DATA CENTER

Edition

ISSUE 01
OCTOBER 2026



The world behind computing.

Inside the issue.

Stories, field photographs and visual explanations.

People & systems 3

The journey of heat 18

Power & performance 42

The water account 56

Land & networks 69

Testing the building 85

Planning for the load 102

Sources, image credits and editorial notes close the issue.

A hand in the frame.

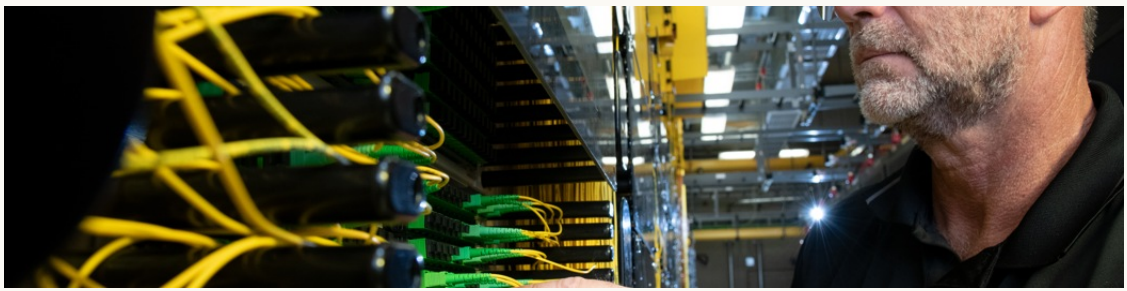
On our cover, a technician at a rural cooperative’s data center in Tennessee.

The hand on our cover belongs to Eddie Blankenship, a central office technician at North Central Telephone Cooperative, or NCTC. In September 2018 a USDA photographer found him fitting a fibre jumper in the cooperative’s data center in Lafayette, Tennessee. [1]

The building has a wide job. According to USDA’s description, it hosts the soft switch behind NCTC’s telephone service and supplies cloud computing, off-site backup, hosted servers and disaster recovery. A

network operations center on the same site watches the cooperative's fibre transport, IP network and television channels.

The cooperative has a local history too. NCTC has borrowed from USDA's Rural Utilities Service since 1951 to serve subscribers in Tennessee and Kentucky. After tornadoes destroyed much of the area in February 2008, it built a facility



Eddie Blankenship at NCTC, Tennessee, 2018. USDA / Lance Cheung. Public domain.

that the description calls EF-5 tornado resistant.

Data centers are usually described through their largest numbers. This issue keeps returning to smaller ones: the first phase of a campus, a few degrees of water temperature, the heat still left for air to carry, and the people who connect, test and run the equipment. These are questions that cross borders.



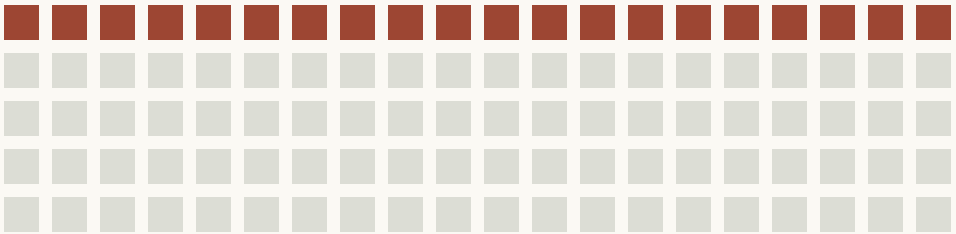
Rear of a rack at NERSC, Oakland, 2011. Derrick Coetzee / CC0.

One campus. Different numbers.

A worked example separates the scale of a plan from the part ready for service.

100 MW

Whole-campus IT design



20 MW

First phase in service

20%

Share of the design

Each mark is 1 MW of design capacity. An illustrative campus, not a reported facility or a floor plan.

One campus. Different numbers.

Imagine a campus designed for 100 MW of IT capacity. Its first 20 MW phase has completed testing and entered service. Both numbers describe the same development, but they answer different questions. The first is the scale of the plan. The second is the part ready for operation.

This is an illustrative scenario, not a reported facility. It deliberately leaves out a third number: the electricity being drawn. Twenty megawatts of operational capacity does not mean twenty megawatts are being consumed. Racks may be empty,

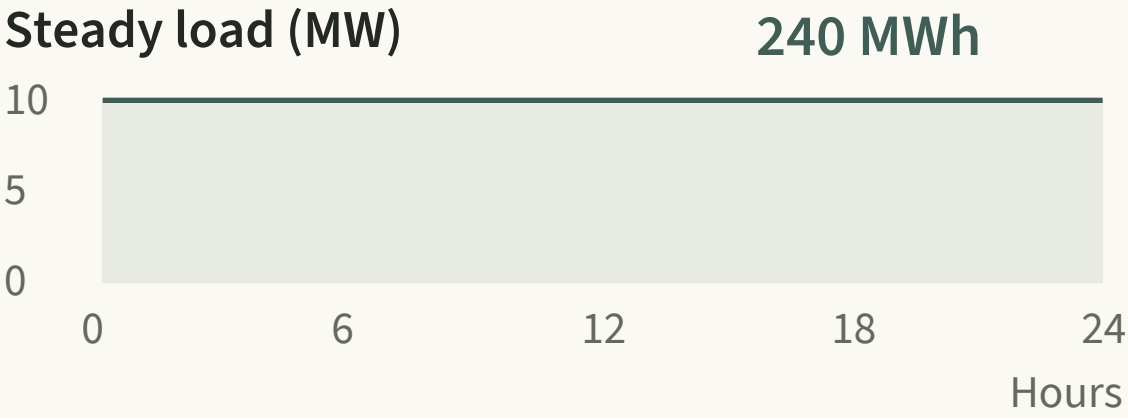
equipment may be arriving, or workloads may change through the day.

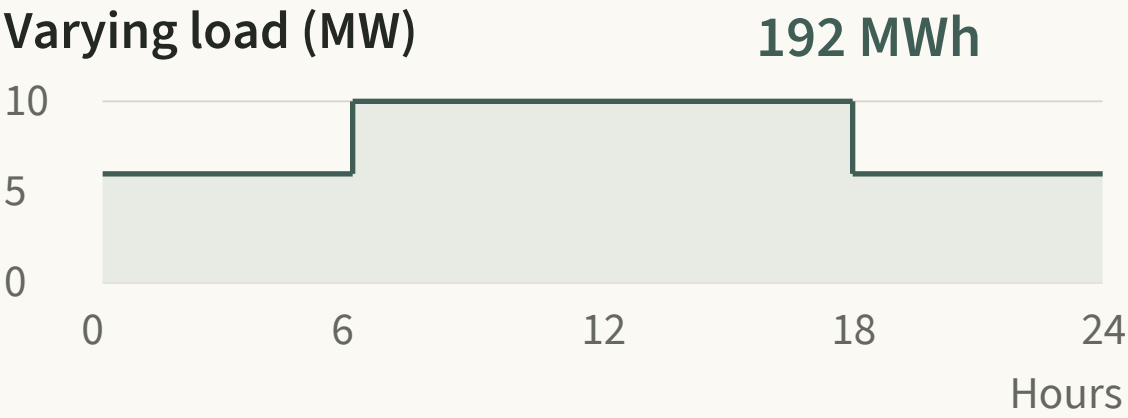
A customer preparing a move needs the particular room, its available power and cooling, and the date those services can be used. The campus total cannot supply those answers. A plan becomes useful when it is tied to the exact service being offered.

The distinction travels well. Wherever a project is built, keep planned capacity, capacity in service, commercial commitments and measured demand separate. The calendar belongs beside the number.

The rate and the total.

Power is a rate. Energy accumulates over time.





Both profiles peak at 10 MW. The shaded area adds up to the energy used over 24 hours.

The rate and the total.

Power is a rate; energy is the total it adds up to over time. A steady 10 MW load uses 240 MWh in a day. A site that idles at 6 MW overnight shares the same 10 MW peak but uses 192 MWh. A capacity figure gives the size of the job; only a record over time gives the energy.

Power held over time becomes energy.

10 MW	24 hours

240 MWh

A different day outside.

The computers can keep doing the same work while the cooling plant has to adapt.



CAMPUS, Frauenfeld, Switzerland, May 2022. Kecko / CC BY 2.0.

A different day outside.

The computers can spend a day doing the same work while their building has a very different day. A cool night gives way to a hot afternoon; the heat keeps coming from the racks all the while. Outside, the air available to carry it away has changed.

Cooler outdoor air lets a dry cooler shed heat more easily: for the same water temperature, the difference between water and air is larger. When the air is cool enough, an appropriately designed system can reject heat without running refrigeration compressors. Fans and pumps still have work to do.

Evaporation gives a cooling tower another route. Some water leaves as vapour, taking heat with it and cooling the remaining water. [2] The air's humidity now matters alongside its temperature. More humid air at the same temperature leaves less scope for evaporative cooling.

A building's cooling equipment may change modes as those conditions change. The annual result emerges from their succession: many mild hours, difficult hot hours, changing loads. The electricity bill accumulates all those afternoons and nights.

Behind the rack.



Rear of a rack at NERSC, Oakland, 2011. Derrick Coetzee / CC0.

Frontier's virtual rehearsal.

At Oak Ridge, a model offered a place to explore a consequential change.



Frontier at Oak Ridge, June 2022. OLCF / CC BY 2.0.

Frontier’s virtual rehearsal.

A working supercomputer is an expensive place to discover that a cooling change went too far. Warmer incoming water can create opportunities to reduce cooling energy, but the computers still need their heat carried away. The attraction is obvious; so is the reason to investigate carefully.

