

Data Centre Cooling Systems: A Practical Guide

Understanding Air Cooling, Liquid Cooling, Thermal Management, and Infrastructure Efficiency

GSDC PROFESSIONAL REFERENCE GUIDE

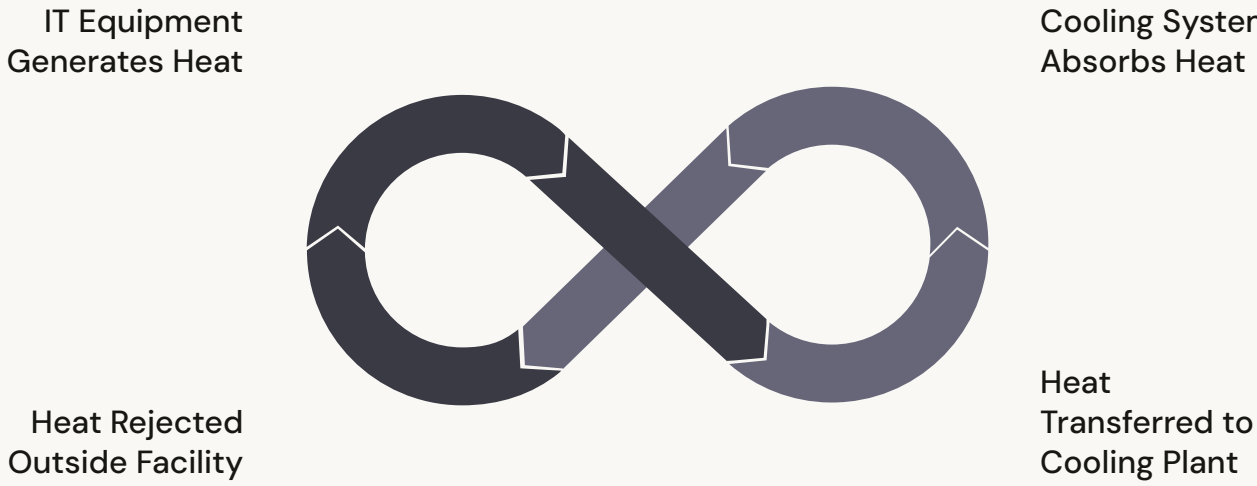
Understanding Data Centre Cooling

Cooling is one of the most operationally critical functions in any data centre. Every server, processor, storage device, and networking component generates heat as a by-product of electrical operation. Without a continuous and reliable means of removing that heat, internal component temperatures rise rapidly, leading to thermal throttling, hardware failure, and in severe cases, complete system shutdown.

The relationship between IT power consumption and heat generation is direct and predictable. For every watt of electrical power consumed by IT equipment, an approximately equivalent amount of heat energy is produced. A data centre drawing 1 megawatt of IT load will therefore generate close to 1 megawatt of heat that must be continuously extracted from the facility. As compute densities have increased — driven by virtualisation, artificial intelligence, and high-performance computing — the volume of heat generated per square metre of data centre floor space has grown correspondingly.

Modern servers operate within defined thermal envelopes. ASHRAE (American Society of Heating, Refrigerating and Air-Conditioning Engineers) publishes recommended inlet air temperature ranges for IT equipment, typically between 18°C and 27°C for most server classes, with some equipment tolerating higher inlet temperatures. When these thresholds are exceeded, processors reduce their operating frequency to lower heat output — a process called thermal throttling — which directly degrades application performance.

Cooling system failure, even partial failure, poses significant operational risk. A sudden loss of cooling capacity in a high-density environment can raise rack temperatures by several degrees within minutes. Without intervention, this progression leads to emergency shutdowns triggered by equipment firmware, hardware damage, and potentially data loss. For mission-critical facilities — financial institutions, healthcare providers, cloud platforms — even brief thermal events can have serious operational and financial consequences. Cooling infrastructure must therefore be designed with appropriate redundancy, monitoring, and response procedures.



