Data Center Physical Infrastructure - Cooling

Why data centers overheat and the full toolkit for keeping them cool — architectures, air and water cooling technologies, water treatment, airflow management, rack cooling, monitoring, and surviving a power outage

A data center is, thermodynamically, a room full of machines that turn electricity into heat. Left unmanaged, that heat damages chipsets, drives condensation and corrosion, and eventually takes the whole facility down — so cooling is not an accessory to a data center, it is part of what keeps it alive. This module works from the ground up: why cooling matters and how it protects reliability, efficiency, and availability; how to size the load against the ASHRAE temperature and humidity window; the three architectural choices (how heat is removed, how air is circulated, and where the cooling units sit); the specific technologies that move heat out of the room — air-cooled, glycol, condenser-water, chilled-water, and liquid systems, plus the heat-rejection gear behind them; how water-based systems are kept clean with water treatment; the airflow management that separates a well-run room from a hot one; the strategies that cool ever-denser racks; the sensors, DCIM, and routine checks that keep it all honest; and finally, what happens to temperature when the power fails, and how to keep cooling alive through an outage.

LEARNING OBJECTIVES

- Recognize why a data center needs a cooling system and how it improves reliability, efficiency, and availability
- Determine the cooling load and heat load of a data center
- Recognize cooling architectures: heat removal methods, air circulation, and cooling-unit placement
- Describe cooling system technologies and identify direct expansion (DX) systems and heat-rejection systems
- Differentiate air-based, free cooling/economizer, and water/liquid-based cooling systems
- Describe a data center water treatment system and its two types
- Recognize the benefits of a better airflow management strategy and calculate basic heat transfer
- Identify reasons for poor airflow, approaches to lower server temperature, and airflow best practices
- Identify the strategies for improving rack cooling in a data center
- Recognize ASHRAE guidelines and the method to calculate and measure heat
- Identify sensor types, DCIM components, and the parameters of a routine check
- Identify factors causing temperature rise, predictive heat models, and strategies to manage cooling during a power outage

8.1 Recognizing the need for cooling

It is normal for stacked equipment to heat up, but with rising energy use, IT loads, and server density, an efficient cooling system becomes essential to hold an optimal temperature. Data centers are highly energy-intensive and prone to overheating: as the number of servers grows, so does the heat they emit, and a rise in temperature damages IT equipment and causes downtime. A cooling system removes that heat and balances the air temperature — protecting the equipment and, in turn, improving the data center's **reliability, efficiency, and availability**.

WHAT A COOLING SYSTEM DOES

- Maintains server temperatures at a safe operating level.
- Removes the excessive heat generated by densely stacked servers.
- Improves the data center's reliability, efficiency, and availability.
- Prevents downtime and maintains ideal working temperatures.

THE ASHRAE COMFORT WINDOW

- **Temperature: 64.4°F to 80.6°F** (18–27°C) for ideal IT equipment operation.
- **Dew point: 15.8°F to 59°F** (−9 to 15°C).
- **Relative humidity: below 60%** — noncondensing.
- These 2016 ASHRAE technical-committee figures assume the latest equipment; consider local regulations for data centers outside the US.

TABLE 8.1 How a well-implemented cooling system pays off

ATTRIBUTE	HOW COOLING DELIVERS IT
Reliability	Keeps thermal stress away and boosts the life of IT equipment through well-directed airflow
Efficiency	Avoids mixing hot and cold air, reducing pressure on the air-conditioning units and lowering cost and power consumption
Availability	Prevents outages from IT equipment that fails, malfunctions, or is damaged by thermal stress — the leading cause of data center outages

Cooling load is the total amount of heat in a data center that requires cooling to maintain the ideal temperature. Higher temperature means more thermal stress and a larger cooling load; lower thermal stress means a lower load. It depends broadly on four factors, so anything that changes the heat in the room changes the load.

THE COOLING LOAD DEPENDS ON FOUR THINGS

- **Heat generated by data center equipment** — IT gear, power distribution systems, lighting, UPS and its batteries, and other devices.
- **Heat from other elements, businesses, or people** inside the building — significant when the building is poorly insulated and heat transfers through walls.
- **Outside temperature** — enters through poor insulation or the opening and closing of doors.
- **Air ventilation load** — the ventilation equipment and sources that filter outside air.

