Sophia Ize-Iyamu

THERMOFLUIDS

Data Center Fundamentals

Self-directed engineering study · Sole author, 2026

Cooling & thermal analysis | Self-directed study

A self-directed study connecting heat-transfer fundamentals to how a modern data center is cooled, laid out, powered, and kept efficient.

A data center is a heat-management problem first: every watt a server draws becomes heat that must be removed before the hardware throttles or fails. I built this study to connect the heat-transfer and thermofluids methods I use to how real data centers are designed, the airflow, the rack and power layout, the IT hardware that drives the load, and the efficiency metrics that decide what the cooling itself costs. It is a fundamentals study, not field experience, written to bridge a mechanical-engineering background into data center infrastructure.

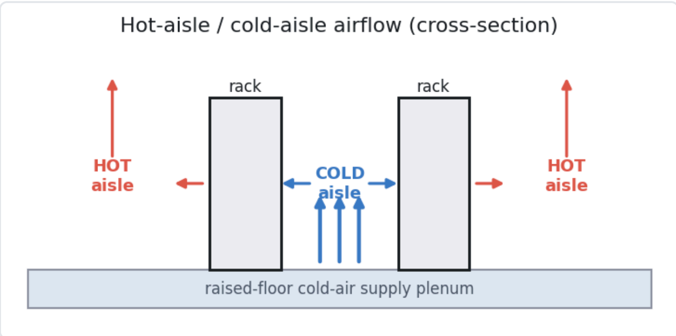

Hot-aisle / cold-aisle airflow: cold supply into rack fronts, hot exhaust captured and re-cooled.

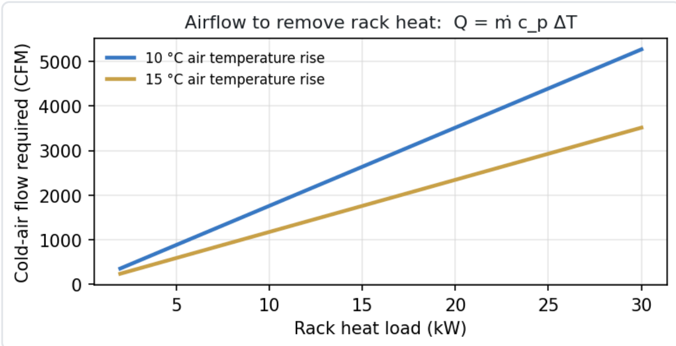

Cold-air flow needed to remove rack heat, from the Q = m-dot c_p ΔT energy balance.

Cooling: the core constraint

A loaded rack dissipates 5 to 30 kW, all of it heat. Racks sit front-to-front and back-to-back so supply and exhaust never mix: cold air rises from a raised-floor plenum into the cold aisle, through the servers, and out as hot exhaust that CRAC/CRAH units re-cool, with aisle containment the biggest lever on efficiency. Sizing follows the same energy balance as my heat-transfer labs, $Q = \dot{m} c_p \Delta T$: at a 10 to 15 °C rise, a 10 kW rack needs roughly 1,000 to 1,500 CFM.

Rack, power, and hardware

Equipment mounts in standard 42U racks (1U = 1.75 in); floor layout sequences rows, aisles, and service clearance. Power runs from utility to generator, UPS, and floor and rack PDUs with N+1 or 2N redundancy, so one failure never drops the load. The load itself is compute servers (CPU and increasingly GPU, the densest heat sources), network switches, and storage, and power and cooling capacity, not floor space, are the usual binding constraints.

Efficiency: PUE and liquid cooling

Power Usage Effectiveness (PUE = total facility energy ÷ IT energy) captures cooling and power overhead: a typical site sits near 1.5, an efficient one near 1.1. As GPU racks outrun what air can carry, direct-to-chip and immersion liquid cooling push that overhead down, the sustainability and cost story behind modern data center design.

SELECTED REFERENCES

- ASHRAE Technical Committee 9.9, Thermal Guidelines for Data Processing Environments, 5th ed. (hot/cold aisle, temperature envelopes).
- The Green Grid, PUE: A Comprehensive Examination of the Metric (White Paper #49).
- Uptime Institute, Tier Standard: Topology (N+1 and 2N redundancy).
- T. Incropera et al., Fundamentals of Heat and Mass Transfer (convection and the energy balance used here).