In September 2024, Oak Ridge National Laboratory described a digital twin of Frontier that its team was still developing. ORNL said the twin was already yielding insights into how workloads affect cooling and into power-conversion losses, and that

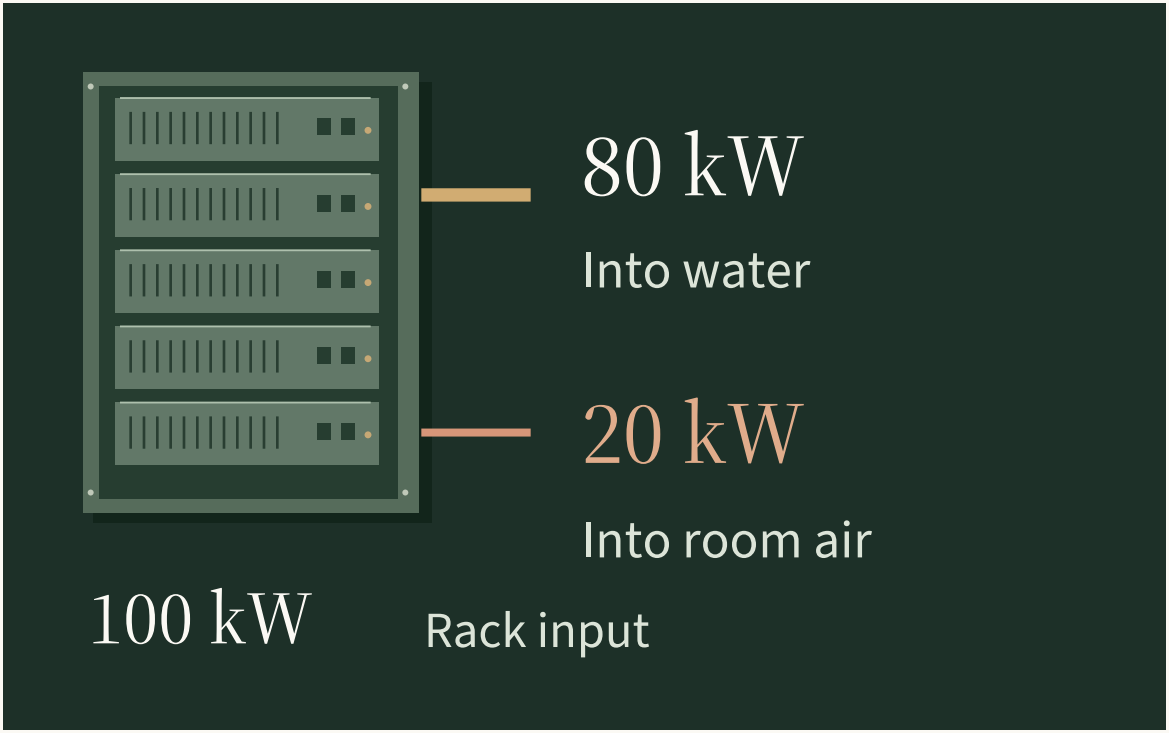
it would let operators try scenarios such as warmer incoming water before changing the real machine. [3]

Feiyi Wang explained the ambition in ORNL's published account: "if we have enough fidelity modeled for the power and cooling aspects of the system, we can experiment." [3]

That freedom rests on fidelity. A promising setting on screen has to survive contact with the real machine. The virtual rehearsal lets the team explore the consequences first.

The afterlife of heat.

A calculation can be finished while its heat is still leaving the room.

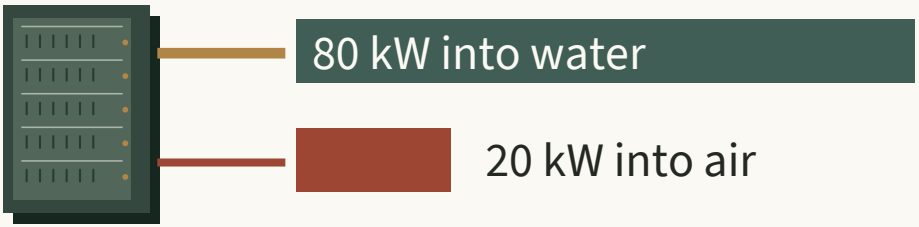


Assumed 80/20 heat split. All 100 kW becomes heat; supporting equipment adds its own heat separately.

The afterlife of heat.

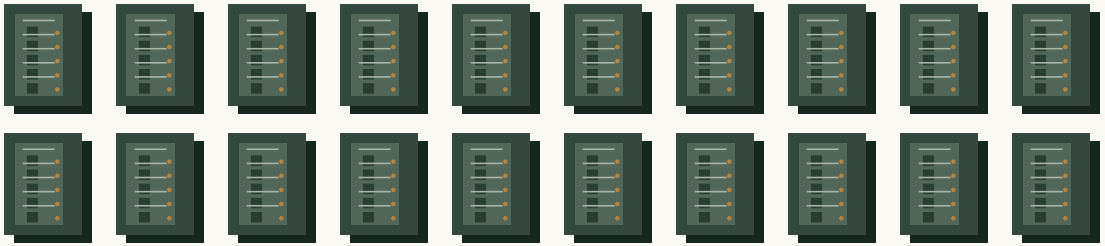
On the screen, the job is over. Beyond the screen, its energy is still working through water, pipes and air. Start with a steady 100 kW rack. All of that electrical input ends up as heat. [4] In this example an assumed 80 kW enters water and 20 kW enters room air; real splits depend on the equipment. External supporting equipment adds heat separately.

One rack. Two heat paths.



Twenty kilowatts stay in the room.

Water takes the heaviest load. Air keeps a substantial share.



20 racks × 100 kW = 2 MW



1.6 MW

Into water

0.4 MW

Into room air

Assumed 80/20 heat split; supporting heat is additional.

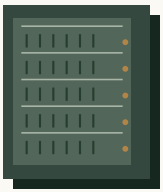
Twenty kilowatts stay in the room.

A pipe leaving a rack makes liquid cooling easy to picture. The heat left behind is less conspicuous. It enters the room through components beyond the cold plates, leaving the air-cooling system with work that can survive the conversion to liquid.

For our 100 kW rack, that remaining share is 20 kW. Twenty racks make it 400 kW, before heat from external support equipment joins the account. Water carries away the other 1.6 MW. Both routes lead back to the same electrical input.

The smaller share has become a substantial job. Bringing water to the processors has changed the distribution of heat without ending the old air-cooling route. Beyond the cold plates, the room still has to keep collecting what the computers release. Most of the heat is now in the water, ready for the next leg of the journey. [4,5]

The room still has a cooling duty.



× 20

400 kW

into room air

20 racks × 20 kW of remaining air-side heat.

Why the return pipe is warmer.

The rise in temperature tells us how much added energy each kilogram carries away.

80 kW carried in both cases

10 K	115 L/min
Temperature rise	About 1.91 kg/s
5 K	230 L/min
Temperature rise	About 3.83 kg/s

Water cp 4.18 kJ/(kg K); density 1 kg/L. Values rounded.

Why the return pipe is warmer.

Water leaving a rack warmer than it arrived is doing its job. The temperature rise records the energy it has collected. Under the same assumed water properties, a kilogram gaining ten degrees carries twice as much added energy as one gaining five.

The return pipe tells that story in degrees. A smaller rise means each kilogram has taken less energy on its journey through the equipment. To collect the same heat, more kilograms have to make the trip each second. Halving the rise doubles the traffic.

A larger permitted rise could reduce the flow, but components have temperature and flow limits. The moving water also encounters resistance in cold plates, hoses, filters and valves. The pump must overcome it, and the flow must reach the individual branches. A small heat balance has opened into a much larger network. [4,6,7]

The same heat, more water.



Same heat and assumed water properties.

The same heat. More traffic.

A smaller load per kilogram needs more kilograms passing each second.

10 K rise

115 L/min

Two vehicles, two parcels each.



5 K rise

230 L/min

Four vehicles, one parcel each.



Same 80 kW in both rows. Each vehicle represents equal water mass per interval; each parcel represents added energy.

The same heat. More traffic.

The upper pair of trucks carries four parcels. The lower fleet carries four too, spread across twice as many vehicles. Each truck stands for the same water mass passing during the same interval; each parcel represents added energy.

The smaller temperature rise has spread the job across more moving water. Both rows deliver the same heat in the same interval. The illustration holds the vehicles and parcels identical so the change is in the number of journeys, rather than the size of the load being removed. [4,6]

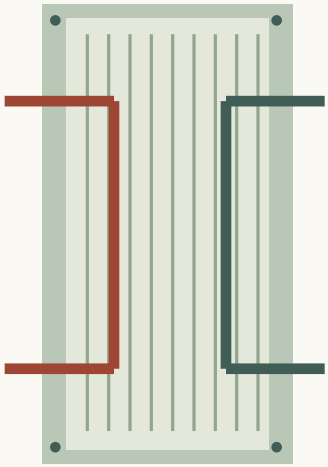
Where the water turns back.

Energy crosses the exchanger. The two liquids keep their separate routes.

Heat crosses. Water stays separate.

40 °C
From racks

30 °C
To racks



35 °C
To building

25 °C
From building

5 K difference at either end

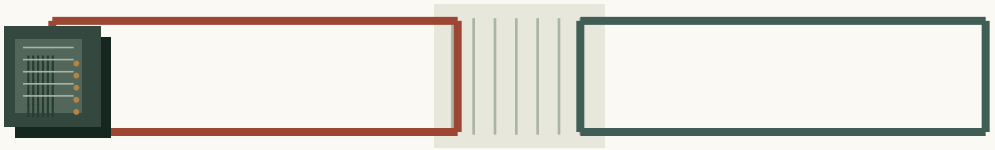
Counterflow example: 80 kW and about 115 L/min in each water circuit. Heat crosses the plates; the liquids do not mix.

