissioned. Personal protective equipment requirements for routine fluid handling and maintenance tasks should also be established as standard practice.

The cultural dimension is equally important. IT and facilities teams that have historically operated independently need to develop a shared understanding of how the integrated system behaves, particularly when a cooling event affects compute performance, or a workload change drives an unexpected cooling demand. Delta's monitoring platforms and operational dashboards support this visibility, giving both teams access to the same real-time data on flow, temperature, pressure and power.



7. Operating a liquid-cooled data center: risk, monitoring and security



7.1 Operational risk management

Leak prevention begins at the design and component selection stage. High-quality fittings, appropriate materials for the fluid and pressure conditions, and redundant containment design all reduce the probability of a leak occurring. Where a leak does occur, the consequences depend heavily on how quickly it is detected and how the system responds.

Leak detection sensors provide early warning before fluid reaches IT equipment. Zoning the detection system allows the source of a leak to be quickly identified and isolated without shutting down the entire cooling loop. Delta's CDUs incorporate configurable leak response logic: on detection, the system can reduce flow, trigger alarms, or initiate a controlled shutdown of the affected zone, depending on the operator's defined response procedure.

Condensation is a risk in facilities where coolant supply temperatures approach the ambient dew point. Managing supply water temperature setpoints, typically keeping supply water above the dew point of the surrounding air, is the primary control measure.

Where this is not always possible, insulation on pipework and careful environmental monitoring in the data hall provide additional protection.

7.2 Monitoring and analytics

Effective operation of a liquid-cooled data center depends on real-time visibility across the key parameters of the cooling system: flow rate, supply and return temperatures, pressure differential across the manifold, pump status and filter condition. These telemetry streams, correlated with IT load data from power meters and server management systems, give operators a complete picture of how cooling performance tracks against compute demand.

Delta's CDU systems provide this telemetry through integrated sensors and a 10-inch color touchscreen interface, with data available for export to building management systems (BMS) or data center infrastructure management (DCIM) platforms.

This integration supports the kind of predictive maintenance approach that reduces unplanned downtime, identifying trends in filter pressure drop, pump performance degradation, or temperature drift before they develop into failures.

7.3 Cybersecurity for cooling and power infrastructure

Cooling control systems are networked infrastructure, and they carry cybersecurity risk.

A compromised CDU controller could affect chip temperatures, trigger unnecessary shutdowns, or provide a foothold for broader network access. As EU NIS Directive obligations extend to critical infrastructure operators, including large data centers, cybersecurity for building systems and operational technology is an area of growing regulatory scrutiny.

Network segmentation between cooling control systems and general IT networks is a foundational measure. Authentication controls, secure remote access protocols, and regular firmware updates for CDU and BMS components are further elements of a sound operational security posture. Delta's monitoring and control platforms are designed with these considerations in mind, supporting secure access and integration within the operator's wider security architecture.

7.4 Water, carbon and environmental performance

Liquid cooling changes the water and carbon profile of a data center in ways that need to be understood and monitored.

Direct-to-chip systems that use closed-loop water circuits without evaporative cooling can reduce water consumption significantly, relative to air-cooled facilities that rely on cooling towers. WUE, the ratio of water consumed to IT load, is a key metric under the EU EED reporting framework, and liquid-cooled facilities are generally better positioned on this measure.

Carbon usage effectiveness (CUE), measuring the greenhouse gas emissions associated with data center energy consumption, is improved indirectly through the energy efficiency gains that liquid cooling enables.

Combining liquid cooling with renewable energy sourcing and waste heat recovery provides a route to demonstrating meaningful progress against ESG commitments, with the monitoring data to support credible reporting.