COMMON PITFALL

Cooling load is not fixed. If 15% of the racks will be shut down for a migration, that reduces heat generation — one of the four load factors — so the cooling plan should account for it. Cooling can be dialed back for that period, saving cost.

8.2 Cooling architectures and systems

Data centers use three types of cooling architecture, and every design makes a choice in each category: the **heat removal technique**, the **type of air distribution (circulation) model**, and the **placement of the cooling units**. Air circulation is a key aspect of the whole system — how supply and return air move decides how effective the other choices are.

There are many **heat removal methods** for cooling mission-critical computing rooms. Some shift the components that handle the refrigeration cycle; others add independent loops of chilled water or cooling liquids to carry heat away. Nine methods appear across the deck.

TABLE 8.2 Heat removal methods

METHOD	HOW IT REMOVES HEAT
Chilled water system	Cool water (46.4°F–59°F) pumped from a chiller to CRAH cold-water coils that absorb room heat; heat passes to a condenser water stream and is rejected through a cooling tower
Pumped refrigerant for chilled water units	A pumped-refrigerant heat exchanger and chiller keep water cool; resolves availability issues and meets higher densities
Air-cooled split system	Air-cooled CRAC plus a condenser — the air-cooled CRAC direct expansion (DX) unit; most refrigeration-cycle parts sit in the CRAC; used in standard IT environments of any size
Glycol-cooling unit	Glycol-cooled CRAC plus a dry cooler in one enclosure; glycol absorbs more heat than air and carries it away via a heat exchanger
Water-cooled CRAH system	Cooling tower plus a water-cooled CRAH; a condenser water loop (instead of glycol) gathers heat and a cooling tower carries it outdoors
Air-cooled self-contained one-piece unit	Self-contained unit plus an air duct; the refrigeration cycle is in one shell; expels exhaust air near 120.2°F away from the IT room
Direct air evaporative cooling	Air duct plus a direct fresh-air evaporative cooler; fans and louvers draw filtered cold outside air straight into the data center (outside air must meet set parameters)
Indirect air evaporative cooling	Air duct plus an indirect evaporative cooler; uses naturally cold outdoor air in winter to reduce the data center air temperature and cut energy cost
Self-contained roof-top system (RTU)	Like the air-cooled self-contained system but roof-mounted, outdoors, and larger; not suitable for modern data centers

METHODS OF AIR CIRCULATION

- **Flooded** — air is supplied from and returned to the air distribution system; flooding a room with cold air leaves a mixture of warm and cool spots.
- **Targeted** — a raised floor plus devices (duct, perforated inlay, or cooling unit) guide supply and return airflow within 10 feet of where IT equipment consumes and dissipates heat.
- **Contained** — IT equipment airflow is fully enclosed for precise control, removing any mixing between the supply and return air streams.

TABLE 8.3 Placement of the cooling unit

PLACEMENT	CHARACTERISTICS
Room-based	CRAH units linked within the room monitor the whole heat load; one or more air conditioners; affected by ceiling height, room size and shape, floor type, rack layout, CRAH placement, and power distribution
Row-based	CRAH units linked with a single row, placed among the racks or fixed overhead; smaller, distinct air paths
Rack-based	CRAC units connected or positioned inside the racks; shortest, clearly defined air paths; full CRAC capacity usable; up to 50 kW per rack
Hybrid	A combination of room-based, rack-based, and row-based cooling

8.3 Cooling system technologies

Routine operation emits heat, and rising temperature and humidity can damage equipment through condensation and cut efficiency — which is why cooling technology is required. Which technology fits depends on the data center's design and architecture. Sizing starts from the **heat load**, which is not only the equipment deployed per area but also lighting, people, environmental temperatures, building architecture, and other factors.

ASHRAE STANDARDS AND COOLING BEST PRACTICES

- **Standards** — dry-bulb temperature 18°C–27°C (64°F–81°F); noncondensing humidity from −9°C to 15°C dew point (15.8°F–59°F DP) with 60% relative humidity.
- **Equipment** — plan no less than 1 ton of cooling per 400 sq ft of IT floor space (1 ton = 12,000 BTU/hr per 3,516 W).
- **Usage** — use a 40–60% nameplate electrical rating as the steady-state power withdrawal.
- **Power density** — the concentration of power in one rack; modern racks can exceed 20 kVA.
- **Economizer mode** — use an alternative cooling medium (e.g. outside air to cool refrigerant instead of running compressors) to cut electrical load and stay green.

