


The Modern Data Center

Physical Systems, the Four-Role Stack, and Campus-Scale Development



The Modern Data Center

Physical Systems, the Four-Role Stack, and Campus-Scale Development

Prepared August 14, 2026 · Updated August 20, 2026

1. Purpose

This brief consolidates three lenses on data center infrastructure that have come up across this project: (1) the **physical systems** that make a modern facility run, (2) the **four-role “stack”** the industry uses to describe who owns what, and (3) the **campus-scale, industrial-ecosystem view** now shaping how these facilities get planned, sited, and built. Individually, each lens is partial — a builder's parts list isn't a site-selection guide, and a single-building cutaway doesn't explain why interconnection queues now drive where projects land. Together, they're meant to leave no stone unturned.

2. Three Lenses on “What Is a Data Center”

A useful way to hold all of this at once is to separate the question “what is a data center” into three distinct questions, each answered by a different part of this brief:

- **How does it physically work?** The power, cooling, storage, and network systems inside one facility — Section 3.
- **Who defines “infrastructure,” and from where?** The builder / operator / network engineer / technician split — Section 4.
- **What does it take to site and build one?** Land, power, water, fiber, and community factors at campus scale — Section 5.

3. Physical Systems — How a Modern Facility Actually Works

This section reflects the updated systems diagram built earlier in this project (included as Figure 1). It moves away from the old model — grid power, a diesel generator for backup, and evaporative cooling towers — toward several shifts already underway at leading sites.

3.1 Power: On-Site Generation First, Grid Second, Generator Last

Grid interconnection queues in core markets now run 2–4 years, so more developers are treating on-site, behind-the-meter generation as the primary power source rather than a backup. In this model:

- **Primary:** solar PV, hydrogen fuel cells (20+ year service life, modular), geothermal (steady baseload), and — as an emerging option — small modular nuclear reactors (SMRs), all feeding an on-site microgrid tied to battery storage.
- **Secondary:** the utility grid becomes a tie-in and export path rather than the sole source.

- **Last resort:** the diesel/gas backup generator stays outside the building and is sized for occasional, short-duration use — not the primary safety net.

A precision on SMR specifically: most nuclear-for-data-center deals announced so far (Microsoft/Constellation's Three Mile Island restart, Amazon/Talen's Susquehanna plant) draw power from **existing large conventional reactors, not SMRs**. True small modular reactors (factory-built, 50–300 MW) are still pre-commercial: as of mid-2026, ~9.8 GW of nuclear capacity is committed across 13 hyperscaler deals, but only ~1.9 GW is operational — none of it SMR. The furthest-along genuine SMR project (GE Vernova Hitachi's BWRX-300, Ontario) targets completion around late 2029. Regulatory readiness (NRC's Part 53 framework isn't expected until late 2027), fuel supply (HALEU), and first-of-a-kind costs point to realistic broad SMR availability in the 2029–2035 window rather than “now.”

3.2 Cooling: Liquid at the Chip

Cold plates mounted directly on GPUs/CPUs pull heat off the chip into a closed coolant loop — now the dominant design for high-density AI racks. A Coolant Distribution Unit (CDU) bridges that rack-level loop to the building's facility water loop, which rejects heat outside via dry/fluid coolers. Water use for this approach is far below older evaporative cooling towers, which matters given water has become a site-selection gate in many markets (see Section 5).

3.3 Energy Storage: Batteries in the Room

Battery Energy Storage (BESS) now sits in the data hall itself, alongside UPS. The UPS still handles instant ride-through; BESS extends that protection from minutes to hours and does double duty outside of outages — shaving peak demand and providing grid services when the site is grid-interactive.

An emerging alternative to lithium-ion: vendors are starting to pitch electrostatic energy storage as a next-generation BESS chemistry for exactly this role — Newport Beach-based Critical Grid Agility markets it with a vendor-claimed >25-year lifespan and no thermal-runaway risk, versus a <10-year service life and known fire risk for lithium-ion, while also sidestepping lithium-ion's rare-earth-mineral dependency. Worth flagging as a vendor claim rather than an independently verified figure until corroborated by a second source or field data — but directionally it's a real reliability lever worth tracking alongside 3.6's redundancy stack, since a storage medium that doesn't degrade or combust changes the maintenance and safety math for in-room BESS. (This same vendor's broader platform claims — power generation and on-site water, not just storage — are picked up in Section 5's site-selection discussion.)

