ChatGPT

Attached Deliverables: Please find the following files for download: **(1)** `bibliography.csv` – a comprehensive spreadsheet of sources with claims and details, **(2)** `fact_cards.csv` – compact claim→fact→source mappings, and **(3)** `top_30_sources.md` – a markdown reference list of the 30 most relevant sources (with "why it matters" and supported claims).

Summary of Coverage, Gaps, and Conflicts:

Claim Category	Coverage (Sources & Alignment)	Gaps / Needs	Conflicts / Differences
Capacity (MW vs kW) – Hyperscale vs Edge vs Enterprise sizing.	*Covered by multiple sources:* JLL [1] and Dgtl Infra provide ranges (hyperscale 50–100+ MW vs edge ~50–500 kW; enterprise/colo in between). McKinsey (2025) reinforces hyperscale growth (200 MW now "normal"). All sources align that *hyperscalers operate at orders of magnitude larger scale* than typical enterprise or edge sites.	*No major gaps:* Could use more real-world examples of **enterprise** DC capacities (most sources generalize ranges). However, current coverage sufficiently establishes typical load profiles for each category.	**Definition variance:** JLL's definition of "edge" (up to 2 MW) [1] vs. Dgtl Infra's stricter 0.5 MW upper bound – indicating *industry terminology varies*. Overall message is consistent (edge << hyperscale).
Redundancy Models (N, N+1, 2N, 2N+1) – Uptime trade-offs.	*Well covered:* CoreSite and 123NET blogs define N vs N+1 vs 2N vs 2(N+1) with examples. All sources agree on what each level means (N+1 = one spare, 2N = full duplication, etc.) and their relative uptime (e.g. Tier IV ~99.995% vs Tier I ~99.671%). These independent sources (colocation providers and Uptime Institute data via CoreSite) **corroborate each other** on redundancy concepts.	*No significant gaps:* Covered definitions and examples. Perhaps could include more on *"distributed redundancy" (e.g. 3N/2)*, but this is beyond the main scope. Overall, redundancy tiers and uptime implications are documented with multiple sources.	**No direct conflicts:** Sources are in agreement. Only nuance: 123NET/ others note N+1 still has common-circuit risk. This nuance complements (doesn't contradict) the general definitions. All sources consistently differentiate *parallel N+1* vs *fully fault-tolerant 2N*.

Claim Category	Coverage (Sources & Alignment)	Gaps / Needs	Conflicts / Differences
UPS Topologies – Double-conversion vs Rotary (DRUPS); Battery types.	*Extensive coverage:* DCD article clearly explains static **double-conversion UPS** vs **rotary UPS** (motors/flywheels). CoreSite (Hatzenbuehler) and Schneider data cover **Lithium-ion vs VRLA batteries**, aligning on Li-ion's ~3× longer life and faster recharge. Sources uniformly highlight that *modern data centers predominantly use double-conversion UPS with VRLA, but Li-ion adoption is rising*.	*Minor gaps:* Could use more vendor-neutral data on **operational experience with Li-ion UPS** (most info comes from vendor blogs). Also, limited neutral data on **DRUPS reliability** beyond expert commentary (e.g., no broad stats on DRUPS vs static UPS failure rates). These areas might be bolstered by future case studies.	**No major conflicts:** All sources agree on technical facts (e.g., Li-ion ~15-year life vs VRLA ~5-year). A slight difference in emphasis: DCD notes rotary UPS are a *niche mostly in Europe*, which doesn't conflict but adds context. Overall, the narrative is consistent: *double-conversion UPS + batteries is standard; rotary/DRUPS are specialized*.
Generator Systems – Diesel, Gas, Hydrogen; Utility feeds.	*Well covered:* Multiple sources (Rolls-Royce mtu via DCD [2], DCF on Generac) confirm **diesel generators** as the dominant backup (2–3 MW units common) and note emerging **natural gas options**. Microsoft's report on a **3 MW hydrogen fuel cell** trial provides a forward-looking example. EPI/Uptime sources cover **utility feed redundancy** (Tier IV requiring two substations). Overall, sources agree that *diesel gensets with N+1 or 2N are standard, with gas/ hydrogen just starting to appear*.	*Gaps:* Lack of independent field data on **natural gas generator performance** in data centers (e.g., latency, reliability vs diesel). Also, hydrogen fuel cell adoption is nascent – we only have a pilot case (Microsoft) – so no neutral stats on reliability/cost for that yet. These are noted as future areas to watch ("no neutral source on hydrogen genset adoption yet").	**No direct contradictions:** All sources acknowledge diesel as current standard. Some variation in **emphasis**: e.g., vendor sources tout gas gensets as viable now, while others note gas's slower ramp as a drawback (implied in Rolls-Royce piece). These aren't outright conflicts – just different focus. There's also a *practical vs theoretical gap*: hydrogen is praised for zero-emission, but **Uptime**/ industry reports caution it's not yet widely adopted (no disagreement, just reality vs future hope).

Claim Category	Coverage (Sources & Alignment)	Gaps / Needs	Conflicts / Differences
Efficiency & PUE – Energy efficiency benchmarks.	*Strong coverage:* Uptime Institute data (via DCK and Uptime's own report) shows **average PUE ~1.55** in recent years – indicating industry-wide efficiency has plateaued. Vertiv and others highlight the **PUE gap between hyperscalers (~1.1) and typical enterprise (~1.7)**. All sources align that *lower PUE is better and hyperscalers are leading, while the global average is higher*. We have both the broad survey data and specific provider examples.	*Gaps:* Could include more segmented data (e.g., PUE by region or by facility size). Also, no source in this pack specifically gives **PUE by tier** or by cooling type – although we have enough to infer trends. In general, current sources cover the main points; detailed breakdowns (if needed) might come from future DOE/EPA reports.	**Conflicts:** Mainly in the form of *claims vs reality*: Hyperscalers (Google/ Facebook) **claim ~1.1 PUE**, whereas the **industry average remains ~1.5+**. This isn't a direct disagreement between sources, but a contrast between *best-in-class vs typical*. It underscores potential skepticism on provider claims – some colos advertise PUE ~1.2, but Uptime data suggests most run higher. No conflicting numbers on the same statistic from two credible sources were found in our research.

1 Why smaller data centers are taking off
https://www.jll.com/en-us/insights/why-smaller-data-centers-are-taking-off

2 Balancing act: How demand for data centers can square with environmental impact - DCD
https://www.datacenterdynamics.com/en/opinions/balancing-act-how-demand-for-data-centers-can-square-with-environmental-impact/