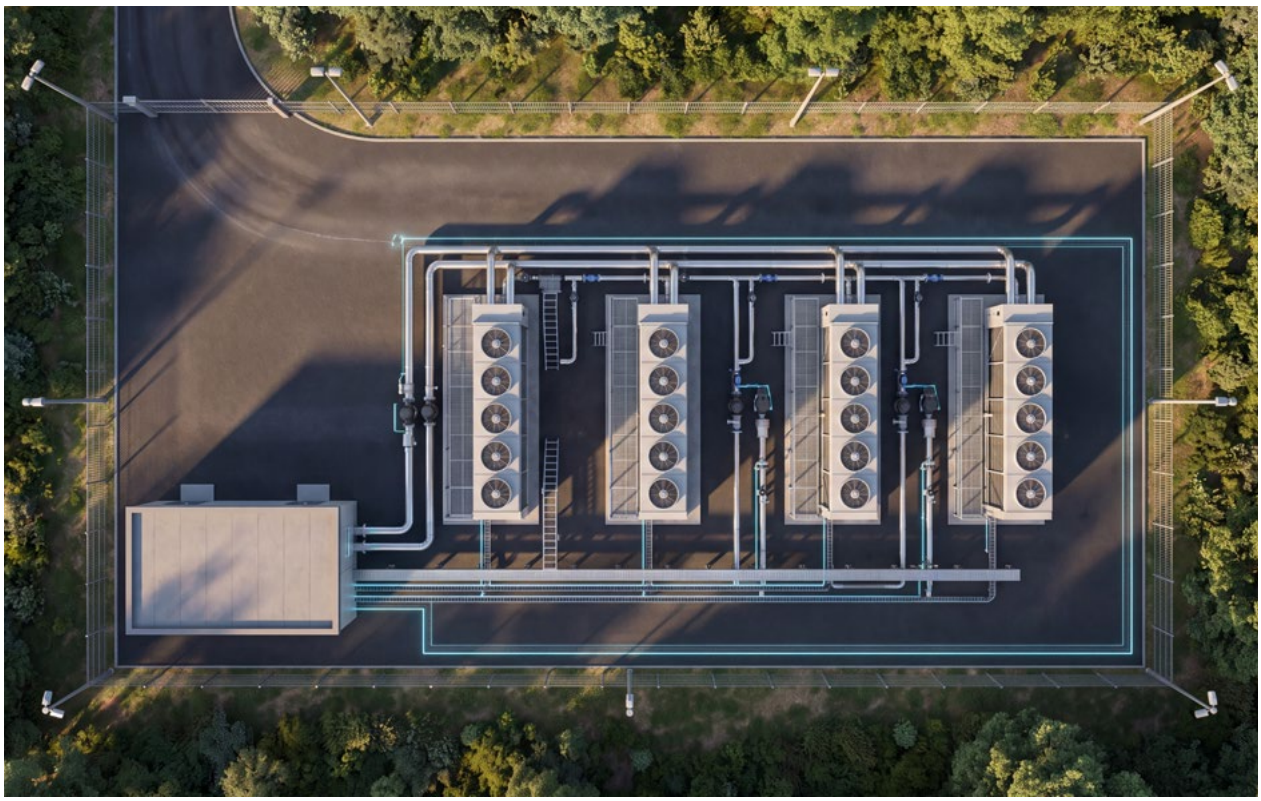

8. The business case: TCO, compliance and heat reuse economics

8.1 TCO comparison framework

The business case for liquid cooling is most usefully framed around TCO (total cost of ownership) over a 10-year horizon.

The upfront investment in CDUs, manifolds, pipework, controls and integration is real, but it also needs to be set against the operating cost profile of the alternative, and the costs of not transitioning.

On the capital expenditure (CAPEX) side, the key variables are the cooling architecture chosen, the scale of deployment, and how much existing infrastructure can be reused. Retrofit deployments that use existing chilled water connections or room cooling infrastructure have lower incremental CAPEX than new-build liquid-to-liquid systems.

On the operating expenditure (OPEX) side, liquid cooling's primary advantage is energy: CDUs serving high-density GPU racks consume considerably less power than the air handling equipment that would otherwise be needed to manage equivalent heat loads.

The primary economic driver is revenue enablement. For most AI operators, the key challenge is maximizing compute efficiency, producing more tokens per kilowatt-hour from their existing or planned infrastructure. Liquid cooling helps remove the thermal limits that would otherwise restrict GPU performance and density, making it essential rather than optional for running current-generation AI hardware at scale.