How a Data Centre Cooling System Works

A data centre cooling system operates as a continuous thermodynamic cycle. Heat generated by IT equipment must be collected from within the data hall, transferred to cooling plant equipment, and ultimately rejected outside the facility — either into the atmosphere or into another thermal medium. Understanding this cycle is essential for anyone responsible for operating or specifying data centre infrastructure.

The cooling cycle begins at the IT equipment itself. Servers and other devices draw cool air across their internal components — or, in liquid-cooled systems, pass coolant through or around heat-generating components. The air or liquid absorbs heat and exits the equipment at a higher temperature. This warmed air or liquid is then directed towards cooling units that extract the heat and transfer it to a refrigerant or water circuit.

Primary Cooling Components

CRAC Units

Computer Room Air Conditioning units use direct expansion refrigerant to cool and dehumidify return air, returning conditioned air to the data hall floor.

CRAH Units

Computer Room Air Handling units use chilled water coils rather than direct expansion refrigerant, requiring a chiller plant to supply cold water.

Chillers

Mechanical refrigeration systems that produce chilled water for CRAH units and other cooling equipment. Available as air-cooled or water-cooled configurations.

Supporting Infrastructure

Cooling Towers

Reject heat from water-cooled chiller systems into the atmosphere through evaporative cooling. Require water treatment and regular maintenance to prevent biological contamination.

Fans and AHUs

Air handling systems distribute conditioned air throughout the facility. Variable speed drives allow fans to adjust output based on actual cooling demand.

Pumps and Heat Exchangers

Circulate chilled water or coolant through the system. Heat exchangers transfer thermal energy between fluid circuits without mixing the fluids themselves.

Heat rejection is the final stage of the cycle. In air-cooled chiller systems, heat is released into the outdoor air via condenser coils and fans. In water-cooled systems, cooling towers evaporate a small proportion of water to carry heat away from the condenser water loop. The cooled refrigerant or chilled water then returns to the cooling units to begin the cycle again. This process operates continuously as long as the data centre is in operation.

Air Cooling Systems

Air cooling remains the dominant cooling technology in the global data centre industry. It works by circulating conditioned air through the data hall, directing cool air across server inlet faces, and collecting the warmed exhaust air for treatment and return. When correctly designed and managed, air cooling is reliable, well-understood, and straightforward to maintain using widely available skills and equipment.

The two principal unit types used in air-cooled data centres are CRAC (Computer Room Air Conditioning) and CRAH (Computer Room Air Handling) units. CRAC units use a self-contained direct expansion refrigerant cycle to cool return air. They do not require an external chiller plant, making them suitable for smaller facilities or as standalone deployments. CRAH units, by contrast, use chilled water supplied by a central chiller plant. This arrangement is more energy-efficient at scale and provides greater flexibility for capacity management.

Hot Aisle / Cold Aisle

Server racks are arranged so that equipment inlets face into cold aisles and exhausts face into hot aisles. This prevents the mixing of conditioned supply air with hot exhaust air, improving cooling effectiveness.

Containment Systems

Cold aisle containment encloses the cold aisle with roof panels and end-of-row doors to prevent cool air from escaping. Hot aisle containment captures exhaust air and directs it back to cooling units more efficiently.

Raised Floor Distribution

Perforated floor tiles allow conditioned air from an underfloor plenum to rise into cold aisles. Tile perforation percentage and placement can be adjusted to manage airflow across different rack densities.

Overhead Distribution

Some facilities deliver conditioned air from overhead ducts rather than raised floors. This approach is often used in retrofit environments or where underfloor space is not available.

Air cooling becomes increasingly difficult to sustain as rack power densities rise. Standard air-cooled data centres are generally effective up to approximately 10–15 kW per rack, depending on airflow management quality. Beyond this threshold, maintaining adequate inlet temperatures without localised hotspots becomes challenging. Blanking panels should be fitted in all empty rack slots to prevent recirculation of hot exhaust air back through the equipment. Regular airflow assessments and thermal surveys help identify gaps in containment that allow hot and cold air to mix.