3.4 Heat Reuse: Warming the Community

In cold climates, the same warm-coolant loop that would otherwise go to dry coolers can instead feed a **community district-heating network** through a heat exchanger. Two operating examples:

- **Stockholm, Sweden:** the Open District Heating scheme had 20 data center suppliers by 2022, recovering enough heat to warm roughly 30,000 modern apartments per year and cutting the district system's emissions by about 50g CO₂/kWh.

- **Mäntsälä, Finland (Nebius):** the site's heat-recovery system reclaims about 20,000 MWh per year — enough to heat roughly 2,500 Finnish homes.




Figure 1. Updated data center systems diagram — renewables-first power, direct-to-chip liquid cooling, in-room BESS, and community heat reuse.

3.5 Physical Security: Drones and Robots for 24/7 Protection

Physical security has become its own automation category, and it pairs naturally with on-site generation (3.1):

- *Autonomous perimeter security*: “drone-in-a-box” systems (e.g., Asylon's Guardian, already deployed at critical infrastructure/energy/data center sites) patrol on schedule or launch on alarm, land and swap batteries in under 3 minutes for near-continuous coverage, and carry thermal/EO-IR sensors for day/night detection — footage routes to a human-in-the-loop SOC rather than replacing guards outright (Asylon alone: 45,000+ automated missions). Ground robots cover what a drone can't hold continuously (fence lines, building envelopes, interior corridors), typically sharing the same command center.
- *Drone-based solar panel cleaning*: dust/soiling is the #2 drag on solar output after weather — global average loss ~4–7%, but 5–50% at dusty/arid sites (the same sunbelt geography a resource-matched site is likely generating solar in). Frequent drone rinsing recovers ~8–15% more yield than seasonal manual cleaning on ground-mount utility arrays, without manual crews' water/labor cost.

A campus already running on-site solar gets a natural pairing: the same drone fleet built for perimeter security can extend into panel inspection and cleaning — turning part of the security budget into a partial generation-yield investment.

3.6 Redundancy & Reliability: N, N+1, 2N, and Uptime Tiers

The power and cooling systems in 3.1–3.2 aren't just sized for normal load, they're sized against failure, using shorthand the industry borrows from electrical engineering:

- **N**: the minimum equipment needed to carry the full load — e.g., three chillers, if three chillers are what the cooling load requires.
- **N+1**: one spare beyond that minimum (a fourth chiller, or a fifth UPS module where four are needed) — so one unit can fail or go down for maintenance while the rest keep carrying the load.
- **2N**: the entire system duplicated as two independent, full-capacity paths — so an entire path can go offline with zero loss of capacity.

The Uptime Institute's Tier system correlates roughly to these levels: **Tier III** facilities are “concurrently maintainable” and typically built to N+1; **Tier IV** facilities are fully fault-tolerant and typically built to 2N or equivalent. N+1 tends to be the practical sweet spot for most sites — most of the reliability benefit of full redundancy at a fraction of the capital cost of 2N.

This brief's own design choices are a redundancy stack in miniature: 3.1's primary/secondary/last-resort power hierarchy layers a fourth path on top of the traditional N+1 utility-plus-generator model, and 3.3's UPS-plus-BESS pairing is redundancy in *time* (instant ride-through, then extended runtime) rather than redundancy in equipment count.

3.7 Controls & Monitoring: The Building Management System (BMS)

None of 3.1–3.6 operate as isolated systems — a **Building Management System (BMS)** integrates the electrical and mechanical sides into one operating environment, continuously monitoring temperatures, power draw, equipment status, and alarm conditions, and triggering automated responses (failing over to a

backup power path, ramping cooling, paging operations staff) faster than a manual response could. In a modern facility the BMS is treated as being as critical as the physical equipment it watches over — it's the layer that turns the redundant hardware in 3.6 into an actual failover, rather than just spare capacity sitting idle.

This also grounds a reference that otherwise appears without explanation elsewhere in this brief: the Synthesis Checklist (Section 7) lists “BMS” under the Operator column for both Power and Cooling, and under Site/Land alongside physical security — that's the same system described here. Section 6.5 makes the case that this monitoring layer needs to extend to millisecond-scale electrical events, not just thermal/power dashboards.