8.2 Compliance as value

The EU EED's mandatory reporting requirements apply to data centers with a total rated power of one megawatt or more, covering PUE, WUE, ERF and other KPIs. Operators that do not have systems in place to collect and report this data face compliance risk. Those that can demonstrate strong performance on these metrics are better positioned with regulators, customers and investors.

Early adoption of liquid cooling simplifies this compliance picture. A facility with liquid cooling in place will typically have lower PUE and better WUE than an equivalent air-cooled facility. The monitoring infrastructure associated with liquid cooling systems also provides much of the telemetry needed for EED reporting as a byproduct of normal operations.

The avoided cost of retrofitting compliance infrastructure later is part of the TCO case for moving sooner rather than later.

8.3 Monetizing waste heat

Revenue or cost-avoidance models for waste heat reuse vary by context. Where a district heating network connection is available, the operator may be able to sell recovered heat under a long-term supply agreement, generating a revenue stream that offsets part of the facility's operating cost. Where heat is reused on-site, for office heating, on-site industrial processes, or server room preconditioning, the benefit appears as avoided energy expenditure.

In both cases, the financial model benefits from being designed in from the start, since the heat exchanger sizing, output temperature targets and interface with external networks all affect the facility design.

8.4 Working with Delta on the business case

Delta supports customers in developing the financial case for liquid cooling transitions, drawing on operational data from its portfolio of deployments.

This includes scenario modeling across air-only, hybrid and full liquid architectures, sensitivity analysis against energy price and carbon price assumptions, and guidance on how to structure the business case for internal approval processes. The output is not a generic model but an analysis grounded in the specific site, workload and regulatory context of the operator's facility.

9. Reference architectures

9.1 Reference architectures

The following reference architecture patterns reflect the design approaches Delta works with in practices.

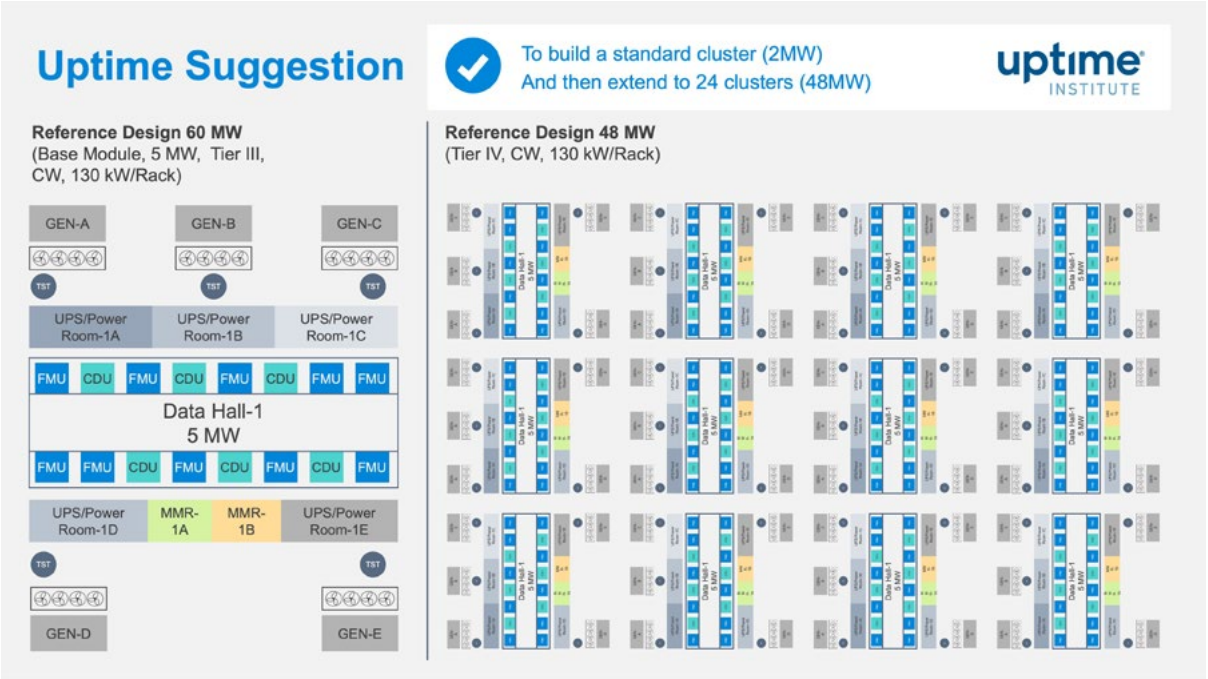

MW AI cluster phased build. The starting point for many operators entering AI infrastructure. Designed around a standard cluster of eight GPU cabinets drawing approximately one megawatt, with a single data hall module, dedicated CDU row, and power distribution sized for expansion.