Where the water turns back.

The rack’s warm water may end this leg of its journey at a coolant distribution unit, or CDU. In a liquid-to-liquid arrangement, heat passes through an exchanger wall into the building’s water circuit. Each liquid then returns around its own loop.

In this counterflow example, computer-side water enters at 40 °C and leaves at 30 °C. Building water enters from

Heat moves. The water returns.



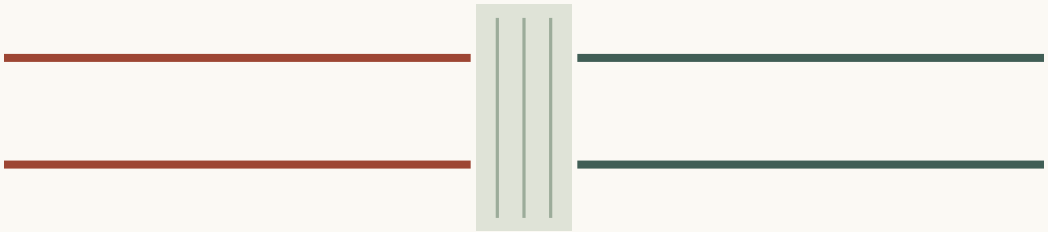
Computer loop

Building loop

the opposite end at 25 °C and leaves at 35 °C. The difference at each end is 5 K, providing the driving force for transfer.

The two circuits can have different fluids, pressures and material requirements. Energy continues through the building, while the water beside the computers turns back for another load. [4,5]

Heat crosses. The liquids stay apart.



Computer side

Building side

The lake joins the circuit.

At CSCS, the cooling route reaches far beyond the computer room.



Illustration by Data Center Edition. Imagined scene; details and attribution in the closing notes.

The lake joins the circuit.

The Swiss National Supercomputing Centre’s 2015 factsheet begins its outer cooling route 45 metres below the surface of Lake Lugano. Lake water travels 2.8 kilometres and climbs 30 metres to the centre. Those dimensions take the story out into the surrounding infrastructure.

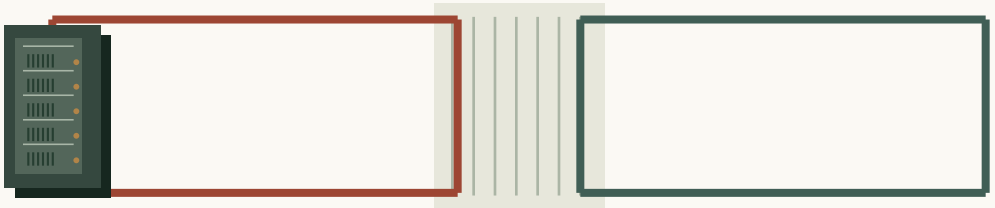
A journey beyond the computer room.

45 m	2.8 km	30 m
Intake depth	Pipe distance	Climb to site

CSCS factsheet, 2015. Infrastructure dimensions.

At the heat exchangers, that water collects energy from a separate internal circuit. The lake water and internal coolant stay apart. These are historical infrastructure dimensions, rather than current operating temperatures for a particular computer. This AI-assisted illustration imagines the setting; it shows neither the CSCS building nor its actual pipe route. [8]

Heat moves. The water returns.



Centre circuit

Lake-water circuit

Warm water can do the cooling.

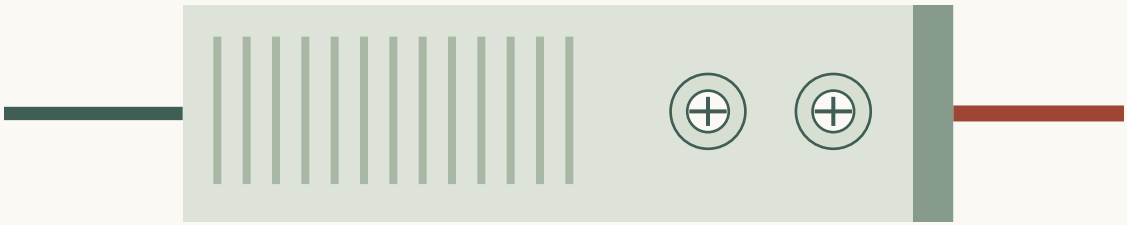
Frontier brings a familiar assumption about cooling into question.

1.0 MW

Cooling duty

0.2 MW

Compressor input



Chiller

1.2 MW

Condenser heat outdoors

Cooling COP 5. Worked example; other auxiliaries extra.

Warm water can do the cooling.

Frontier's cooling story includes a striking absence. ORNL's September 2024 account describes fan-free liquid cooling inside the machine. Its warmer circulating water can reject heat through evaporative towers without refrigerating that water. The machine still belongs to a system with pumps, evaporation and a route outdoors.

Water can feel warm to a person and still accept heat from hotter equipment. Within the equipment's permitted conditions, a warmer coolant can widen the opportunity to release heat without

refrigeration. Coldness is a sensation; heat transfer follows temperature differences.

When refrigeration joins the route, it brings an extra heat load of its own. A refrigerator warms the kitchen while cooling its contents. At data-center scale, the same energy account continues outside, where the weather decides how readily the heat can leave. [3,4]

Refrigeration adds heat outdoors.

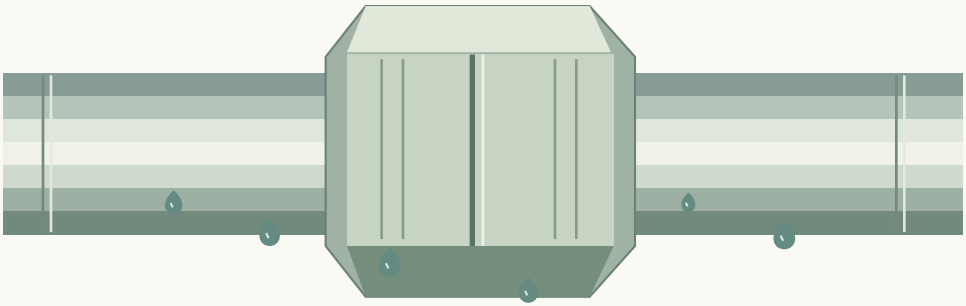
1.0 MW	+ 0.2 MW	1.2 MW
Cooling duty	Compressor	Heat out

Worked example at cooling COP 5; auxiliaries extra.

When a fitting gets wet.

Comfortable room air can still leave water on a cold connection.

Room air: 25 °C / 60% RH



Condensation on the outside

15 °C

Exposed surface

16.7 °C

Approximate dew point

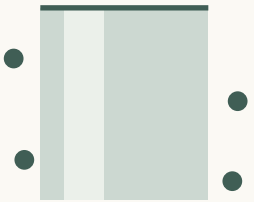
The surface is colder than the dew point. Water vapour in the air can condense on the metal.

When a fitting gets wet.

Air at 25 °C may feel unremarkable. At 60 per cent relative humidity, it has a dew point of approximately 16.7 °C, using the National Weather Service approximation. Bring that air against an exposed surface at 15 °C and water vapour can condense.

The droplets record a meeting between the local air and the surface. A room sensor measures conditions at its own

A cold surface meets moist air.



15 °C surface

16.7 °C dew point

position. The liquid inside a pipe and the outside of its insulation can have different temperatures.

That makes condensation a local story. Humidity, insulation and the exposed fitting all matter. A broad description such as ‘warm-water cooling’ cannot describe the temperature of every connection the air can reach. [9]

The air meets the outer surface.



Local surface conditions determine condensation.

Water from the air.

The glass stays intact. Droplets appear outside.



Illustration by Data Center Edition. Imagined scene; details and attribution in the closing notes.

Water from the air.

The glass has acquired water without giving up its drink. Vapour from the surrounding air meets its cold outer surface and condenses there. The little beads make an invisible ingredient of the room visible.

It is an ordinary sight on a café table. In a server hall, the same physics can happen beside electrical equipment. The air can feel comfortable to a person while carrying enough moisture to leave beads on a cold local surface. [9]

The good news in the bad score.

A constructed refresh delivers the same work with less electricity. Then the efficiency report arrives.

Before the refresh



1.30 MW total / PUE 1.30

After the refresh



1.00 MW total / PUE 1.43

Gold is the unchanged 0.30 MW supporting load. Same work, about 23% less total electricity over equal durations.

The good news in the bad score.

The new computers have done their job. IT demand has fallen from 1.00 MW to 0.70 MW. For this example, the same work still meets the same quality and timing, while supporting demand stays at 0.30 MW.

Total demand drops from 1.30 MW to 1.00 MW. Over equal durations, that means about 23% less electricity. Yet power usage

Same work. Less total electricity.



Green: IT. Gold: unchanged supporting demand.

effectiveness, PUE, rises from 1.30 to about 1.43: the facility total is being divided by a smaller IT total.

The report now contains a successful saving and a worsening ratio. Both are correct. The trouble begins when the score is allowed to tell the whole story. [10,11]

The denominator changes the score.

1.30	1.00
1.00	0.70
PUE 1.30	PUE 1.43

Facility demand divided by IT demand, in MW.

Computing continued through construction.

At Jefferson Lab, the upgrade had to make room for growth while the existing service kept working.

Above 2

PUE before the project

1.27

Reported after the project



Cooling. Power. Containment. Measurement.

Jefferson Lab case reported by DOE. Functional sketch, not a site plan.

Computing continued through construction.