4. The Four-Role Stack — Same Layer, Different Perspectives

This framing (circulating in industry commentary this year) observes that if you ask a builder, an operator, a network engineer, and a data center technician to define “infrastructure,” you get four different but equally valid answers — each person is describing the same stack from where they sit.

Role	What They're Focused On	Example Components
Builder	Constructs the facility — the shell, power plant, and mechanical systems	Transformers, generators, switchgear, chillers
Operator	Operates the facility day to day — keeps it available and efficient	UPS systems, PDUs, cooling, BMS
Network Engineer	Gets data in and out — the whole reason the building exists	Fiber, routers, switches, cross-connects
Data Center Technician	Hands-on with the compute itself	Racks, servers, storage, cabling

Why it matters here: mapping our physical-systems diagram (Section 3 / Figure 1) onto this framework surfaces real gaps. The **Builder** layer is underrepresented — transformers and switchgear are implied but not drawn. The **Network Engineer** layer is the thinnest — a single “network switch” box stands in for fiber entry, core routing, and cross-connects, which are really their own subsystem. The **Operator** and **Technician** layers are comparatively well covered (UPS/cooling/BMS and racks/cabling, respectively). Closing those two gaps is an open item in Section 7.

5. Campus-Scale Development — The Industrial Ecosystem View

Data centers are increasingly developed as large industrial campuses — combining compute, power, cooling, fiber, backup systems, and room for future phases — rather than single buildings. Several forces are driving that shift in 2026:

- **Power-driven, not land-driven, site selection:** utility capacity now outweighs zoning or fiber access as the primary siting filter. Established hubs (Northern Virginia, Chicago, Phoenix, Dallas–Fort Worth, Atlanta) are congested, pushing developers to secondary and tertiary markets with more predictable grid capacity.
- **Interconnection timelines:** utilities now need 2–4 years to connect large loads in core markets, so power negotiations must start much earlier in the site-selection process than before.
- **Behind-the-meter generation and storage:** more projects are self-generating (fuel cells, gas turbines, solar storage, and eventually SMRs) specifically to route around interconnection delays — despite higher upfront capital cost. This is the same shift reflected in Section 3.1.
- **Water as a site-selection gate:** no longer a “sustainability appendix” item. In Texas, data centers are projected to account for 3–9% of state water use by 2040. Residents understand well levels, low pressure, and drought restrictions better than PUE or WUE metrics, which makes water a genuine political risk, not just an engineering one.
- **Fiber connectivity corridors:** proximity to fiber routes remains a core sitting criterion alongside power, though power has overtaken it as the binding constraint in most markets.
- **Master planning for phased growth:** civil engineering plans now build in internal roadway networks, utility corridors, stormwater management, and land banked for future expansion pads from day one, rather than treating each phase as a separate project.
- **Community and permitting risk:** local opposition has become more organized in power-constrained metros. Political scrutiny is now a real site-selection risk, and markets that can demonstrate credible, on-schedule permitting are pulling ahead of competitors — regardless of the incentives on offer.
- **Scale of the buildout:** JLL projects roughly 100 GW of new data center capacity globally between 2026 and 2030, requiring on the order of \$3 trillion in combined investment — an “Infrastructure Supercycle” that rewards execution speed and certainty over incentives alone.

Zero Grid Dependence: When Fiber Becomes the Only Site-Selection Criterion That Can't Be Engineered Around — A more aggressive version of “power-driven site selection” is starting to show up in vendor pitches: if a site's power, water, and storage are all fully self-generated on-site, does the grid or municipal water system need to factor into site selection at all? Critical Grid Agility (the same Newport Beach vendor whose electrostatic battery storage claim is flagged in Section 3.3) markets exactly this position — a single zero-grid-capable platform fusing microgrids, fuel cells, solar, geothermal, and, they say, next-generation SMR nuclear for power (an integrated version of the 3.1 primary/secondary/last-resort stack), paired with on-site water capture from two sources — recovery from data-center waste heat and generation drawn directly from the atmosphere — backed by the same next-generation battery storage already noted in 3.3. Under this framing the local grid and municipal water utility become the fallback, not the dependency, and the site can be sited “anywhere,” including fully off-grid and remote — leaving proximity to fiber-optic backbone infrastructure as the one input that still has to come from somewhere else, since connectivity can't be generated on-site the way power and water can. For a land search, that reframes the checklist: instead of screening parcels against grid capacity, water rights, *and* fiber access, a fully self-sufficient site only needs to clear the fiber-proximity bar.