The Uptime Institute's suggested approach is to build a validated two-megawatt cluster first, then extend to 24 clusters at 48 MW, each additional module replicating the same baseline design. This phased model allows operators to validate the architecture, build operational familiarity, and scale incrementally without redesigning the facility.

50 to 60 MW AI data center new build. At this scale, the facility is designed entirely around AI workloads: all white space allocated to IT racks, CDUs and fan walls; UPS and battery systems consolidated into power train units (PTUs) to reduce delivery time; and chilled water architecture using separate low-temperature chillers for fan wall support and high-temperature chillers for CDU (liquid-to-liquid) support. Power is distributed on a 2 MW base module basis, allowing the facility to be built and commissioned in stages.

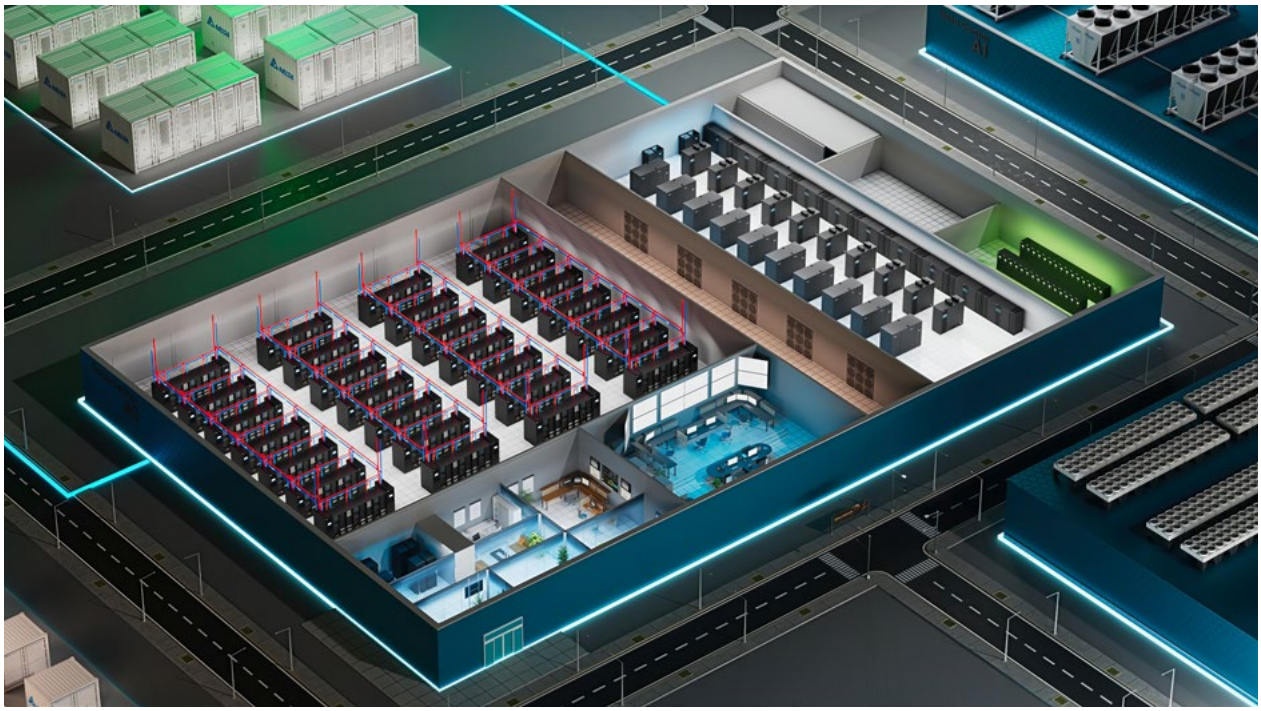

Hybrid enterprise retrofit. An existing facility adding an AI pod within existing white space. The preferred approach uses liquid-to-air CDUs, no raised floor or facility-side piping required, allowing the AI zone to be deployed without disrupting the remainder of the data hall. Existing room cooling infrastructure continues to serve non-AI racks. This is the most common pattern for European operators looking to add AI capability quickly.