Liquid Cooling Systems

Liquid cooling addresses the fundamental limitation of air as a heat transfer medium. Water has a heat capacity approximately 3,500 times greater than that of air by volume, meaning that a relatively small flow of coolant can remove substantially more heat than a large volume of moving air. This thermal advantage makes liquid cooling increasingly relevant as rack power densities associated with GPU computing, AI inference, and high-performance workloads exceed the practical limits of air-only systems.

A liquid cooling system typically consists of a coolant distribution unit (CDU), a network of supply and return piping, manifolds that connect to individual servers or racks, and heat exchangers that transfer heat from the IT coolant loop to the facility water or chiller circuit. The CDU is the central component: it conditions the coolant to the correct supply temperature and pressure, monitors flow rates, and provides alarms and controls for the liquid circuit.



Coolant Types

Most direct liquid cooling systems use deionised water or water-glycol mixtures. Dielectric fluids such as mineral oils or engineered fluorocarbon fluids are used in immersion cooling where electrical isolation is required.



Heat Transfer at Source

By bringing the cooling medium directly to or around the heat-generating component — rather than relying on airflow to carry heat across a server chassis — liquid cooling achieves far higher heat removal rates per unit volume.



Facility Requirements

Liquid cooling requires piping infrastructure within the data hall, appropriate leak detection systems, secondary containment where required, and staff trained in fluid system maintenance and leak response procedures.

Liquid cooling is particularly well-suited to AI training clusters, high-frequency trading environments, genomics computing, and other workloads that place sustained, high thermal loads on individual processors. It is not universally necessary for all data centre environments: standard enterprise servers, storage arrays, and networking equipment continue to operate effectively within well-managed air-cooled facilities. The decision to introduce liquid cooling should be driven by actual thermal requirements rather than trend adoption.

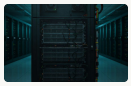
Types of Liquid Cooling Technologies

Liquid cooling is not a single technology but a family of approaches, each suited to different equipment types, facility configurations, and operational requirements. The three most widely deployed liquid cooling methods in commercial data centres today are direct-to-chip cooling, rear-door heat exchangers, and single-phase immersion cooling.



Direct-to-Chip Cooling

Cold plates — precision-engineered aluminium or copper blocks — are mounted directly onto CPUs, GPUs, and other high-heat components. Coolant circulates through internal channels within the cold plate, absorbing heat at the component surface and carrying it away to the CDU. The surrounding server chassis continues to use air cooling for lower-power components such as memory, storage, and power supplies. This approach, sometimes called hybrid liquid cooling, can support rack densities well above 50 kW in GPU-dense configurations. It is compatible with standard 19-inch rack infrastructure and does not require the server room to be resealed or isolated.



Rear-Door Heat Exchangers

A rear-door heat exchanger (RDHx) is a panel-mounted heat exchanger that replaces or supplements a standard rack rear door. Chilled water or coolant circulates through the door's internal coil matrix, absorbing heat from the warm exhaust air as it exits the rear of the rack. The air leaving the door is returned to the data hall at a lower temperature, reducing or eliminating the need for room-level cooling in that zone. RDHx systems are particularly useful for increasing the cooling capacity of existing racks without requiring server hardware changes, making them a practical retrofit option for facilities experiencing localised hotspots or density limitations.



Single-Phase Immersion Cooling

In single-phase immersion cooling, server hardware is fully submerged in a tank containing a non-conductive dielectric fluid — typically a synthetic or mineral oil-based liquid. Heat generated by the hardware is absorbed directly by the fluid, which circulates through the tank and passes through a heat exchanger to reject the heat to a facility water circuit. The fluid remains liquid throughout the process (hence "single-phase"), distinguishing it from two-phase systems where the fluid boils and condenses. Immersion cooling can support extremely high rack densities, eliminates the need for server fans, and provides effective cooling for components that are difficult to cool with cold plates alone. However, it requires hardware to be removed from standard rack enclosures, complicating maintenance access and hardware upgrade procedures.