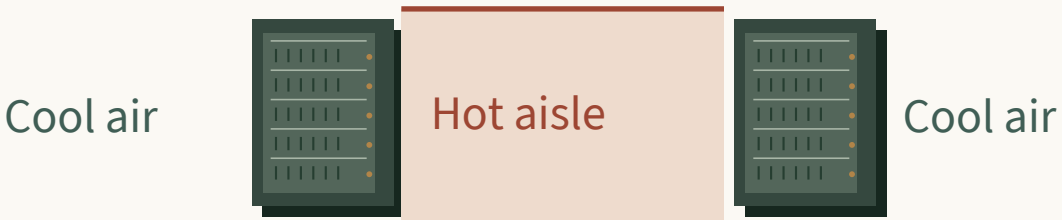
Computing continued while the renovation went on. At Thomas Jefferson National Accelerator Facility, that was part of the difficulty: two computing systems occupied different areas, supported by different cooling and power arrangements. Neither arrangement could accommodate expansion.

The project brought the computing systems into one area and changed the infrastructure around them. Sealed hot aisles separated rack exhaust from cool supply air. Cooling changes reduced

overcooling, while more efficient power equipment cut losses.

Staff and construction teams had to fit the work around an operating facility. DOE's account of the 2018 award-winning project credits coordination across those teams with maintaining its functionality through construction. The renovation had a growth

Separate the hot and cool air paths.



Functional sketch of the containment principle.

problem and a continuity problem to solve together.

Temperature sensors, electrical meters and flow meters made the changed operation easier to follow. The account reported PUE falling from more than 2 to 1.27. Behind that result were a new physical arrangement, changes to air and power, and work carried out while the computing service continued. [10,12]

Follow the changed operation.



Temperature



Electricity



Flow

The average that changes the answer.

Two half-year records can produce a misleading annual summary.

First half-year

150 / 100 MWh — PUE 1.50

Facility / IT energy; bar shows IT energy.

Second half-year

1,080 / 900 MWh — PUE 1.20

Facility / IT energy; bar shows IT energy.

1.23

Combined PUE

1,230 / 1,000 MWh

Combine energy totals, not the two ratios.

The average that changes the answer.

The first of two half-year reporting periods records 100 MWh of IT electricity and 150 MWh for the whole facility: PUE 1.50. The second records 900 MWh and 1,080 MWh: PUE 1.20. Their simple average is 1.35.

But the meters tell a combined story of 1,230 MWh for the facility and 1,000 MWh for IT. Divide those totals and PUE is 1.23.

The energy totals carry the weight.



100 MWh 900 MWh

IT energy: 10% first period; 90% second.

The simple average gave each period half the influence, although the first accounted for only a tenth of combined IT energy. A chart of the separate ratios can still reveal changes in operation. Turning that chart into one annual number takes the energy totals underneath it. [10,11]

Add the energy. Then divide.

1,230 MWh facility

1,000 MWh IT

1.23

Combined PUE for the two example periods.

When the electricity bill arrives.

An invented operating year puts a small change in PUE into money.

43,800 MWh

IT energy: 5 MW × 8,760 hours

Facility electricity



6,570 MWh

Avoided facility electricity

INR 52.56 million at INR 8/kWh. Other project costs excluded.

When the electricity bill arrives.

An efficiency score is dimensionless. The electricity bill arrives in currency. The distance between them depends on how much computing electricity the facility uses and what each kilowatt-hour costs.

Hold an illustrative IT load at 5 MW for 8,760 hours. It uses 43,800 MWh. With that IT energy unchanged, improving annual PUE from 1.40 to 1.25 removes 0.15 kWh of facility consumption for each IT kWh. Across the year, the difference is 6,570 MWh.

At an assumed flat INR 8 per kWh, those avoided units are worth about INR 5.3

crore (about INR 53 million) in energy charges. This calculation leaves out demand charges, tax, capital cost, maintenance and financing. A real project's economics would have to carry those costs too.

There is a further number an electricity meter cannot supply: the value of the

From avoided units to energy charges.

6,570 MWh

× 1,000 kWh/MWh × INR 8/kWh

INR 52.56 million

Assumed flat tariff. Other project costs excluded.

computing delivered. A job completed too late can consume exactly the same energy as one completed on time. Workload, quality and timing give the energy account its purpose. DOE's design guidance places that operating context alongside the facility measures. The saving means more when the service behind it is clear. [13]

Energy serves a computing task.



Workload



Quality



Timing

An energy total alone cannot describe the service.

The server water comes back.

Elsewhere on site, water can be leaving the cooling system.

A returning pipe gives a closed loop its reassuring shape. Water collects heat near the computers, gives it away and comes round again. The same liquid is ready for another journey.

At the exchanger, the heat finds a new carrier. A second circuit can take it towards an open cooling tower, where evaporation helps release it to the air. Some tower water leaves as vapour. The computer-side liquid continues circulating.

Both things can happen at once: a sealed server circuit and a facility that needs a continuing fresh supply. ‘Liquid cooled’ describes the first part of the route. The outdoor plant determines much of what happens next.

The water account therefore follows several travellers. Some water returns to the computers; some leaves the tower; heat passes between them without requiring their liquids to mix. [2]

Heat moves. The water returns.



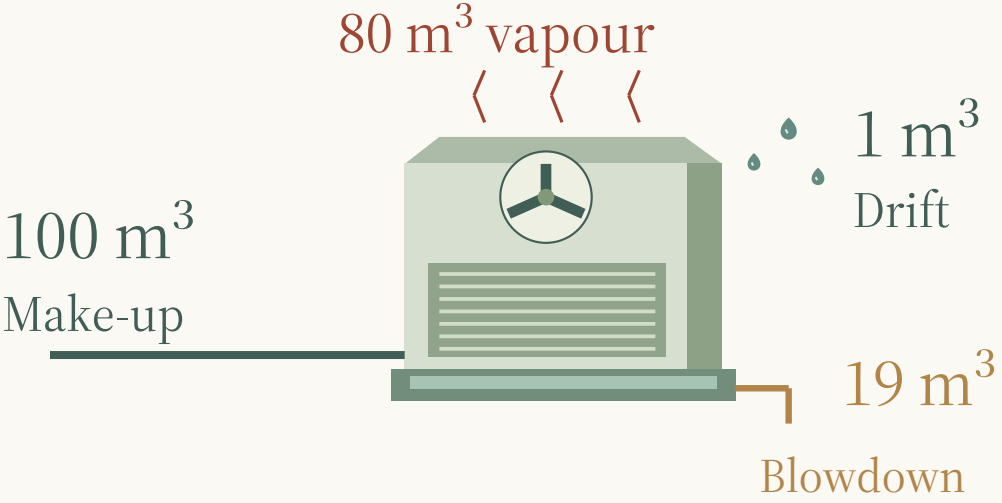
Computer loop

Building loop

One hundred in. One hundred out.

The tower’s replacement water has several departures to cover.

One supply. Three departures.



100 = 80 + 19 + 1 m³

Worked balance, unchanged inventory. Other losses omitted. Liquid drift enlarged for visibility.

One hundred in. One hundred out.

A tower loses water along different routes. Evaporation carries vapour away. Drift carries liquid droplets out with the air. Blowdown removes some water whose dissolved material has become more concentrated. The incoming make-up supply replaces these losses.

For an invented day, let 80 cubic metres evaporate, 19 leave as blowdown and one as drift. With stored volume unchanged and no other losses, 100 cubic metres must arrive. The drift figure is exaggerated so that it shows; DOE notes that drift loss is

small compared with evaporation and blowdown. [14]

The destinations matter as much as the sum. Eighty cubic metres left as vapour; twenty remained liquid as they left through other routes. Calling the entire supply evaporation would erase part of the water story. Leaks or overflows would add further departures.

Supply balances all three exits.



80 m^3	19 m^3	1 m^3
Evaporation	Blowdown	Drift

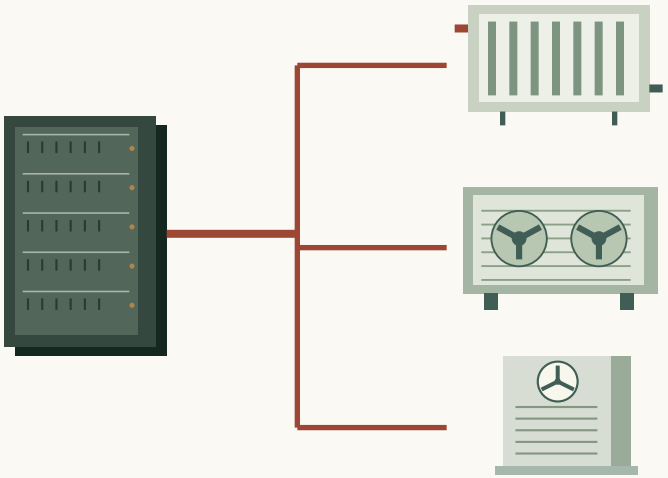
100 m^3 make-up; drift exaggerated for visibility.

A year in Golden, Colorado.

A historical NREL record captures one cooling system through a full year.

0.7	888	1.034
L/kWh site water	kW average IT	Reported PUE

One system. Several ways to release heat.



Heat recovery

Warm the building

Dry rejection

Release heat to air

Cooling tower

Reject heat outdoors

NREL, Golden. Operating year: Sep 2016 to Aug 2017.

A year in Golden, Colorado.

For a full year, heat at NREL in Golden, Colorado, had more than one way out. It could warm the building, reach the outdoor air through a thermosyphon dry cooler when conditions allowed, or go to a cooling tower. The installation had given the heat a choice of destinations.

The operating record ran from 1 September 2016 to 31 August 2017, the first year of the thermosyphon. NREL's March 2018 presentation reported site water use of 0.7 litres per IT kWh. Alongside it was a tower-only reference of 1.42 litres per kWh, with estimated avoided-water

components. The comparison paired a measured operating year with a reference route, rather than a second measured year.

The interesting result belongs to the combination: heat recovery, dry rejection and the tower taking their places as conditions changed. A single annual number preserves the outcome of that year in Colorado. It captures a cooling system with several answers to the weather. [2,15]

Heat reuse

Dry rejection

Cooling tower

What evaporation leaves behind.