TABLE 8.4 Cooling technologies and DX classification

TECHNOLOGY	HOW IT WORKS	DX?
Air-cooled	Refrigerant cycle rejects heat to air	Yes
Glycol-cooled	Glycol carries heat to a dry cooler	Yes
Condenser water-cooled	Condenser water loop carries heat to a tower	Yes
Chilled water	Water (not refrigerant) is chilled and piped to air-handling coils	No
Liquid-cooled	Liquid carries heat, often to the chip level	No

HEAT REJECTION SYSTEMS

- As important as the main cooling equipment; called **self-contained** when located inside and **split** when located outside.
- **Cooling towers and dry coolers** are the most common heat-rejection systems in the industry.
- Examples: cooling towers with water-cooled chillers, dry coolers, glycol-cooled chillers, and air-cooled chillers.
- Water- and glycol-cooled chillers usually stay inside; air-cooled chillers sit outside.

8.4 Differentiating the types of cooling

The IT infrastructure emits excessive heat, so cooling systems must maintain an optimal heat level and expel the hot air to the outdoors. By mode of operation, data center cooling systems fall into three families: **air-based cooling systems**, **free cooling or economizer modes**, and **water- and liquid-based cooling systems**.

An **air-based cooling system** uses a heat swapper (exchanger) to take in hot air and pass it outside. Its structure has three parts — the **indoor heat-swapping unit**, the **fluid haul**, and the **outdoor heat-swapping unit** — and it comes in four types.

TABLE 8.5 Air-based cooling system types

TYPE	DESCRIPTION
One-piece, air-cooled independent (packaged AC unit)	Combines an air duct and an air-cooled self-contained air conditioner; refrigeration elements in a single enclosure; expels heat as hot-air exhaust; mounts on a drop ceiling or roof
Two-piece, air-based (air-cooled CRAH / CRAC DX / split system)	A condenser makes it two-piece; components split between indoor and outdoor units; the refrigeration cycle sits in the CRAC, where the compressor compresses the vapor and a pump moves the refrigerant
Evaporative cooling with fresh air	An economizer coupled with an air duct; fans, louvers, and dampers draw outside air over a wet mesh, evaporating water and absorbing heat; works best in dry, low-humidity climates
Evaporative cooling with indirect air	An economizer for when outside temperature is lower than inside; a heat swapper isolates inside and outside air and is showered with water to cool incoming air; cost-effective, favors low humidity

AIR-BASED COOLING – ADVANTAGES	AIR-BASED COOLING – DISADVANTAGES
cheap and simple ▪ Low initial investment and operational costs ▪ Requires less maintenance than water-cooled systems ▪ Air-cooled units mainly just need air-filter changes	*sensitive to layout* ▪ Long piping reduces reliability ▪ Lower heat-transfer coefficient forces longer piping than fluid systems ▪ Poorly estimated distance or placement drastically cuts performance ▪ Evaporative cooling can make indoor air humid

A **water- and liquid-based cooling system** is a combination of cooling towers, cooling fluid, pumps, chillers, and condensers. It is highly cost-effective, needs little installation space, absorbs heat efficiently, and suits high-density data centers — but it requires large ductwork to move air, filter maintenance to keep out pollutants, and attention to server humidity requirements. Four types cover the ground.

TABLE 8.6 Water- and liquid-based cooling types

TYPE	HOW IT WORKS	NOTES
Liquid cooling method	A heat exchanger cools the air dispersed from the IT devices	High heat transfer, low operational cost, low thermal resistance; but high initial/maintenance cost and leak risk
Icy-cold / chilled water	A CRAH is connected to a chiller; a refrigeration cycle carries chilled water to the CRAH, whose coils absorb heat; warmed water is evaporated, condensed, and recycled	Cost-effective, high heat absorption; better with an economizer; high installation cost
Refrigerant-based chilled water	Refrigerant pumped by a compressor is tied to an air-based chiller; icy-cold water circulates through pipes to absorb heat; a heat swapper recools the refrigerant	A distinct chilled-water variant driven by refrigerant
Cooling tower	Ejects heat by spraying warm condensed water onto a sponge material	More efficient than dry coolers, high savings, long life; but clogging/scaling costs, weaker in humid air, high install labor

Free cooling or economizer modes run year-round where the weather is cool all year, and seasonally elsewhere when the outdoor temperature is favorable. They **reduce cost and save energy** without a large investment. The trade-offs differ by side: **air-side** economizers need large ductwork, filter maintenance, and humidity control, while **water-side** economizers have limited operating hours because of their temperature requirement.