Liquid Cooling vs Air Cooling

Selecting between air cooling and liquid cooling — or a combination of both — requires careful consideration of workload characteristics, existing infrastructure, hardware compatibility, and operational capabilities. Neither technology is universally superior: each has specific strengths that make it appropriate for different data centre contexts. The table below provides a structured comparison across the dimensions most relevant to infrastructure decision-making.

Dimension	Air Cooling	Liquid Cooling
Heat Removal Method	Conditioned air circulated across server components via fans and room-level cooling units	Coolant or dielectric fluid in direct or close contact with heat-generating components
Cooling Capacity	Practical limit typically 10–15 kW per rack in well-managed environments	Capable of supporting 50 kW+ per rack depending on technology used
Rack Power Density	Suited to low-to-medium density standard enterprise workloads	Required for high-density GPU, AI, and HPC environments
Infrastructure Requirements	Cooling units, containment systems, raised floor or overhead distribution	CDUs, piping, manifolds, leak detection, secondary containment where required
Hardware Compatibility	Compatible with all standard server and networking equipment	Requires server models with cold plate support or immersion-compatible hardware
Maintenance Complexity	Well-understood; skills widely available; filter, fan, and refrigerant maintenance	Requires fluid management, leak monitoring, and specialist maintenance skills
Energy Consumption	Fan and compressor energy significant at high load; efficiency improves with containment	Lower fan energy; pumps required; potential for free cooling and higher efficiency at scale
Retrofitting	Straightforward in most existing facilities; containment can be added incrementally	Requires significant infrastructure investment; not all existing facilities can accommodate
Suitable Workloads	Standard compute, storage, virtualised environments, enterprise applications	AI training, GPU clusters, HPC, high-frequency trading, dense compute nodes

It is important to note that many modern data centres operate hybrid cooling environments — using air cooling for the majority of standard workloads whilst deploying liquid cooling selectively in zones or racks housing high-density GPU or AI servers. This approach allows facilities to manage infrastructure investment progressively rather than undertaking wholesale replacement of existing cooling systems.

Cooling High-Density and AI Data Centres

The rapid expansion of artificial intelligence training, large language model inference, and GPU-accelerated computing has fundamentally changed the thermal profile of many data centre environments. Where a standard enterprise server rack might draw 5–8 kW, a rack populated with modern AI accelerators — such as NVIDIA H100 or comparable hardware — can consume 40, 60, or more than 100 kW in a single rack footprint. This represents a qualitative change in cooling requirements, not simply an incremental increase.

GPUs and AI accelerators generate heat in a highly concentrated form. Unlike a rack of general-purpose servers where heat is distributed relatively evenly across a row, a cluster of GPU nodes produces intense, localised thermal loads. Conventional air cooling systems struggle with this concentration: even with optimised containment, the volume of conditioned air required to maintain acceptable GPU inlet temperatures at extreme power densities may exceed what is physically feasible in a standard data hall without substantial modifications to airflow infrastructure.

Hypothetical Example

A facility originally designed for 8 kW average rack density introduces a row of 10 GPU server racks, each drawing 40 kW — a total additional load of 400 kW in a small footprint. The existing CRAH units serving that zone were sized for 80 kW total. The resulting thermal deficit cannot be resolved through airflow management alone. The facility must either introduce in-row or rack-level liquid cooling, relocate the GPU row to a purpose-built high-density zone, or accept performance degradation from thermal throttling.

Key Considerations for AI Deployments

1 Assess Total Rack Power

Calculate actual per-rack power draw at full GPU utilisation, not nameplate maximum, to avoid over- or under-specification of cooling systems.

2 Plan for Hotspots

Use computational fluid dynamics (CFD) modelling or thermal surveys to identify where hotspots will develop before hardware is installed.

3 Verify Hardware Compatibility

Confirm whether GPU servers support liquid cooling cold plates and whether the required manifold connections are available from the server manufacturer.