How Delta can help



Delta's position in the data center infrastructure market is grounded in the breadth of what it designs and manufactures in-house.

Power systems, liquid cooling infrastructure, racks, busways, PDUs, UPS and monitoring platforms are all developed within the same organization, which means they are designed to work together.

This matters in AI deployments, where the interaction between power delivery, cooling performance and compute efficiency has a direct bearing on both reliability and operating cost.

The company employs **12,000 engineers** and holds **15,000 patents**, with **8%** of revenue reinvested in R&D.

When Delta collaborates with major technology partners, including GPU manufacturers and hyperscale operators, it retains intellectual property ownership of what it develops.

This R&D depth underpins products like the GoCool-225 L2A CDU.

In Europe, Delta Electronics EMEA operates 15 R&D centers, 33 sales offices and four manufacturing plants, with its EMEA headquarters in Hoofddorp, Netherlands.

This local presence supports the kind of on-the-ground engagement that large infrastructure projects require, from initial site assessment through to commissioning and ongoing operations.

Engagement journey with Delta

Delta's engagement with customers typically follows four stages.

1. Discovery and assessment

Delta conducts a site audit covering power capacity, existing cooling architecture, space and structural constraints, and regulatory position. CFD modeling maps current airflow and identifies how the facility would perform at higher rack densities. The output is a readiness assessment and a range of scenarios for how a liquid cooling transition could be approached.

2. Co-design and solution selection

Working from the readiness assessment, Delta helps the customer select a cooling architecture, define the phasing plan and develop the financial model. This includes TCO analysis comparing air-only, hybrid and liquid approaches over a 10-year horizon.

3. Delivery and commissioning

Delta manages the project through installation, pre-commissioning checks, control loop tuning and staged load testing, with full documentation and handover on completion.

4. Operations and lifecycle partnership

Delta's monitoring platforms provide ongoing visibility of cooling and power performance, supporting predictive maintenance and operational efficiency. Service agreements cover fluid quality management, component inspection and firmware updates across the system lifecycle.