8.5 Data center water treatment systems

Water-based cooling needs clean water — contaminated water degrades performance and damages the system. The contamination shows up as three challenges, and a **water treatment system** is the solution that removes microbes and solid contaminants and enhances water quality before it enters the cooling system.

TABLE 8.7 Challenges of contaminated water

CHALLENGE	WHAT IT DOES
Scaling	Minerals (calcium, magnesium, carbonate) from the water supply build up; impacts quality and efficiency, reduces heat-control capacity, and forces the system to draw more power and water. Controlled by chemical and nonchemical methods
Microbes	Grow in warm, moist conditions; block and clog filters, pumps, and other parts, reducing airflow. Removed with chemicals or by water treatment
Corrosion	Chemicals and scaling disintegrate metal layers, causing leaks, water leakage, and wastage. Galvanized metal parts and water treatment reduce it

THE SEDIMENTATION PROCESS, IN ORDER

Water treatment often begins with sedimentation: **(1)** chemicals are added to the tank to decontaminate the water; **(2)** the chemicals bind with finer particulate matter; **(3)** the heavy particles settle at the bottom of the tank; **(4)** the heavy particles are removed from the bottom; and **(5)** the water is disinfected through a chlorination process to remove bacteria and microbes.

Two major water treatment systems are used in data centers: the **pulse-power** system and the **cooling tower blowdown** system.

PULSE-POWER WATER TREATMENT

- The most popular and widely used system; uses **electromagnetic pulses** to kill germs and microbes and combat scaling.
- Uses **no pumps or extra devices** and is **noncorrosive** — it does not harm the cooling system body.
- Two components: a **controller** (electronic circuit generating electromagnetic waves of varying frequency) and a **reaction chamber** (stainless steel or PVC pipe wrapped in a solenoid coil with no direct coil-to-water contact).
- It does not filter minerals; the waves clump them so they don't deposit, and the clumps leave during blowdown.
- **Pros:** controls scaling, corrosion, and microbes; saves water and power; safe (no chemicals). **Cons:** poor with soft or distilled water; consumes power.

TABLE 8.8 Cooling tower blowdown – reverse osmosis (RO) process

STEP	ACTION
1	Remove dissolved solids from the water in the first phase
2	Create permeate water — water with zero dissolved solids
3	Send a small amount of permeate water back to the cooling tower
4	Send the remaining permeate water to the polishing RO filter
5	Combine the first-phase and polishing water for a second RO phase
6	Remove suspended solids and send the water on for more treatment

Blowdown water is the concentrated water in a cooling tower that carries dissolved and suspended solids; it must be replaced and discharged to the treatment system for efficiency. The RO-based blowdown system offers **lower maintenance cost, fewer chemicals and less corrosion, and stable water quality**, at the price of a **costly initial investment, more floor area, and higher energy use** from the RO process.

8.6 Cooling infrastructure and airflow management

Efficient cooling comes from reducing airflow waste, and poor airflow management is usually the reason a room runs hot. A well-defined airflow management strategy **reduces the amount of air that must be supplied —** lowering energy consumption and avoiding waste — and **improves the distribution of temperature across cabinets**.

The physical mechanism is heat transfer: IT equipment uses internal fans to move its heat load (air that carries heat from the internal circuits and processors) to the **hot aisle**, and the **air handling unit (AHU) or CRAC** removes that heat by passing the hot air through a cooling coil with its own fans.