Three caveats worth separating from the pitch before treating this as a general site-selection rule. First, it's currently one vendor's integrated-platform claim, not an industry-wide capability — most of the projects behind this section's other bullets are still sizing on-site generation as a hedge against interconnection delay (Section 6.2), not as a full grid replacement, and SMR nuclear specifically is not yet available on the timeline 3.1 already lays out (2029–2035 at the earliest). Second, atmospheric water generation at data-center-relevant volumes is a new claim beyond anything else in this brief's water discussion (Section 5's “Water Rights vs. Effective Alternatives” below is about reducing withdrawal via closed-loop/immersion cooling, not manufacturing water from air) and deserves the same corroboration flag already applied to the storage-lifespan claim in 3.3 before it's treated as a planning input rather than a marketing claim.

Third, and specific to the fuel-cell leg of the platform: fuel cells consume a natural gas or hydrogen fuel stream, and where that stream comes from changes what “fiber is the only requirement” actually means. Every fuel cell deployment already documented in this project — Bloom Energy at Equinix, FuelCell Energy at Bridgeport/Pfizer/Hwaseong, the Cheyenne hydrogen system — runs on pipeline-delivered natural gas or hydrogen, not gas extracted on-site. Under that pattern, gas pipeline proximity is a second “connect to it, don't own it” requirement alongside fiber — not a land-ownership issue, but a real second dependency this section's headline claim currently omits. The exception is a stranded/flared-gas model — the approach Crusoe Energy (named among the behind-the-meter operators in Section 6.2) uses at active oil/gas fields, generating power directly from gas that would otherwise be flared on-site. If this platform is ever sited that way — consuming gas produced from under the parcel itself, rather than piping it in — mineral rights become a real site-selection gate: in most U.S. oil/gas-producing states (Texas, Oklahoma, Wyoming, and others), subsurface mineral rights are legally severable from surface rights (a “split estate”), so owning or leasing the surface doesn't automatically grant the right to the gas beneath it. That right has to be secured separately — often from a different owner entirely — and is its own title search, negotiation, and lease process before any extraction can happen. Worth deciding explicitly which fuel-sourcing model the platform is actually built around (pipeline delivery vs. on-site/stranded-gas extraction) before finalizing “fiber is the only requirement” for external use — the claim holds cleanly for the solar/geothermal/storage legs, but needs a footnote or a second dependency called out wherever fuel cells are the generation source.

Water Rights vs. Effective Alternatives — Water rights are a legal/political constraint, not just an engineering variable: a well permit, municipal allocation, or drought contingency plan can kill or delay a project regardless of facility efficiency. This brief's own diagram (Section 3.2, direct-to-chip + dry coolers) already sits well below old evaporative-tower water use, but some operators push further, toward effectively zero water after commissioning:

- *Closed-loop liquid cooling* (Microsoft): no additional water intake after initial fill — reportedly saves 125+ million litres/year per facility vs. evaporative designs.
- *Closed-loop air cooling* (Evolution Data Centres): built for high-temp, humid climates (SE Asia), eliminating evaporative water dependency entirely.
- *Immersion cooling* (Iceotope, LiquidStack): servers submerged in non-conductive fluid, passive/two-phase circulation — no pumps, fans, or water.

- *Hybrid air + closed-loop* (Vertiv/Bridge/Chindata “X-Cooling”): ~1.2M tons of water saved/year per 100 MW at a Hebei, China deployment vs. up to 1.5M litres/day for traditional evaporative systems; up to 80% WUE improvement reported.

Open question for any specific site: is local water risk severe enough to justify moving from this brief’s “low water” baseline to “near-zero water” (immersion/fully closed-loop), given the added capital cost?

Dual-Use Land: What Existing Asset Can We Make More Productive? — Land is one of the biggest inputs to a campus; the most overlooked site-planning question is what the site already has that could do double duty, rather than what greenfield acreage needs to be newly disturbed:

- *Parking canopies*: Dell’s Franklin, MA data center campus paired a 1.60 MWdc solar canopy over its parking lot with a 0.87 MW ground-mount system — power without consuming additional acreage, plus vehicle shade and stormwater management.
- *Floating solar (floatovoltaics)*: stormwater retention ponds/reservoirs a campus already needs can carry floating solar arrays — generation on water surface area that would otherwise sit unused.
- *Brownfield and landfill reuse*: the EPA’s RE-Powering America’s Land initiative promotes siting renewable generation on contaminated sites and landfills — avoiding the friction of converting undisturbed greenfield or farmland.