Less water. The same amount of nonvolatile dissolved material.



Illustration by Data Center Edition. Imagined scene; details and attribution in the closing notes.

What evaporation leaves behind.

The two kettles hold the same eight gold markers. The second holds less water. Those markers stand for nonvolatile dissolved material: as water evaporates, the material remains and becomes more concentrated.

The little kitchen comparison explains why a tower needs more than a top-up. Its water is changing as evaporation continues. Blowdown carries some concentrated liquid away; make-up replaces losses. The actual treatment depends on the site’s water chemistry. [14]

The score falls. The litres rise.

Growth can outrun an improving water-use ratio.

Year one



10 million L

20 GWh IT

0.50 L/kWh

Year two



12 million L

40 GWh IT

0.30 L/kWh

20% more water

Despite a lower water-use ratio

Worked example. Pair annual litres with the intensity measure.

The score falls. The litres rise.

Two reports on the same facility can seem to disagree. One shows a 40 per cent drop in site water usage effectiveness, or WUE. The other shows a 20 per cent increase in annual litres. In our invented example, both are correct.

Year one pairs 10 million litres with 20 GWh of IT electricity: 0.50 litres per kWh. Year two pairs 12 million litres with 40 GWh: 0.30 litres per kWh. WUE divides annual site water use by annual IT electricity. The denominator grew faster than the litres.

For the surrounding community, the calendar matters too. Annual litres can conceal a heavy demand in dry months. Sources and liquid return flows complete the picture: water delivered to a site, water withdrawn upstream and water consumed through routes such as evaporation are different records. [16,17]

A lower ratio can accompany more water.

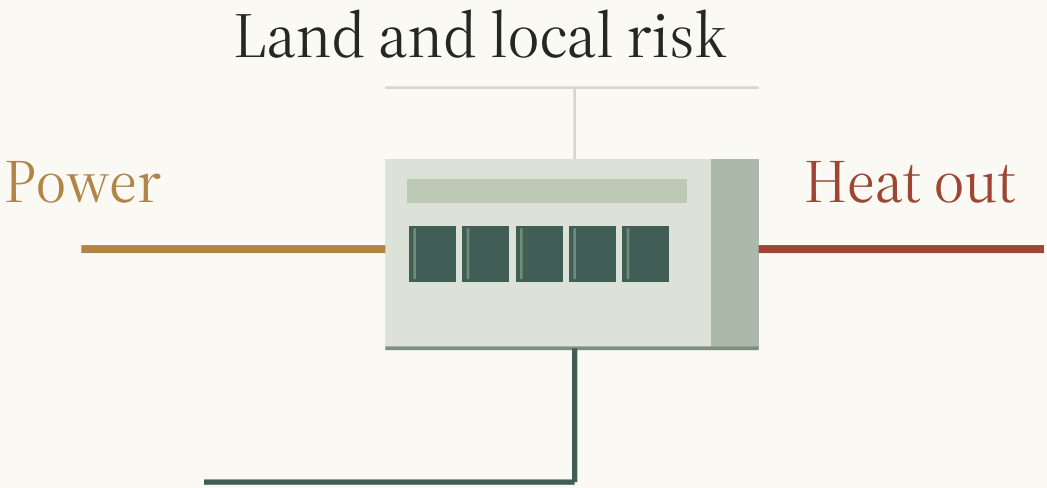
10 million L	0.50 L/kWh
12 million L	0.30 L/kWh

IT energy rises from 20 to 40 GWh in this example.

The site is a system.

Land, power, cooling and connectivity have to work together.

The building is only the meeting point.



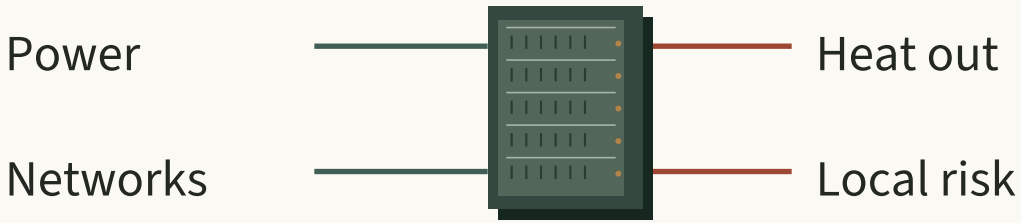
Functional relationships. Each connection needs site-specific evidence.

The site is a system.

A site is more than a parcel on a map. A useful shortlist asks whether electricity can arrive when needed, heat can leave under local weather conditions, and networks can connect the workload to the people and systems it serves.

Water risk is one part of that search. WRI describes Aqueduct as a screening tool using catchment-level information, not

A site connects several systems.



real-time measurements. Its results help identify where deeper local investigation is needed; they do not certify the water available to an individual project. [18]

The questions are global. Their answers are local: a connection agreement, a drainage survey, a cooling design and a credible route into service. No regional label can substitute for that evidence.

From an address to a working rack.

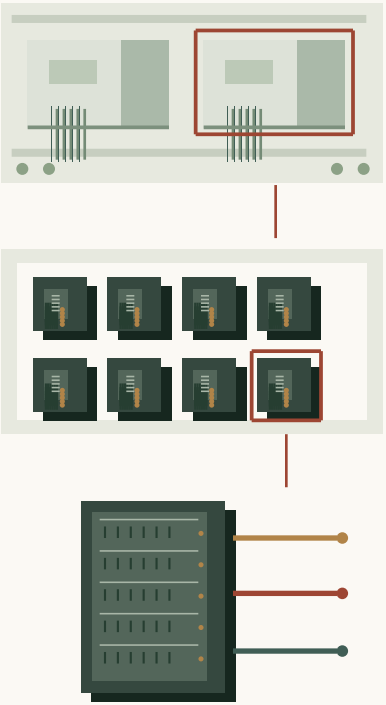


The room, connections and service date matter.

At the end of the move.

The address narrows from a region to the place where a rack will stand.

From campus to rack.



Region / campus

Where the project sits

Phase / room

The space ready for equipment

Rack position

Power, cooling and networks meet here

A sequence of decisions, not a map or a floor plan.

At the end of the move.

For the people moving computers, the address gets smaller as the project gets closer: the region, the campus, the phase, the rack. In our example, the first 20 MW phase is one fifth of the planned 100 MW campus. It gives the development a beginning; it does not prove every future phase is available.

For the customer, the part serving its racks is where the investment becomes tangible.

Power

Cooling

Network

Power must reach those positions. Heat must leave them. Network paths must reach the required destinations. Each interface needs an owner and an acceptance condition.

Those questions should be settled before equipment travels. A completed shell, a utility connection and a room accepted for service are related milestones, but they are not the same milestone.

From an address to a working rack.



A promise is not a load.

Commercial commitments and electrical demand follow different calendars.

40 MW

Illustrative building capacity



30 MW

75% commercially committed

Demand needs a meter reading.

Each mark is 1 MW of capacity. A contract is not measured load.

A promise is not a load.

Consider a second worked example: a building with 40 MW of capacity, of which customers have contracted 75 per cent before opening. Thirty megawatts are spoken for. That is useful commercial information, but it says nothing by itself about the opening-day meter reading.

The example is illustrative, not an operator announcement. Equipment might arrive in stages. Workloads may ramp gradually. The building also consumes electricity for cooling, electrical losses and other supporting services, so

the measurement boundary must be stated.

A commitment describes an agreement. Capacity describes what is available or planned. Load describes electricity being drawn at a particular moment. Comparing the three is useful only when the words remain attached to the numbers.

An agreement is not a meter reading.



40 MW building capacity

75% contracted. Actual electricity demand is separate.

Four kinds of number.

One vocabulary for comparing projects around the world.

100 MW

Campus design

The whole plan

20 MW

In service

The first operating phase

75%

Committed

Of a separate 40 MW building

12 MW

Measured load

A meter reading

Separate examples. Keep the boundary, status and date with each number.

Four kinds of number.

Our illustrative campus has a 100 MW plan and a 20 MW phase in service. Our separate 40 MW building is 75 per cent committed. A sample meter reading is 12 MW. None of these figures is presented as a real facility disclosure.

The same distinctions apply to published announcements. First identify the boundary: a rack, a building, a campus, a

Keep each number with its meaning.

100 MW

Campus plan

20 MW

Phase in service

portfolio or an entire grid. Then identify the status and date. Finally ask whether the figure is capacity, a contract or measured demand.

A design can take years to fill, a commitment is not yet a load, and one reading is not a year of consumption. Each needs its date.

Keep each number with its meaning.

75%

Committed

12 MW

Meter reading

Separate worked examples, with different boundaries.

A line needs a clock.

A worked series shows demand changing; energy still needs a time interval.

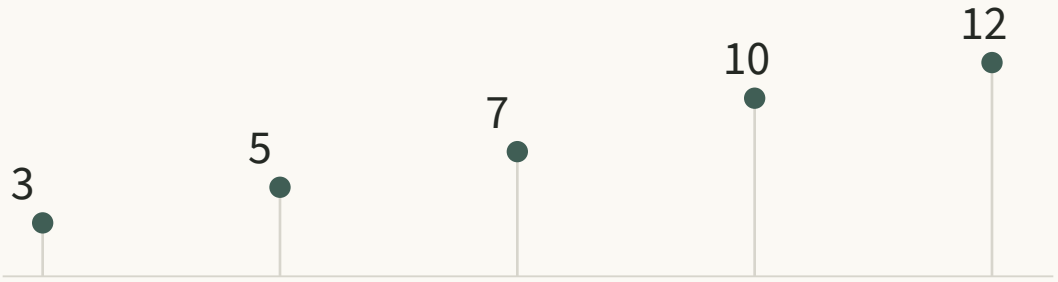
In this illustrative series, a meter reads 3, 5, 7, 10 and 12 MW at five monthly observations. The final reading is four times the first. That does not mean consumption over the whole period is four times larger: five snapshots leave the hours between them unknown.