CALCULATING HEAT TRANSFER FOR AIR AT SEA LEVEL

Heat transfer is `Q = 1.085 × ΔT × CFM`, where `Q` is the heat transferred (BTU/hr), `1.085` is the constant for the specific heat and density of air, `ΔT` is the temperature rise (°F), and `CFM` is airflow in cubic feet per minute. The same form holds for the AHU and for the IT equipment, so heat in must equal heat out: `Q(AHU) = Q(IT)`, which reduces to `ΔT(AHU) × CFM(AHU) = ΔT(IT) × CFM(IT)`. In reality, small extra heat loads (lighting, people, AHU motors) are added to `Q(AHU)`. The upshot: the AHU airflow needed is proportional to the IT-generated airflow times the ratio of the IT temperature difference to the cooling-equipment temperature difference.

BYPASS AIRFLOW	RECIRCULATION AIRFLOW
cold air that skips the load ▪ Conditioned air returns to the AHU without passing through IT equipment ▪ Happens when cooling units oversupply air, or air escapes through wrongly placed perforated tiles and cable openings ▪ The wasted air mixes with warm room air and returns to the AHU	*hot exhaust re-inhaled* ▪ Cold air supply to the equipment is inadequate ▪ Server fans start pulling in the warmer air circulating nearby ▪ Raises inlet temperatures and endangers the servers

APPROACHES TO LOWER SERVER TEMPERATURE

- **Cold air at server inlets** — place cold-air perforated tiles directly in front of the rack; keeps inlet temps within limits but has no effect on the proportion of bypass air, and servers still see recirculation.
- **Decrease the AHU supply-air temperature** — lower it until the hot-aisle temperature falls below the highest recommended server temperature, so recirculated air stays cool.
- **Add more cold air** — turn on more CRAH units; often ineffective, since it raises CFM(AHU), lowers ΔT, and supplies more air than the server fans need.
- **Containment** — arrange hot and cold aisles; removes both bypass and recirculation, forces CFM(AHU) to equal CFM(IT), and reduces server energy use.

BEST PRACTICES FOR AIRFLOW CONTROL

- **Create cold and hot aisles** — parallel rows with server inlet sides facing each other to prevent a well-mixed thermal environment.
- **Install blanking panels** in every open slot to stop internal bypass and recirculation inside cabinets.
- **Restrict perforated tiles to cold aisles** — tiles elsewhere increase bypass air and waste under-floor pressure.
- **Close cable-cutout openings with restrictors** — cutouts severely hurt cooling capacity.
- **Seal gaps** between walls, raised floors, and columns (found by simple visual inspection).
- **Use appropriate tile selection** — high-capacity, adjustable grates near hot spots to cool higher-density racks.
- **Manage perforated-tile placement** by calculated load — too few tiles cause recirculation, too many cause bypass.
- Doing all of this achieves a variance of **25% or less** between CFM(AHU) and CFM(IT).

> **COMMON PITFALL**
>
> Adding more cold air feels like the obvious fix, but turning on more CRAH units just raises CFM(AHU), lowers ΔT, and oversupplies air the server fans don't need. Containment — not more air — is what removes bypass and recirculation and brings CFM(AHU) into balance with CFM(IT).

8.7 Improving rack cooling

A data center houses servers and IT equipment in **racks** to conserve space. Racks follow a row-based layout, are positioned as stacks with numerous shelves, and hold equipment connected by cables — an arrangement that increases capacity through dense servers. But density brings heat: higher-density equipment releases volumes of heat, and rising temperature and relative humidity cause irreparable damage, hamper productivity and availability, and make servers less efficient.

> **DENSITY IS CLIMBING**
>
> - Per Colocation America, a rack consumed about **6 kW in 2008**, rising to almost **16.5 kW in 2020**.
> - To prevent hardware damage, a data center needs a planned rack-cooling strategy — and continuous improvement of it.
> - The key is selecting the strategy most efficient for the specific data center and consistently reassessing the choices.