4 Design for Hybrid Operation

Many AI clusters use direct-to-chip cooling for GPUs whilst retaining air cooling for networking, storage, and management servers within the same zone.

Data Centre Thermal Management

Effective thermal management is the ongoing operational discipline of monitoring, controlling, and adjusting temperature conditions throughout a data centre to ensure that IT equipment operates within safe and efficient thermal boundaries. It goes beyond the initial design of cooling systems to encompass continuous measurement, alerting, and response to changing conditions as workloads, hardware configurations, and environmental factors evolve.

Server Inlet and Outlet Temperatures

Measuring air temperature at the front (inlet) and rear (outlet) of server racks provides direct evidence of whether cooling is reaching equipment adequately. A large differential between inlet and outlet temperatures — typically above 20°C — may indicate insufficient airflow or localised congestion within the rack.

Humidity Monitoring

Relative humidity must be maintained within defined limits — typically 40–60% RH for most data centre environments. Excessively low humidity increases the risk of electrostatic discharge. Excessively high humidity can lead to condensation on equipment and within cable runs. CRAC units typically include dehumidification functionality.

DCIM Systems

Data Centre Infrastructure Management (DCIM) platforms aggregate data from temperature sensors, power meters, cooling unit controllers, and environmental monitors into a unified management interface. DCIM enables operators to identify developing thermal issues before they become operational events, and to plan cooling capacity changes in advance of hardware deployments.

Airflow Management

Proper airflow management includes ensuring blanking panels are fitted in all empty rack slots, containment systems are intact, and perforated floor tiles are positioned to deliver adequate airflow to each rack based on its actual power density — not distributed uniformly across a row.

Practical Thermal Management Checklist

1. Verify server inlet temperatures are within ASHRAE recommended ranges at all operational rack positions
2. Confirm all empty rack slots have fitted blanking panels; remove any cable runs that bypass containment barriers
3. Review DCIM or BMS temperature trend data for gradual temperature rises that may indicate cooling degradation
4. Conduct thermal walk-throughs or IR surveys quarterly to identify hotspots not visible in sensor data alone
5. Confirm cooling unit alarms and environmental thresholds are correctly configured and tested

Cooling Efficiency: PUE and WUE

Cooling systems are among the largest consumers of energy in a data centre, often accounting for 30–40% of total facility power consumption in conventionally designed facilities. Two industry-standard metrics — Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE) — provide quantifiable measures of how efficiently a data centre uses its total energy and water resources relative to the IT load it supports.

Power Usage Effectiveness (PUE)

PUE measures the ratio of total facility energy — including cooling, lighting, power distribution losses, and building systems — to the energy consumed solely by IT equipment.

$$PUE = \frac{\text{Total Facility Energy (kWh)}}{\text{IT Equipment Energy (kWh)}}$$

A PUE of 1.0 would represent a theoretically perfect facility where all energy is consumed only by IT equipment. In practice, most data centres operate between 1.2 and 2.0. A PUE of 1.5 means that for every 1 kWh consumed by IT equipment, an additional 0.5 kWh is consumed by facility infrastructure — predominantly cooling.

ⓘ Hypothetical Example: A facility consumes 2,500 kWh total, of which IT equipment accounts for 1,800 kWh. PUE = 2,500 ÷ 1,800 = **1.39**. This indicates a reasonably efficient facility, though there is scope to reduce cooling and infrastructure overhead further.

Water Usage Effectiveness (WUE)

WUE measures the volume of water consumed by the facility — primarily by cooling towers and evaporative cooling systems — per unit of IT energy delivered.

$$WUE = \frac{\text{Annual Site Water Usage (Litres)}}{\text{IT Equipment Energy (kWh)}}$$

Water-cooled chiller systems and evaporative cooling towers consume water through evaporation during the heat rejection process. Facilities in water-scarce regions, or those subject to water use regulations, should monitor WUE alongside PUE to gain a complete picture of their environmental resource consumption. Improving PUE through evaporative cooling may worsen WUE — the two metrics measure different aspects of efficiency and must be considered together.