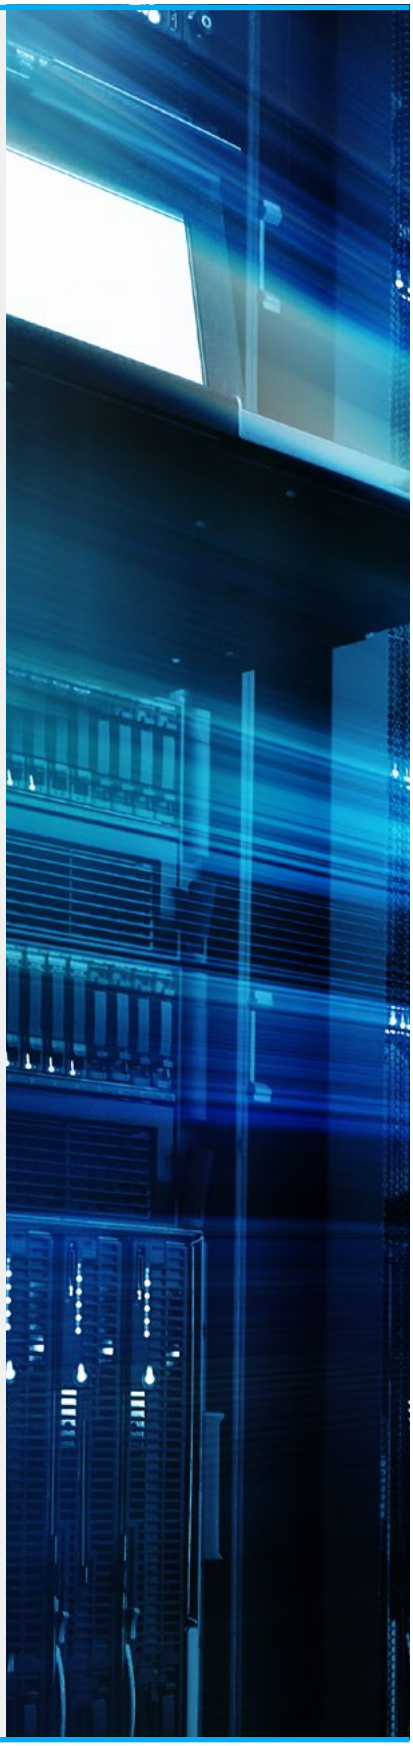

Unlocking value with liquid cooling

The transition to liquid and hybrid cooling is an urgent, present-day requirement for European data center operators.

As AI and HPC workloads continue to drive rack densities beyond the reach of air cooling, and as EU regulatory obligations tighten around energy efficiency, water use and heat recovery, the question for operators is how to choose the best path.

The good news is that the technology is mature, the business case is clear, and the pathway to operation is well established. Done right, liquid cooling can be a strategic investment that improves TCO, reduces carbon exposure, unlocks new revenue from waste heat and positions a facility for the next decade of AI infrastructure demand.

Delta's integrated approach to power, cooling, monitoring and heat reuse gives operators a single, accountable partner for this journey, one with the engineering depth, European presence and proven track record to help them move from planning to performance with confidence.



To discuss your AI data center cooling needs, please contact us:

Delta EMEA Headquarters

Netherlands B.V.

Zandsteen 15, 2132MZ, Hoofddorp
The Netherlands
Tel: +31 (0)20 800 3900

www.delta-emea.com

Delta Telecom Power - EMEA Locations

Delta Electronics (Germany) GmbH

Ferdinand-Porsche-Str. 45, 60386
Frankfurt am Main, Germany
+49-69-42002-0

Delta Electronics (Switzerland) AG

Freiburgstrasse 251, 3018 Bern-
Bümpliz, Switzerland
+41-31-998-53-11

Delta Solutions (Finland) Oy

Rajatorpantie 8, FI-01600 Vantaa,
Finland
+358-9-849-660

Delta Electronics (France) SAS co.

2 Rue du 19 Mars 1962, Zi Bastillac
Nord, 65000 Tarbes, France
+33 562 34 09 30

Delta Electronics (Italy) S.R.L

Building Spaces Eur Arte, unit 508 -
5F, Viale dell'Arte, 25, 00144 Rome,
Italy
+39 06 9931 0867

Delta Electronics (Poland) Sp. z.o.o.

23 Poleczki Str. 02-822 Warsaw
Poland
+48-22-335-2600

Delta Electronics (Czech Republic), spol. s r.o

Průmyslová 1306 /7, 102 00 Praha 10,
Czech Republic
+420 272 019 330

Delta Electronics Solutions (Spain) SLU

Ctra. De Villaverde a Vallecas, 265
1º Dcha Ed. Hormigueras – P.I. de
Vallecas 28031 Madrid
+34 91 223 74 20

Delta Electronics (UK) Ltd.

Hemel Hempstead, Hertfordshire
HP2 7EY,
United Kingdom
+44 (0)1442 219355

Delta Energy Systems MEA (South Africa)

Tuinhof Office Park, Unit C401, Karee
Building, 265 West Avenue, Centurion,
0157 South Africa

Delta Energy Systems (Sweden) AB

P.O. Box 3096 SE-350 33 Växjö
Sweden
+46 470 706 800

Delta Energy Systems AG - Dubai BR

P.O. Box 185668, Gate 7, 3rd Floor,
Hamarain Centre, Dubai, United Arab
Emirates
+971 4 2690148

Delta Electronics (Norway) AS

Terminalen 12-16, 3414 Lierstranda
P.O Box: 2340 Stromso, 3003
Drammen,
Norway
+47 3220 3200