TABLE 8.9 Rack cooling strategies

STRATEGY	WHAT IT DOES
Cable management	Neatly lays cables on ladders and organizes the rest with clips/holders (removing dangling cables) so bulky cable webs don't hinder cold-air circulation; also helps spot damaged cables
Blanking panels	Fill empty rack slots to block hot-aisle air from recirculating into the cold aisle and direct cold air through the servers
Thermal insulated duct	Transports hot or cold air and expels IT hot air to the HVAC or CRAC plenum; compatible with hot/cold-aisle configurations and fire rules
Active heat removal	Ceiling- or roof-mounted ventilation fans pull hot air up and out and reduce negative pressure near racks; far cheaper than air conditioners
Passive heat removal	Natural ventilation through two vents (top near the ceiling, bottom near the floor) — hot air rises out the top as cool air enters the bottom; best paired with active removal
In-rack cooling	CRAC units mounted to or within racks give short, precise air paths and up to 50 kW per rack, using less CRAC-fan power
Closely-coupled cooling (in-row coolers)	Air conditioning placed within or among rack rows; easy to install; a backup that cools specific hot spots and shortens the IT-to-cold-air distance
Chimney containment	A lightweight metal chimney on the top or exhaust side ducts heat to the return plenum; may include fans; saves energy but can build pressure if unmanaged
Liquid-based chip-level cooling	A chilled-liquid coil on the motherboard heatsink absorbs heat at the chip; difficult to install, so used only by extremely high-density data centers

8.8 Measuring, monitoring, and routine checks

Data centers consume lots of electricity and produce volumes of heat, and both heat and humidity can seriously damage the workforce and servers: **high heat** burns the IT equipment's chipsets, **high humidity** causes condensation (and corrosion), and **low humidity** causes electrostatic discharge that can damage equipment. Avoiding these means identifying the cooling requirement and equipping the room with a well-planned cooling and humidity-control system.

TABLE 8.10 ASHRAE equipment classes

CLASS	TEMPERATURE	MAX RH
A1	15°C to 32°C	80%
A2	10°C to 35°C	80%
A3	5°C to 40°C	85%
A4	5°C to 45°C	90%
B	5°C to 35°C	80%
C	5°C to 40°C	80%

Most data center equipment falls in classes **A1 to A4**, for which the recommended operating range is **18°C to 27°C**, a dew point of **−9°C to 15°C**, and **relative humidity below 60%**. The recommended range is deliberately wide for better performance and reliability with the latest equipment; older equipment can make the optimal setpoint harder to pin down.

Heat calculation identifies and regulates thermal stress and helps size the cooling system. Heat comes from IT equipment, the workforce, electrical devices, and power, and the total is the sum of five individual calculations.

TABLE 8.11 Heat-source calculations

SOURCE	CALCULATION
IT equipment load	IT equipment load power
UPS and its battery	(0.04 × power rating) + (0.05 × overall IT load)
Power distribution	(0.01 × power rating) + (0.02 × overall IT load power)
Electrical lighting	1.33 × floor area (sq ft), or 14.32 × floor area (sq m)
People / workforce	Maximum number of workforce × 100

TABLE 8.12 Heat measurement – units and conversions

FROM	FORMULA	TO
BTU	BTU/hour × 0.293	Watts
Tons	Tons × 3,530	Watts
Watts	Watts × 3.41	BTU
Watts	Watts × 0.000283	Tons

SENSORS USED IN DATA CENTERS

- Temperature monitoring, humidity monitoring, and water-flow detection sensors.
- Power-supply monitoring, water-pressure, and air-pressure sensors.
- Sensors give real-time monitoring of environmental attributes, precise readings, and control over the cooling system's airflow — because too-high or too-low temperature and humidity can break the system.

DCIM – COMPONENTS	DCIM – BENEFITS
Data Center Infrastructure Management ▪ Proper planning and adept design ▪ Identifying and analyzing operations ▪ Monitoring the whole data center's infrastructure and equipment ▪ Predictive analysis of future growth trends	*a guided, real-time view* ▪ Better understanding of the infrastructure ▪ Quick decision-making and resource planning ▪ Enhanced uptime ▪ Cost and energy efficiency; decreased risk

PARAMETERS OF A ROUTINE CHECK

- **Temperature** — apply ASHRAE recommendations; install a good sensor that alarms in emergencies.
- **Humidity** — install a humidity sensor to catch very high or very low levels.
- **Water** — check for leakage or standing water to avoid corrosion and condensation.
- **Cooling system conditions** — never neglect a flaw; call a dedicated engineer.
- **Fire-extinguishing equipment** — examine regularly even though it is rarely used.
- **Electrical system** — schedule regular checks for voltage fluctuations and short circuits.
- **Physical security** — 24/7 surveillance of the premises protecting sensitive data.

8.9 Factors affecting transient temperature rise

Many conditions push a data center's temperature up, and a good design anticipates them. The primary factors span the building itself, the load, the equipment sizing, and — critically — what happens when the power fails.