1.2

Best-in-class PUE

Achieved by hyperscale operators using advanced cooling and power distribution

1.58

Industry Average PUE

Uptime Institute global average across reported data centres

Choosing the Right Cooling Approach

Selecting a cooling technology is one of the most consequential infrastructure decisions a data centre professional will make. An inappropriately specified cooling system — whether under-sized, over-engineered, or incompatible with planned workloads — creates operational problems that are costly and disruptive to correct after deployment. The evaluation should be methodical, grounded in actual workload data, and account for both current and projected requirements.

Current and Projected Rack Power Density

Start with actual measured rack power density across the facility. Identify any racks or zones planned to house GPU, AI, or HPC hardware. Project how density requirements will change over the next three to five years based on known hardware refresh or expansion plans.

Workload and Hardware Compatibility

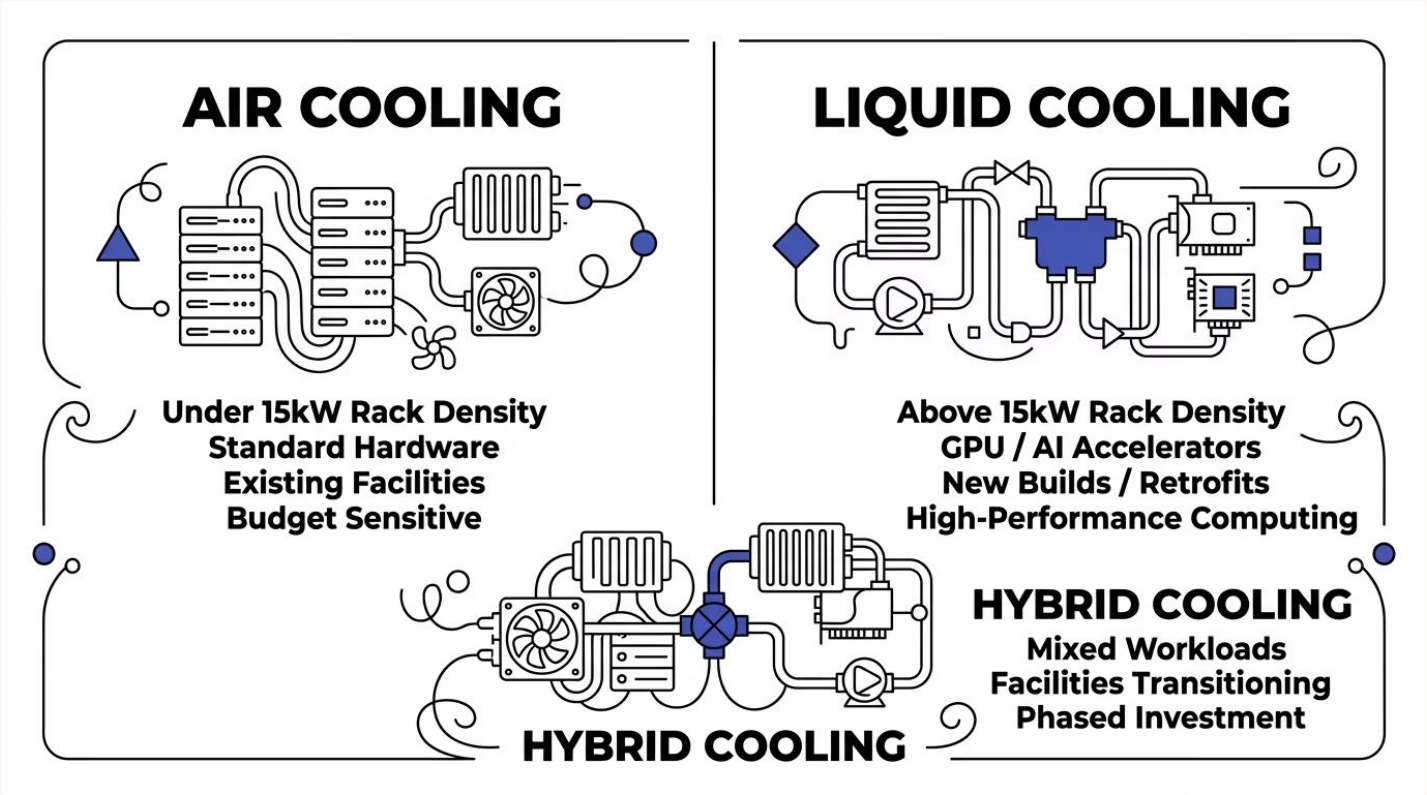
Determine whether planned servers support liquid cooling cold plates. Check with hardware vendors whether direct-to-chip or immersion cooling is validated for specific server models. Not all IT equipment is compatible with every liquid cooling implementation.