If a 12 MW load were held steadily for 24 hours, it would use 288 MWh. If the load moved, the total would come from adding the power used through every interval.

The chart on page 9 shows that difference using two daily profiles.

Before comparing two sites, keep their measurement boundaries and periods together. IT electricity, total facility electricity and grid imports describe different things, even when every value is printed in megawatts.

Five observations leave hours unseen.



MW at five monthly observations; not an energy total.

A larger electricity story.

A global outlook, and the physical routes that make it possible.

485 TWh

IEA estimate for 2025



950 TWh

IEA central projection for 2030



An outlook, not an opening schedule.

IEA, April 2026. The future value is a projection, not a recorded result.

A larger electricity story.

The IEA's April 2026 central outlook puts data-center electricity use at an estimated 485 TWh in 2025 and a projected 950 TWh in 2030. The latter is a projection, not a recorded result or a guarantee that every proposed project will open. [19]

Electricity is only one connection. ITU's July 2026 cable-resilience report highlights physical risks, repair delays and concentrated routes. Network diversity and repair readiness belong in a site discussion alongside power and cooling. [20]

The screen turns green.

An imagined cooling test begins with the power coming back.



NASA Discover, 2010. NASA / Pat Izzo, CC BY 2.0. Separate from the imagined test.

The screen turns green.

The control screen looks reassuring again. Electrical supply has returned. The cooling pump shows RUN. After the interruption, those are welcome signs of movement back towards normal.

One branch, however, still has a closed valve. The pump is turning, but the equipment on that branch is receiving no flow. Water can be moving elsewhere in the system while that local cooling duty remains unmet. The room temperature has not yet given anyone much reason to worry.

This is the gap commissioning has to cross. A collection of healthy component indications can coexist with a service that has not recovered. The PNNL/ASHRAE/NEMA framework widens the testing from individual components to the interactions among systems, with the future operators involved.

In the scene on this page, the green light is accurate about the pump's indicated state. The cooling question sits farther along the pipe. Until that part of the story is visible, the apparently finished event still has something left to reveal. [21]

The reading that changes the result.

The imagined test continues. A second reading changes the meaning of the green light.

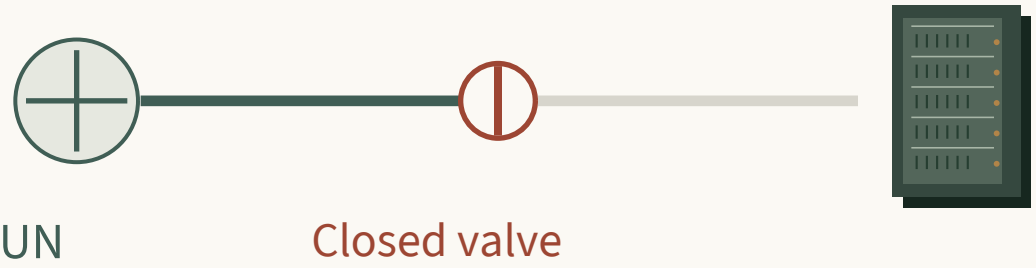
Two readings. The pump still shows RUN. Alongside that indication, the branch-flow measurement shows no delivered flow. The two observations describe different parts of the system: the machine is turning, but the shut branch cannot carry its cooling duty.

The correction. The test team locates the closed valve. Restoring the branch to its intended operating state gets water moving again. That intervention stays in

the event record, alongside the green indication that had first looked reassuring.

The repeat. The planned test is then repeated with the corrected configuration documented. This time, the branch receives flow. The recorded flows and temperatures demonstrate the return to the agreed operating conditions at the test load. DOE's commissioning guidance

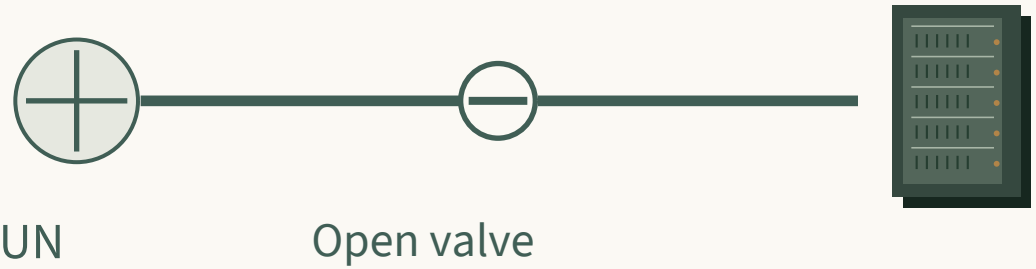
Initial attempt: no delivered branch flow.



places retesting after correction for this reason.

Both records. The first attempt included manual correction. The successful repeat establishes the response of the arrangement actually tested. The unchanged RUN indication can sit beside both records; the branch measurement explains why their outcomes differ. [21,22]

Repeated test: branch flow restored.



After the test team leaves.

Still the imagined test: the next shift inherits the corrected building and the record of how it got there.

The clean version. The final report could begin with the successful repeat. It would look cleaner. It would also leave the next shift with a gap in the building's history.

The full record. The earlier record explains what was corrected and why: the branch had not delivered its cooling duty. It records the manual help as well. The final drawings, operating arrangement, settings and sequence identify the configuration that produced the accepted result. Nobody has to infer that the

original problem cleared itself automatically.

Deferred work. Some work may still be waiting for a season the project has not experienced. A test that depends on different weather remains unfinished when the construction team packs up. WBDG's commissioning guidance keeps deferred testing in the record alongside

Leave a record the next shift can use.

Correction and retest

Accepted configuration

Open items and owners

operating information, training and completed tests.

The handover. The next shift receives both what has passed and what remains. A person is responsible for each open item; the accepted configuration can be traced to its test. The original team can leave without taking the explanation of the corrected branch, or the ownership of unfinished work, with it. [22,23]

Keep unfinished work visible.

Deferred test

Named owner

Agreed next step

From one machine to the whole building.

A common five-level sequence widens from factory checks to systems working together.

L1 **Factory**
Acceptance

L2 **Site**
Delivery and installation

L3 **Checks**
Pre-functional

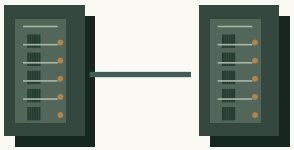
L4 **System**
Performance

L5 **Integrated**
Systems working together

Traditional five-level sequence. Earlier block testing is also possible.

From one machine to the whole building.

PNNL, ASHRAE and NEMA describe a traditional five-level sequence: factory acceptance, delivery and installation, pre-functional checks, system performance tests and integrated systems tests. Their guidance also allows more granular turnover and earlier testing of smaller blocks. The sequence is not a rule that every test must wait until delivery. [21]

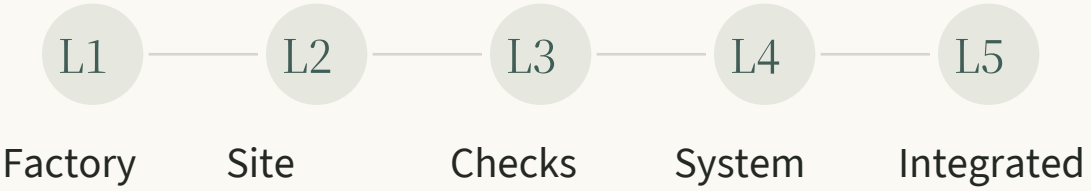


Test the connections, too.

Performance depends on the system.

The imagined branch test belongs to system and integrated testing. Factory results do not by themselves verify site interfaces or the complete facility. The branch, controls and operating team must deliver the agreed cooling duty together; the guidance recommends involving operators during commissioning. [21]

Testing reaches the whole building.

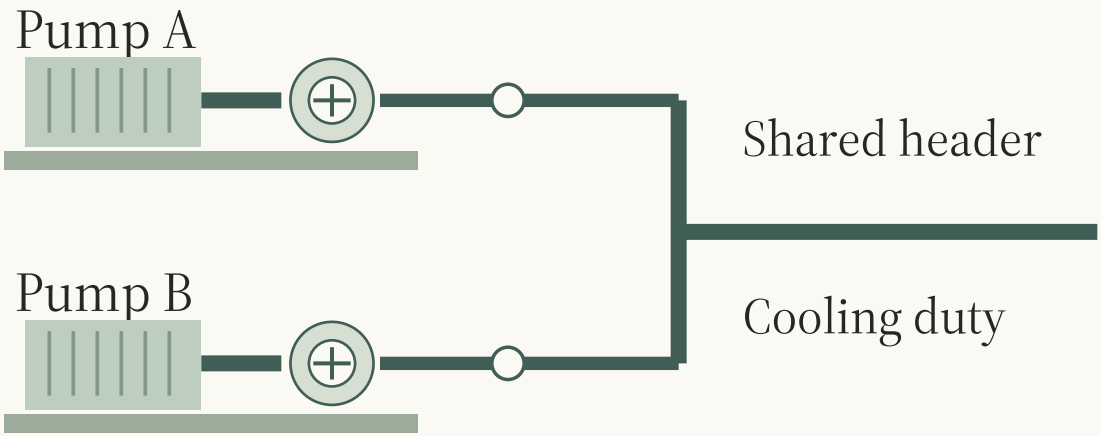


Traditional sequence; earlier block testing is possible.

Where the two lines meet.

The spare pump has capacity. The shared parts still have a story.

Either pump can carry the duty.



Shared control

Two pumps still share dependencies.