FACTORS THAT CAUSE TEMPERATURE TO RISE

- The **architecture and layout** of the data center, plus **unplanned IT load**, which can cause cooling loss.
- A **power-density surge** or an **increase in supply temperature**.
- **Improper cooling-equipment size** or poor **rack air containment** practices.
- A **sudden power failure**, which spikes temperature without UPS-backed cooling.
- **Little focus on cooling-system design during power failures** — organizations plan for power redundancy but overlook cooling optimization during an outage.

FOUR CONDITIONS WHERE THE HEATING RATE VARIES

- A **heat surge** situation.
- **Using more IT equipment than planned.**
- A **power failure**.
- **No backup power system.**
- In a cooling-failure situation, the heating rate is balanced by heat absorbed by the **data center air and the building material**.

FOUR WAYS TO HANDLE COOLING

- Position critical cooling equipment on the **backup power supply**.
- Select equipment with **short restart times**.
- Keep **sufficient cooling capacity in reserve**.
- **Commission thermal storage.**

WELL-MIXED AIR PROTOTYPE MODEL	COMPUTATIONAL FLUID DYNAMICS (CFD)
fast and simple	*detailed and mathematical*
▪ Treats the data center as a single entity — CRAH, CRAC, IT loads, and seamless, well-mixed air ▪ Assumes one uniform air temperature at any time ▪ Fast and straightforward; good for predicting an overall temperature surge over time	▪ Uses math to predict the flow of liquids and gases, including data center airflow ▪ Handles detailed, non-uniform conditions ▪ Often used when a data center loses cooling power and is not restored by backup power

8.10 Managing cooling during power outages

Managing cooling during a power outage is a critical task. An emergency power source helps handle the heating rate, but the right strategies keep the room within limits when power suddenly drops. Four strategies apply, and they work best together.

TABLE 8.13 Four outage cooling strategies

STRATEGY	HOW IT WORKS
Maintain sufficient (excess) cooling capacity	Normally capacity equals the load, but during an outage there is excess heat on top of the IT load, so capacity = load cannot restore the room. Oversize the plant (a 1 MW-rated chilled-water plant may carry only 100 kW of load) based on redundancy and efficiency needs so the room recovers quickly
Link backup power to cooling equipment	Between the outage and generator start, CRAH fans and chilled-water pumps are down and temperature spikes. Put CRAH/CRAC fans (and chilled-water pumps) on a UPS to keep airflow, limit hot-air recirculation, and let thermal mass absorb heat until the generator starts
Select quick-restart equipment	Advances in chiller technology cut restoration time to about 5 minutes (down from up to 15). Quick chillers keep a low-density data center within tolerable limits and reduce the risk of crossing into unacceptable temperatures
Use thermal storage	Supplementary chilled-water storage — with pumps and CRAH fans on the UPS — bridges cooling until the chiller resumes. Low-pressure plastic storage containers cost less than putting chillers on a UPS and are recommended for high-density data centers (mind space and load-bearing limits)

The common thread is preparation before the lights go out: reserve capacity so there is headroom to catch up, keep the air moving and the pumps running on UPS power through the generator gap, choose chillers that come back fast, and bank cold in thermal storage so cooling can be delivered until full power returns.

Module summary

- Cooling protects reliability (less thermal stress, longer equipment life), efficiency (no hot/cold air mixing, lower power and cost), and availability (fewer thermal outages). The ASHRAE window is 64.4°F–80.6°F, dew point 15.8°F–59°F, and relative humidity below 60%.
- The cooling load is the total heat needing removal; it depends on data center equipment heat, heat from other elements/people, outside temperature, and the air ventilation load — and it changes when equipment is added or shut down.
- Every cooling architecture chooses a heat-removal method (nine exist), an air-circulation model (flooded, targeted, contained), and a unit placement (room, row, rack, or hybrid — rack gives the shortest paths and up to 50 kW).
- Air-cooled, glycol-cooled, and condenser water-cooled technologies are direct expansion (DX) systems; chilled water uses water instead of refrigerant; liquid cooling goes to the chip. Cooling towers and dry coolers are the common heat-rejection systems.
- Cooling systems split into air-based (four types), free cooling/economizer modes, and water/liquid-based (four types); air-based is cheap but layout-sensitive, while water/liquid is efficient and dense but needs ductwork and humidity control.
- Water-based cooling faces scaling, microbes, and corrosion; water treatment fixes them with pulse-power systems (electromagnetic pulses, no chemicals) or cooling tower blowdown systems (a six-step reverse-osmosis process).