Existing Infrastructure Capacity

Assess whether the current cooling plant can support planned density increases with modifications, or whether the thermal load will exceed capacity under any configuration. Use DCIM data, cooling load calculations, and thermal modelling.

Energy, Water, and Cost

Model the total cost of ownership across each shortlisted cooling option, including capital expenditure, installation, ongoing energy and water consumption, and maintenance costs over a realistic operational period — typically five to ten years.



Organisations should resist the temptation to adopt liquid cooling as a default response to density concerns without first verifying that thermal loads genuinely exceed air cooling capability. For facilities where rack densities remain below 15 kW and workloads consist of standard enterprise computing, well-designed air cooling with proper containment will continue to provide reliable, cost-effective thermal management.

Cooling System Maintenance and Common Challenges

Cooling systems are mechanical and electromechanical infrastructure with many components subject to wear, contamination, and degradation over time. A structured maintenance programme is essential to sustain cooling performance, prevent unplanned failures, and extend the operational life of cooling plant and equipment. Deferred maintenance in cooling infrastructure rarely results in cost savings — it typically results in higher repair costs and increased risk of thermal incidents.

Common Operational Issues

→ Blocked Air Filters

Accumulated dust restricts airflow through CRAC and CRAH units, reducing cooling output and increasing fan energy. Filters should be inspected and replaced on a schedule based on the facility environment — typically every 3–6 months.

→ Cooling Unit Failures

Compressor failures in CRAC units, coil fouling in CRAH units, and control board faults can reduce cooling capacity. Redundant cooling unit configurations (N+1 or 2N) provide resilience during single-unit failures.

→ Coolant Leaks and Fluid Quality

Liquid cooling systems require regular inspection of connections, fittings, and flexible hoses. Coolant quality — pH, conductivity, biological contamination — should be tested periodically to prevent corrosion and biological growth within piping and heat exchangers.

→ Inadequate Cooling Capacity

As IT loads grow, cooling capacity that was originally sufficient may become marginal. Capacity headroom should be reviewed whenever significant hardware changes are planned, not only during annual facility assessments.

Preventive Maintenance Programme

- [ ] Inspect and replace CRAC/CRAH air filters on scheduled intervals (minimum every 6 months)
- [ ] Test cooling unit compressors, fans, and control systems monthly
- [ ] Check pump operation and flow rates in liquid cooling circuits quarterly
- [ ] Inspect coolant connections, manifolds, and flexible hoses for signs of leakage or degradation
- [ ] Test coolant water quality: pH, conductivity, inhibitor concentration, biological contamination
- [ ] Calibrate temperature and humidity sensors against reference instruments annually
- [ ] Review DCIM cooling trend data for gradual performance degradation between planned inspections
- [ ] Conduct full cooling system performance review following any significant IT load change
- [ ] Inspect cooling tower fill, drift eliminators, and water treatment systems at least quarterly
- [ ] Record all maintenance activities, faults, and corrective actions in a maintained log

Data Centre Cooling Assessment Checklist

The following checklist is designed to help data centre professionals conduct a structured review of their existing cooling infrastructure. It can be used as part of an annual facilities assessment, in advance of a hardware expansion project, or as a baseline audit when taking over responsibility for an existing facility. Each section identifies specific conditions that should be verified and documented.