The same principle applies to generation choice itself: the best resource for a site is usually whatever is genuinely abundant in that geography (solar in the sunbelt, geothermal where geology supports it, wind where it’s steady) — drawing on a local surplus rather than competing with the community for something scarce.

6. The Industrial Turn — Rack Density, Speed to Power, and Medium/High-Voltage Distribution

Source: “Your Data Center Is Now a Factory. Nobody Sent the Memo.” (Parts 1–4), Drew Allen, CEO, Grace Technologies.

Where Sections 3 and 5 describe a data center’s steady-state systems and site-selection factors, this section covers a faster-moving shift underneath both: rack power density has outrun the industry’s operating practices, and the electrical plant is being redesigned — and increasingly built like a factory product — to keep up.

6.1 Rack Density Has Left the Building’s Original Design Behind — The Uptime Institute’s 2025 survey put industry-wide average rack density at 7.6 kW. The same year, NVIDIA shipped the DGX GB200 NVL72 as a standard product: one liquid-cooled, rack-scale system pulling ~120 kW — roughly 16x the industry average. The follow-on GB300 NVL72 reference architecture lists up to 142 kW per rack, and NVIDIA and partners (ABB, Eaton, Schneider Electric, Vertiv) are building an 800 VDC power ecosystem for megawatt-class Kyber/Rubin Ultra systems targeted for 2027. A 120 kW rack draws roughly 144 A per phase at 480 V three-phase (~167 A at 415 V) — current that used to describe an entire row, now landing on one rack’s whips.

Three hundred such racks put a single data hall at ~36 MW of IT load before losses and redundancy; a full campus typically interconnects at transmission voltage with a grid relationship closer to a steel mill's than an office park's.

6.2 Speed to Power Is Rivaling PUE as the Metric That Matters — Two decades of chasing PUE to the second decimal place are giving way to a blunter question: how long from site selection to energized IT load? Enverus reports projects reaching first power in 2025 averaged 2,100+ days from interconnection queue entry to commercial operation — nearly six years, ~60% longer than in 2017. Lawrence Berkeley National Laboratory counted 2,289 GW of generation and storage in US interconnection queues at end-2024 (preliminary end-2025 figures ran closer to 2.1 TW) — more capacity than the country's entire installed base, of which historically only a small fraction ever gets built. Demand isn't waiting: OpenAI has announced ~7 GW of planned Stargate capacity and \$400B+ of planned investment (its Abilene campus ran early workloads in 2025 while the 1.2 GW first phase was still under construction); Meta's Hyperion campus in Louisiana lists 2+ GW with leadership signaling a path toward 5 GW. The IEA projects global data center electricity consumption nearly doubling, from ~485 TWh in 2025 to ~950 TWh by 2030 — approaching 3% of all electricity generated on Earth.

This directly reinforces Section 5's “behind-the-meter generation and storage” bullet: Enverus has identified 40+ GW of *announced* behind-the-meter generation across Meta, Amazon, Microsoft, Google, Oracle, xAI, Crusoe, and others — gas turbines, fuel cells, grid-interactive batteries, anything that energizes load while interconnection paperwork proceeds in parallel. “Announced” is doing real work in that sentence; only a fraction is actually financed, permitted, or running. A few operators are reportedly buying retiring power-plant sites specifically for their interconnection rights.

6.3 Prefabrication: The Electrical Room Arrives on a Truck — When schedule dominates the economics, construction tends to become manufacturing — the pattern shipbuilding and petrochemicals went through earlier. Prefabricated/modular data center construction is estimated at ~\$34.8B in 2025, growing ~17%/year, with modularized projects reporting 30–50% schedule reductions versus conventional builds (vendor-sourced figures, but directionally consistent). Mechanically, it's parallelism: while site crews pour concrete, factory lines assemble switchgear, UPS, and controls on a skid — wired, tested, and certified before the truck leaves the plant — so “build the electrical room” becomes “set the module and connect it” (Vertiv's prefabricated power rooms and Schneider Electric's AI-ready prefab/liquid-cooling-skid line are the visible examples). Beyond schedule, this shifts three things worth tracking on this project: factory QA replaces field variance (workmanship defects — the seed stock of future failures — drop); the labor problem shrinks, since scarce electricians and controls techs work near-home factory shifts instead of chasing field projects (nearly two-thirds of operators report staffing difficulty); and safety hardware — voltage indicators, test points, absence-of-voltage testers, continual thermal monitoring — becomes a specified product feature installed on the factory floor rather than a per-site retrofit after an arc flash study flags the gap.