Functional sketch. A shared failure can interrupt both routes; no availability percentage is implied.

Where the two lines meet.

The spare is ready. Either pump in this sketch can carry the required duty, so an extra machine has clear value when the problem is confined to one pump.

Then the two lines meet. Both pumps depend on a common header and control. A failure affecting a shared part can interrupt the duty even though a second pump is available.

The pair contains duplication and dependence at the same time. That is why counting machines cannot produce an availability percentage. Uptime Institute’s public Tier descriptions draw the same

line at facility scale. At Tier II, with redundant capacity components, distribution failures still affect the site. Tier III asks that every capacity component and distribution path can be maintained without shutting the service down, and Tier IV adds tolerance of faults. The service rests on the complete arrangement, including the parts that have only been drawn once. [24]

Two pumps. One shared dependency.



Common header and control

The crossing both roads need.

An imagined landscape makes a shared dependency visible.



Illustration by Data Center Edition. Imagined scene; details and attribution in the closing notes.

The crossing both roads need.

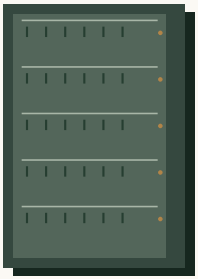
From the near bank, there is a choice of roads. Both can bring traffic towards the river. At the crossing, the choice narrows to one bridge.

If that bridge is unavailable, neither approach reaches the far bank. The second road has removed one dependence while leaving another in place. In the cooling sketch, a common controller, supply or header can occupy that same position. The point to find is where the apparently separate routes become one. [24]

The load moves closer together.

The total stays at 2 MW. At each rack position, the job changes.

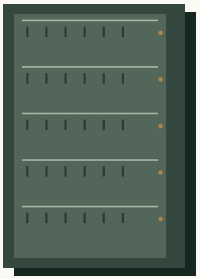
100 racks



20 kW

at each rack

20 racks



100 kW

at each rack

2 MW in either arrangement

Five times the duty at each rack. Counts and loads do not establish a floor plan or prove that equipment fits.

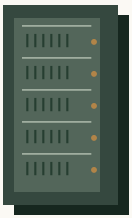
The load moves closer together.

A hundred racks draw 20 kW each. Twenty racks draw 100 kW each. Both arrangements use 2 MW, but the second concentrates five times as much electrical load at each rack.

Power distribution, heat removal and physical access now face that concentration. Any space gained from

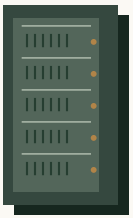
---

Same total. Five times the duty per rack.



20 kW

100 racks



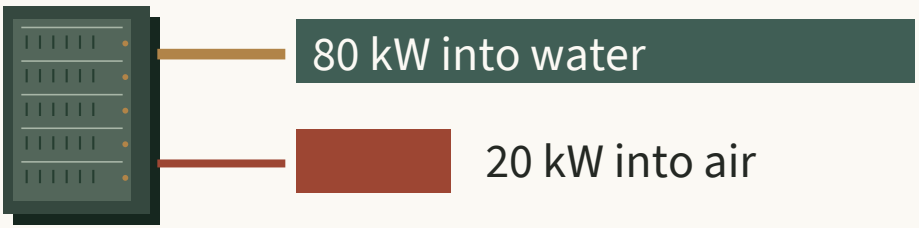
100 kW

20 racks

fewer cabinets does not, by itself, establish that the new arrangement fits the hall.

Using our assumed split, the twenty 100 kW racks send 1.6 MW into water and 0.4 MW into room air, before separately counted supporting loads. There are fewer rack positions, and much more duty at each position. The building's total has stayed the same while its local work has changed.

One rack. Two heat paths.



The price outside the package.

In this invented comparison, the lower equipment price starts a more expensive route.

INR crore	A	B
Equipment	4.8	5.3
Facilities work	0.8	0.4
Integration	0.2	0.1
Acceptance	0.2	0.2
Five-year service	1.0	0.7
Expanded total	7.0	6.7

Invented comparison. Energy, water, tax, financing and discounting are excluded.

The price outside the package.

Option A opens at INR 4.8 crore; B at INR 5.3 crore. The first line favours A by half a crore.

Facilities work, integration, acceptance and five-year service change the picture. With those assumed allowances included, A reaches INR 7.0 crore and B INR 6.7 crore. The ranking reverses because the

The wider scope, in INR crore.



Green: equipment. Gold: included wider scope.

first line left part of the work outside the price.

On a project, an exclusion can reappear in another quotation or another department’s budget. This table makes those connections visible. It still excludes energy, water, tax, financing and discounting: an expanded scope comparison, rather than a full lifetime cost or a supplier verdict.

The ranking reverses as scope grows.

Package	A 4.8 < B 5.3
Wider scope	A 7.0 > B 6.7

Four questions for the next meeting.

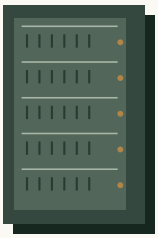
For the moment a promising offer has to become a defined project.

Which phase can receive the equipment, and when? Identify the room being offered and the work still standing between today's installation and the move-in date. Give every unfinished item a person to answer for it.

What arrives at each rack? Put the power, liquid-cooling duty and remaining room-air duty in the same discussion. Distribution, access and handling have to suit the actual equipment arrangement.

What result will establish recovery? Connect the proposed test to its load, settings, observations and accepted operating conditions. Make any human intervention visible, along with the operating team’s role and any deferred work.

Define what a successful test must show.



- Load and settings
- Flows and temperatures
- Agreed operating conditions

Record interventions and outstanding work as well.

What does the price include? Place facilities work, integration, acceptance and service beside the equipment. Record exclusions. For performance claims, retain the measured period, IT and facility energy totals, and the conditions that produced them. Comparable offers expose the complete work needed to deliver the same service. [10,21,22,23]

The ranking reverses as scope grows.

Package	A 4.8 < B 5.3
Wider scope	A 7.0 > B 6.7

Invented comparison, INR crore. Exclusions remain.

A few words, defined.

The shorthand becomes easier once you can picture the job each term does.

CDU: a coolant distribution unit. In a liquid-to-liquid arrangement, its heat exchanger transfers heat between separate computer-side and building-side cooling circuits. The liquids stay on their own sides.

PUE: total facility electricity, including the IT equipment, divided by IT electricity over the same period and reporting boundary. A PUE of 1.30 means 0.30 kWh of supporting electricity for each kWh

used by IT equipment. It does not measure the useful computing delivered.

WUE: annual site water use in litres divided by annual IT electricity in kWh. It is an intensity measure; the annual volume of water gives another part of the picture.

kW and MW: units of power, the rate of energy use or transfer. One MW is 1,000

Rate and accumulated use are different.

MW	MWh
Power at a moment	Energy over time

kW. MWh and GWh describe energy accumulated over time.

Commissioning: the process of establishing that installed systems meet the agreed requirements. Testing, operating information and the record of what passed help the receiving team understand the building it is taking over.
5,10,16,22,23

Keep unfinished work visible.

Deferred test

Named owner

Agreed next step

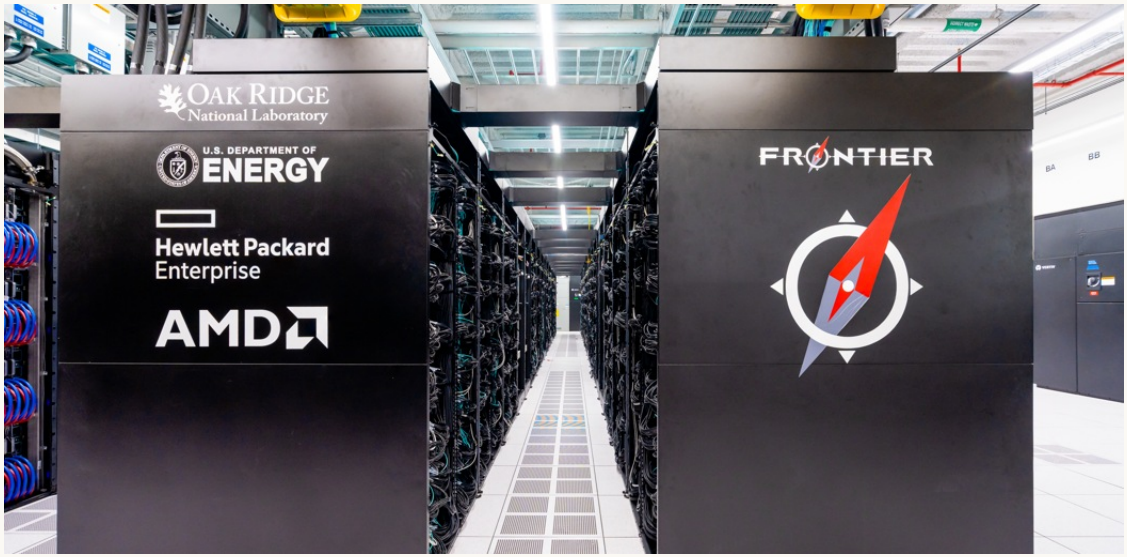
One connection wide.

A closing thought for our first issue.

On the cover, Eddie Blankenship's hand rests on a row of fibre connections in a cooperative's data center in Tennessee. Around it are cables, cabinets and the depth of a working installation. The hand gives the photograph its scale.

The industry is often described from much farther away. Announcements arrive in megawatts; forecasts reach across countries and years. Those views matter. They let us see the ambition and the resources gathering around computing.

Come closer and the work becomes particular. A connector has a socket. A rack has a place. A cooling branch serves equipment whose heat must leave. An operating shift takes over a building with a history of changes, tests and unfinished work.



Frontier at Oak Ridge, June 2022. OLCF / CC BY 2.0.