SECTION 1

Cooling Capacity

- [ ] Is total installed cooling capacity (kW) documented and known?
- [ ] Does the current cooling system support the existing IT load with adequate headroom (minimum 20% recommended)?
- [ ] Can the cooling system support planned rack expansion within the next 12–24 months?
- [ ] Has a formal cooling capacity assessment been conducted within the last 12 months?

SECTION 2

Airflow and Thermal Conditions

- [ ] Are hot and cold aisles clearly defined and consistently maintained across all rack rows?
- [ ] Are containment systems (cold aisle or hot aisle) intact with no broken panels, missing doors, or unsealed cable penetrations?
- [ ] Are blanking panels fitted in all empty rack slots throughout the data hall?
- [ ] Are recurring temperature hotspots documented? Have root causes been investigated?
- [ ] Have floor tile positions been reviewed for alignment with actual rack power density?

SECTION 3

Liquid Cooling Readiness

- [ ] Have high-density or GPU servers been assessed for liquid cooling compatibility with the hardware vendor?
- [ ] Has the facility been assessed for CDU placement, piping routes, and secondary containment requirements?
- [ ] Are leak detection systems in place or planned for all liquid cooling zones?
- [ ] Are maintenance staff trained or being trained in liquid cooling system operation?

SECTION 4

Energy and Water Efficiency

- [ ] Is PUE monitored at least monthly using metered facility and IT energy data?
- [ ] Is WUE monitored for facilities using evaporative cooling or cooling towers?
- [ ] Has cooling energy consumption been benchmarked against comparable facilities or industry averages?
- [ ] Are opportunities for free cooling or economiser operation assessed and implemented where feasible?

SECTION 5

Maintenance and Monitoring

- [ ] Are CRAC/CRAH units inspected and maintained on a documented schedule?
- [ ] Are temperature and humidity sensors calibrated and verified annually?
- [ ] Are cooling unit alarms integrated into a building management system (BMS) or DCIM platform with appropriate alert escalation?
- [ ] Are maintenance records (inspections, faults, corrective actions) maintained and accessible?
- [ ] Has cooling system redundancy been tested through simulated unit failure within the last 12 months?

SECTION 6

Future Requirements

- [ ] Has a formal assessment of AI or high-performance computing workload requirements been conducted?
- [ ] Is there a documented cooling upgrade or expansion plan aligned to the IT hardware refresh roadmap?
- [ ] Have future rack power density assumptions been reviewed with the IT and applications teams?
- [ ] Is the facility's cooling infrastructure able to accommodate planned redundancy or availability tier upgrades?

GSDC Certification

GSDC PROFESSIONAL CERTIFICATION

Develop Your Knowledge of Data Centre Infrastructure

The **GSDC Certified Data Centre Professional** certification provides a structured and comprehensive foundation in data centre infrastructure, operations, and management. The programme covers power systems, cooling infrastructure, capacity planning, redundancy design, infrastructure management, and operational efficiency — equipping professionals with the knowledge required for data centre management and infrastructure decision-making roles.

Whether you are an experienced practitioner formalising your expertise or an infrastructure professional preparing for increased responsibility, the GSDC certification provides a recognised and practical qualification aligned to real data centre operational requirements.

Explore the GSDC Certified Data Centre Profession...

What the Certification Covers

- Data centre infrastructure and operations
- Power systems and UPS design
- Cooling systems and thermal management
- Capacity planning and redundancy
- Infrastructure management and efficiency
- Availability tiers and resilience design

www.gsdcouncil.org/certified-data-center-professional

GSDC
Global Skill Development Council

Globally Recognised Certification Body
190+ countries · 2,50,000+ Certified Professionals

IAS
ACCREDITED
Personnel Certification Body

Certified Data Centre Professional



ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org