6.4 Medium Voltage Moves Into the White Space — Serving a 36 MW hall at 480 V requires roughly a dozen 3,000+ kVA transformers and feeder ampacities in the tens of thousands of amps. Physics and economics push the other way: bring medium voltage as close to the load as code and maintenance philosophy allow.

Newer high-density designs put MV switchgear and cast-resin transformers next to — sometimes inside — the data hall, with NEC 110.26/110.34 working-space rules shaping room layouts; some designs distribute at 4.16 kV or higher to skid-mounted power modules serving individual halls. Inside the rack, AC is running out of mechanical headroom (a 120 kW rack's ~167 A at 415 V doesn't scale cleanly to megawatt racks), and two DC architectures are racing to replace it: **±400 VDC** (OCP's "Mt. Diablo" spec, co-authored by Meta/Microsoft/Google, v0.5 in May 2025) — a disaggregated sidecar power rack feeding ±400 VDC to IT racks, scaling 100 kW–1 MW, riding the EV supply chain for converters/contactors/fuses/cable (Microsoft cites ~3% end-to-end efficiency gain); and **800 VDC** (NVIDIA's ecosystem with ABB, Eaton, Schneider Electric, Vertiv) — sidecar pods first, with a proposed end state of facility-perimeter solid-state transformers rectifying medium voltage straight to an 800 VDC bus feeding native-DC Kyber/Rubin Ultra racks from 2027 (Vertiv has said parts of its 800 VDC portfolio arrive in H2 2026).

6.5 What This Changes for Electrical Safety — This is the section's underlying thesis: *the equipment industrialized faster than the practices did*. Fault current in modern switchgear specs is moving into the 65–100 kA range, and arc energies into territory refineries and generating stations learned to manage decades ago under safety regimes built for exactly that — while much of the data center industry's operating culture (calendar-based or absent maintenance, hot work justified by uptime, monitoring that can't catch a millisecond event, PPE-based rather than engineering-control-based safety programs) still belongs to the low-density IT-building era. On-site generation and BESS (Sections 3.1, 3.3) change fault current, protection coordination, and arc flash boundaries in their own right — every "temporary" power configuration is a real electrical configuration that studies need to cover. DC distribution (6.4) adds its own fine print: DC arcs have no current zero and don't self-extinguish the way a 60 Hz arc can, so protection depends on DC-rated breakers or fast semiconductor devices; stored energy in sidecar converters and rack buffers means "breaker open" no longer means "bus dead"; grounding schemes change what a verification sequence looks like (±400 V designs use high-resistance grounding at the pole midpoint); and product standards for DC switchboards at these ratings are still maturing — data center practice is currently driving standards development in this area rather than following it. NFPA 70B's move from a recommended practice toward an enforceable standard raises the stakes on all of this, and ties directly to the BMS discussion in 3.7: monitoring built for millisecond-scale electrical events, not just thermal/power dashboards, is what would actually catch the failure modes described here.

Open framing question raised by this source: organizations that treat this as a deliberate industrial-safety redesign — not just a bigger version of the old IT building — are positioned to own the next decade of this build-out; those that carry IT-facility habits into industrial-scale plants are set up to relearn heavy industry's older lessons the expensive way.

7. Synthesis — A "No Stone Unturned" Checklist

Cross-referencing the physical systems against the four-role stack and campus-scale considerations gives a working checklist for reviewing any site or design against all three lenses at once:

System	Builder	Operator	Network Eng.	Technician	Campus-Scale Consideration
Power	Transformers, switchgear, on-site generation plant (6.3–6.4)	UPS/ATS operation, BMS setpoints, load balancing, DC protection & grounding verification (6.5)	—	PDU whips, breaker panel labeling	Interconnection queue timing, behind-the-meter capacity, substation siting, speed-to-power scheduling (6.2)
Cooling	Chiller/CDU plant, piping runs, dry coolers	CDU operation, water treatment, BMS alarms	—	Cold-plate install & maintenance	Water rights, drought risk, community water perception
Network	Conduit and duct banks	Uptime monitoring, capacity planning	Fiber routes, cross-connects, routers	Structured cabling, patch panels	Fiber corridor access, carrier proximity
Storage / Racks	Floor loading, seismic bracing	Capacity planning, DCIM, rack density planning (7.6 kW avg vs. 120–142+ kW AI racks, 6.1)	—	Rack installation, cabling	Expansion-pad phasing for future racks
Site / Land	Grading, roads, stormwater	Physical security, BMS	—	—	Zoning, permitting, community relations, land banked for future phases