That closeness is what we want this magazine to keep. There is room for the campus announcement and the global forecast, and for the person working at the cabinet. Between those scales lie the details that turn capacity into a useful service. At the moment captured on our cover, the industry is one connection wide.



Eddie Blankenship at NCTC, Tennessee, 2018. USDA / Lance Cheung. Public domain.

Sources.

Primary records, dates and the scope of each citation.

1. USDA Rural Development, photograph 20180927-RD-LSC-1040 and its description, 27 September 2018
NCTC data center, Lafayette, Tennessee. Photo by Lance Cheung; public domain. Source of the technician’s role, the building’s services and the cooperative’s history.
www.flickr.com

2. National Laboratory of the Rockies (formerly NREL), High-Performance Computing Data Center Cooling System Energy Efficiency
Web page on the ESIF data center’s heat routes. NREL was renamed in December 2025.
www.nlr.gov

Sources.

3. Oak Ridge National Laboratory, Computer engineers at ORNL pioneer approaches to energy efficient supercomputing
10 September 2024; operator description; historical conditions retained.
www.ornl.gov

4. U.S. DOE, Fundamentals Handbook: Thermodynamics, Heat Transfer, and Fluid Flow, Volume 1, DOE-HDBK-1012/1-92, June 1992
Module 1, First Law and steady-flow energy balances.
www.energy.gov

5. Open Compute Project, Liquid Cooling Integration and Logistics White Paper, revision 1.0
PDF pages 6 and 7: terminology; manufacturing, transport and commissioning sections. The document carries no date.
www.opencompute.org

Sources.

6. NIST Chemistry WebBook (SRD 69), Water: liquid-phase heat capacity
Liquid-phase Shomate equation and coefficients, from the NIST-JANAF Thermochemical Tables, fourth edition, 1998. Undated web page; checked 26 October 2026.
webbook.nist.gov

7. DOE, Improving Pumping System Performance: A Sourcebook for Industry, second edition, May 2006
Hydraulic power, pump/system operation and efficiency context.
www.energy.gov

8. Swiss National Supercomputing Centre (CSCS), Lake water to cool supercomputers
June 2015; operator description; historical conditions retained.
www.cscs.ch

Sources.

9. NOAA/NWS, Relative humidity and dew-point equations
Vapour-pressure equations and dew-point inversion. Undated document; checked 26 October 2026.
www.weather.gov

10. The Green Grid, Power Usage Effectiveness (PUE)
Metric definition. Undated web page; checked 26 October 2026.
www.thegreengrid.org

11. The Green Grid, PUE: A Comprehensive Examination of the Metric, White Paper 49 (2012)
Printed pages 7 to 10: context, virtualization effect, partial PUE, measurement boundaries
datacenters.lbl.gov

Sources.

12. U.S. DOE FEMP, 2018 Award-Winning Champion Lessons Learned: Thomas Jefferson National Accelerator Facility Data Center Optimization
Project changes and reported PUE outcome. Undated web page; checked 26 October 2026.
www.energy.gov

13. U.S. DOE FEMP / NREL, Best Practices Guide for Energy-Efficient Data Center Design (2024)
Section 8, Benchmarking, printed pages 29 to 31
www.energy.gov

14. U.S. DOE FEMP, Best Management Practice #10: Cooling Tower Management
Evaporation, blowdown, drift and make-up water; cycles of concentration. Undated web page; checked 26 October 2026.
www.energy.gov

Sources.

15. NREL (now National Laboratory of the Rockies), DCDM1: Lessons Learned from the World's Most Energy Efficient Data Center, March 2018
Slide 37: operating year 1 September 2016 to 31 August 2017.
docs.nlr.gov

16. U.S. DOE FEMP, Cooling Water Efficiency Opportunities for Federal Data Centers, 9 January 2019
Cooling towers, WUE and water-efficiency measures.
www.energy.gov

17. U.S. Geological Survey, Water Science Glossary
Definitions of water withdrawal and consumptive use. Published 15 June 2018; checked 29 October 2026.
www.usgs.gov

Sources.

18. World Resources Institute, Aqueduct FAQ, accessed 29 October 2026
Catchment-level screening and the limits of baseline, non-real-time information.
www.wri.org

19. IEA, Key Questions on Energy and AI, executive summary, April 2026
Updated electricity-demand outlook
www.iea.org

20. ITU, International Advisory Body approves report to strengthen submarine cable resilience, 10 July 2026
Physical risks, repair delays, geographical concentration and route diversity.
www.itu.int

Sources.

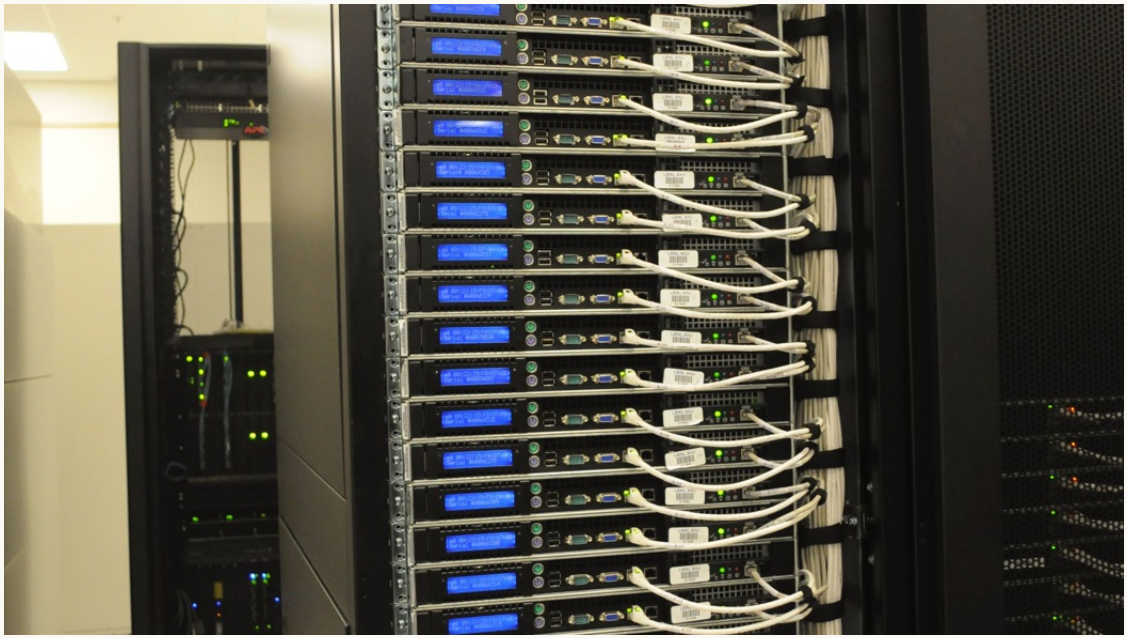
21. ASHRAE, NEMA and PNNL, AI Data Center Energy Performance Framework: Commissioning & Performance Validation, June 2026
Impact section: commissioning levels L1 to L5; recommended practice 11.
www.ashrae.org

22. U.S. DOE FEMP, Commissioning for Federal Facilities
Commissioning process chapter: owner requirements, design intent, correction and retesting. The document carries no date.
www.energy.gov

23. WBDG, Commissioning Documents: Process, Contents, and Acceptance
Systems Manual; Seasonal Testing; Final Commissioning Report. Undated web page; checked 26 October 2026.
legacy.wbdg.org

Sources.

24. Uptime Institute, Tier Certification
Tier II: distribution failures affect the site. Tier III: every capacity component and distribution path maintainable without shutdown. Tier IV: fault tolerance. Undated web page; checked 26 October 2026.
uptimeinstitute.com



Rear of a rack at NERSC, Oakland, 2011. Derrick Coetzee / CC0.

The pictures and the reporting.

This is desk-researched explanatory journalism prepared with AI assistance. It uses public primary sources and original worked examples. No site visit, interview, equipment test or external specialist peer review is represented. The opening story draws on USDA's published description of its 2018 photograph.

The lake, glass, kettles and bridge are imagined scenes made with AI image tools. They explain relationships; they are not documentary views, facility layouts or

measurements. Charts, diagrams and type are drawn separately.

Four companion experiments and three films are online at datacenteredition.com/lab and datacenteredition.com/watch. The films belong to their named publishers. The printed pages stand on their own.

Photographs are cropped for layout and do not imply endorsement. Typography: Source Serif 4 and Source Sans 3, SIL Open Font License.

Photographic credits.



CAMPUS, Frauenfeld, Switzerland, May 2022. Kecko / CC BY 2.0.

Kecko. 2022-05-27. CC BY 2.0.

Original photograph and rights record

Photographic credits.



NASA Discover, 2010. NASA / Pat Izzo, CC BY 2.0. Separate from the imagined test.

NASA / Pat Izzo. 2010-05-04. CC BY 2.0; also public domain NASA work in the US.

Original photograph and rights record

Photographic credits.



Frontier at Oak Ridge, June 2022. OLCF / CC BY 2.0.

OLCF at ORNL. 2022-06-01. CC BY 2.0.

Original photograph and rights record

Photographic credits.



Rear of a rack at NERSC, Oakland, 2011. Derrick Coetzee / CC0.

Derrick Coetzee (Commons credit; Flickr account now Chiara Coetzee). 2011-10-14. CC0 1.0.

Original photograph and rights record

Photographic credits.



Eddie Blankenship at NCTC, Tennessee, 2018. USDA / Lance Cheung. Public domain.

USDA / Lance Cheung. 2018-09-27. Public domain (US federal government work).

Original photograph and rights record

Keep exploring.

The story continues online.

Stories

Illustrated reporting and explanations.

Films

Three films by their original publishers.

Lab

Four interactive explorations.

Contribute

Ideas, expertise and stories from the field.