- Heat transfers as $Q = 1.085 \times \Delta T \times CFM$, balancing when Q(AHU) = Q(IT); bypass and recirculation airflow waste cooling, and containment plus best practices drive the CFM(AHU)/CFM(IT) variance to 25% or less.
- Rack density rose from about 6 kW (2008) to 16.5 kW (2020); nine rack-cooling strategies — from cable management and blanking panels to in-rack, chimney containment, and liquid chip-level cooling — keep dense racks safe.
- Monitoring relies on ASHRAE classes (most equipment A1–A4), five-source heat calculations, unit conversions (watt is standard), six sensor types, DCIM tools, and seven routine-check parameters.
- Transient temperature rise comes from architecture, unplanned load, supply-temperature increases, poor sizing/containment, and power failures; well-mixed air and CFD models predict the heating rate, and four strategies — excess capacity, backup power on cooling, quick-restart equipment, and thermal storage — manage cooling through an outage.

Review questions

1. **How does a well-implemented cooling system improve reliability, efficiency, and availability?**

   Answer: Reliability: it keeps thermal stress away and extends the life of IT equipment. Efficiency: it avoids mixing hot and cold air, reducing pressure on the AC units and lowering cost and power. Availability: it prevents outages caused by equipment that fails or is damaged by thermal stress — the leading cause of data center outages.

2. **What is the difference between a CRAC and a CRAH?**

   Answer: A CRAC (computer room air conditioner) contains the refrigeration cycle and compressor and is used in DX/split systems. A CRAH (computer room air handler) has no compressor and instead uses chilled-water coils fed by a chiller to absorb heat.

3. **Trace the chilled-water cooling loop from chiller to outdoors.**

   Answer: A chiller cools water to 46.4°F–59°F; the water is pumped to CRAH units; the CRAH cold-water coils absorb heat from the room air; the warmed water returns and passes its heat to a condenser water stream; and a cooling tower rejects that heat to the outdoor environment.

4. **Name the three challenges of contaminated water and what a water treatment system does about them.**

   Answer: Scaling (mineral buildup that cuts efficiency and heat control), microbes (block and clog filters and pumps), and corrosion (leaks and water wastage). A water treatment system removes microbes and solid contaminants and enhances water quality — using pulse-power (electromagnetic pulses) or cooling tower blowdown (reverse osmosis).

5. **What are bypass and recirculation airflow, and what actually fixes them?**

   Answer: Bypass airflow is conditioned air that returns to the AHU without passing through IT equipment; recirculation airflow is server fans pulling in warm nearby air because cold supply is inadequate. Containment (hot/cold aisles) plus blanking panels, sealed gaps, and correct tile placement removes both and balances CFM(AHU) with CFM(IT). Simply adding more cold air does not fix bypass.

6. **Write the heat-transfer equation and explain the balance that containment aims for.**

   Answer: $Q = 1.085 \times \Delta T \times CFM$. Heat in must equal heat out, so Q(AHU) = Q(IT), which reduces to $\Delta T(AHU) \times CFM(AHU) = \Delta T(IT) \times CFM(IT)$. Containment forces CFM(AHU) to equal CFM(IT); following the best practices holds the variance to 25% or less.

7. **Why does putting CRAH fans and chilled-water pumps on a UPS matter during an outage?**

   Answer: Between the outage and the generator starting, the CRAH fans and chilled-water pumps would otherwise stop and the room temperature would spike. Placing them on a UPS keeps airflow going, limits recirculation of hot exhaust into inlets, and lets the thermal mass of walls, plenums, and chilled water absorb heat until the generator takes over.

8. **During an outage, why isn't cooling capacity equal to the heat load good enough?**

   Answer: During an outage there is excess heat on top of the ongoing IT load, so a plant sized exactly to the load cannot return the room to its original state. Excess/oversized capacity — plus quick-restart chillers and thermal storage — provides the headroom to cool the room back down quickly.