8. Sources

Power, cooling, and heat reuse:

[PVB — Data Center BESS: UPS Integration, Backup Runtime, and Cost Guide](#)

[Data Center Knowledge — Why Data Centers Are Turning to Behind-the-Meter Power](#)

[EU Covenant of Mayors — Stockholm, Sweden: Heat Recovery from Data Centres](#)

[Nebius Newsroom — Nebius to Triple Capacity at Finland Data Center to 75 MW](#)

[Critical Grid Agility — off-grid renewable generation microgrids and electrostatic energy storage \(Newport Beach, CA\)](#)

Introl — SMR Nuclear Power for AI Data Centers: Feasibility and Implementation Timeline

<https://introl.com/blog/smr-nuclear-power-ai-data-centers-implementation>

Axis Intelligence — Nuclear Energy for Data Centers 2026: 9.8 GW Committed, SMRs by 2029 <https://axis-intelligence.com/nuclear-energy-for-data-centers/>

Campus-scale development and site selection:

[Area Development — 2026 Data Center Outlook: What Actually Changes Next](#)

[RSP Engineers — Data Center Campus Master Planning: A Civil Engineering Guide](#)

[Build.inc — Data Center Water Use Requirements: Cooling, Power and Community Risk](#)

[Build.inc — Data Center Transmission Planning in 2026](#)

P&W Solar — Solar Canopy and Ground-Mount Systems Drive Dell Data Center Expansion

<https://www.pandwsolar.com/solar-canopy-and-ground-mount-systems-drive-dell-data-center-expansion/>

US EPA — What Is RE-Powering America's Land? <https://www.epa.gov/re-powering/what-re-powering>

Cowboy State Daily — Wyoming Is Perfect for Data Centers, So Why Do They Cluster Only in Cheyenne?

<https://cowboystatedaily.com/2025/04/28/wyoming-is-perfect-for-data-centers-so-why-do-they-cluster-only-in-cheyenne/>

Landgate — Microsoft's 3,200 Acres in Cheyenne: Data Center Site Analysis

<https://www.landgate.com/news/microsoft-s-3-200-acres-in-cheyenne-data-center-site-analysis>

Data Centre Magazine — How Are Companies Pioneering Data Centre Zero-Water Cooling?

<https://datacentremagazine.com/news/how-are-companies-pioneering-data-centre-zero-water-cooling>

Redundancy and controls:

MEP Academy — Data Center Redundancy Explained: N, N+1, and 2N Systems

<https://mepacademy.com/data-center-redundancy-explained-n-n1-and-2n-systems/>

MEP Academy — How Data Centers Work: Power, Cooling, and Infrastructure Explained

<https://mepacademy.com/how-data-centers-work-power-cooling-and-infrastructure-explained/>

Course lecture, “Data Center Infrastructure” (mechanical & electrical) — recording and notes provided by the project owner; used to corroborate the redundancy and controls/monitoring framing above.

Physical security:

Asylon Robotics — Guardian Security Drone <https://asylonrobotics.com/solutions/guardiandrone/>

EAUAV — Solar Cleaning Solutions (C60V Drone, DJI T100 Kit, SolarSweeper)

<https://eauav.com/solutions/solar-cleaning/>

Rack density, speed to power, prefabrication, and medium/high-voltage distribution (Section 6):

“Your Data Center Is Now a Factory. Nobody Sent the Memo.” (Parts 1–4) — Drew Allen, CEO, Grace Technologies; slide deck provided by the project owner, primary source for all of Section 6.

Figures cited within that source, not independently re-verified here: Uptime Institute (2025 rack density survey), Enverus (interconnection queue and behind-the-meter generation data), Lawrence Berkeley National Laboratory (US interconnection queue capacity), IEA (global data center electricity demand), OCP “Mt. Diablo” ± 400 VDC specification v0.5 (Meta/Microsoft/Google).