

# Research Gaps and Evidence Map

Eight gaps. Eight opportunities. A graded assessment of what the literature supports and what remains to be established.

*"Gaps where the literature is thin or absent are not weaknesses to conceal.  
They are research territory to own."*

<b>Series</b>	Synestra Research Intelligence Series   Vol. III   Part III of III
<b>Follows</b>	The Hidden Cost of Operational Hyperscale and The Economic Recovery Case
<b>Date</b>	June 2026
<b>Status</b>	Public — For Distribution to Industry Analysts and Investors

# PURPOSE

---

## An honest accounting of the evidence.

This brief maps eight research gaps that are material to Synestra’s commercial and intellectual positioning. For each gap, I document: what the gap is and why it matters to the thesis, the strongest existing evidence from peer-reviewed literature, national laboratory research, and tier-one consulting analysis, an honest evidence strength rating, and the Synestra module the research validates.

Gaps where the literature is thin or absent are not weaknesses to conceal. They are research opportunities to own. The first company to publish credible analysis in an underdocumented area controls the framing for every investor conversation, partnership discussion, and competitive comparison that follows.

Five of the eight gaps have strong to moderate existing evidence that Synestra can cite today. Three gaps represent genuine first-mover research opportunities where no peer-reviewed literature currently exists. Those three are where the intellectual high ground is available and unclaimed.

Evidence Strength	Description
<b>STRONG</b>	Peer-reviewed production data or multiple independent sources confirming the finding
<b>MODERATE</b>	Supporting evidence exists but gap-specific empirical data is limited or indirect
<b>THIN</b>	Market context exists but no peer-reviewed research addresses the specific claim

## Workload to Infrastructure Coupling

Evidence Strength: STRONG

Validates: Resource Intelligence

### Why It Matters

Synestra's most distinctive claim is that workloads actively determine infrastructure economics, not just consume power. AI training, inference, and other workload types create measurably different thermal signatures, power demand profiles, and facility-level economic outcomes. If true, workload awareness is the prerequisite for campus-scale coordination. The research on this is stronger than most assume.

### What the Research Shows

Chen, Wang, Colacelli et al. (Texas A&M; / Harvard, arXiv:2509.07218, November 2025) provide the most comprehensive review to date. AI data center electricity demand is characterized across four distinct stages: model preparation, training, fine-tuning, and inference, each with fundamentally different load profiles, temporal variability patterns, and cooling implications. This is a peer-reviewed, multi-institution confirmation that workload type is a primary driver of infrastructure behavior, not an incidental variable.

Latif, Newkirk, Carbone et al. (Brookhaven National Laboratory / LBNL, arXiv:2412.08602, IEEE Access 2025) provide the critical empirical data point. On an 8-GPU NVIDIA H100 HGX node, the maximum observed power draw during real AI training workloads was 8.4 kilowatts: 18 percent lower than the manufacturer-rated 10.2 kilowatts, even at near-full GPU utilization. Holding model architecture constant and increasing batch size from 512 to 4,096 images reduced total training energy by a factor of four. A scheduling decision changed energy consumption 4x on identical hardware producing identical output.

NVIDIA (Developer Blog, March 2026) frames the commercial consequence: AI factories operate within hard power envelopes, and performance per watt is the defining metric. NVIDIA's DSX orchestration software explicitly co-optimizes workloads, power, and cooling together. This is the largest GPU vendor publicly validating that workload-infrastructure coupling is the architecture of the AI factory era.

### The Gap That Remains

*No peer-reviewed study has quantified, at production campus scale, the economic delta between coordinated and uncoordinated workload-infrastructure scheduling across all subsystems simultaneously. Synestra can own this research space.*

Key Citations: Chen et al., arXiv:2509.07218, Texas A&M; / Harvard, Nov 2025. Latif et al., arXiv:2412.08602 / IEEE Access 2025, Brookhaven NL / LBNL. NVIDIA, Developer Blog, March 2026. McKinsey, Dec 2025.

## Economic Value of Earlier Commissioning

Evidence Strength: STRONG

Validates: Operations Management

### Why It Matters

If a 2.5 GW campus is delayed by one week, what does it cost? That is not a rhetorical question. It is the ROI anchor for Synestra's Operations Management module. If commissioning is the primary source of delay and delay can be quantified, the value of faster commissioning is directly calculable. The research here is clean.

### What the Research Shows

STL Partners (March 2025) provides the cleanest number: delays in data center construction cost developers an average of \$14.2 million per month for a typical 60 megawatt facility. The IRR impact is decisive. A project delivering on time achieves 17.1 percent IRR. A one-month delay drops this to 15.5 percent. A three-month delay collapses it to 12.6 percent: a 26 percent IRR deterioration. The root cause is fragmented, manual reporting that creates a lag between issues arising and being detected.

Exto's hyperscale commissioning analysis adds the granular breakdown. A three-month delay on a 50 megawatt hyperscale campus produces \$15 to \$40 million in financing charges, \$10 to \$60 million in deferred revenue, and \$5 to \$30 million in liquidated damages: a total exposure of \$30 to \$150 million per three-month slip. The root cause is 15 to 25 contractors working across mechanical, electrical, controls, and life safety systems with no unified coordination platform.

Scaling to 2.5 GW: applying STL Partners' \$14.2 million per month per 60 megawatts linearly, a 2.5 GW campus delay costs approximately \$590 million per month in direct economic losses, excluding IRR compression on the capital stack.

### The Gap That Remains

*No published study has modeled commissioning delay costs at GW scale, or the specific commissioning complexity increase from multi-building, multi-phase campus architectures. Synestra can develop and publish this model.*

*Key Citations: STL Partners, March 2025 (\$14.2M/month; IRR 17.1% to 12.6%). Exto, The Hidden Cost That Keeps Data Center Executives Awake at Night (\$30 to \$150M per 3-month delay). PEAK Technical Staffing, Sept 2025.*

## Operational Knowledge Loss

Evidence Strength: MODERATE

Validates: Lifecycle Analytics

### Why It Matters

Every data center runs on institutional knowledge: the mental models that experienced engineers carry about how systems interact, why certain configurations were made, and how to diagnose anomalies that sensors cannot explain. When those engineers leave, that knowledge disappears. The question is whether this loss is quantifiable and severe enough to justify a dedicated knowledge retention system.

### What the Research Shows

Hong, J.H. (Oxford Academic, June 2025) provides the most directly applicable study. Across 94 US federal agencies managing IT projects from 2015 to 2021, employee turnover consistently caused project delays. Centralized authority structures and task standardization significantly buffered this effect. The key principle: organizational structures that retain task knowledge mitigate the operational damage of turnover. Synestra's Lifecycle Analytics module is precisely such a structure, applied to data center operational knowledge.

The Eurasia Review synthesis on workforce turnover and organizational memory identifies the mechanism: it is not just technical skill loss, but the fragmentation of knowledge networks that underpin collaborative decision-making. In a hyperscale data center, this maps directly to cross-domain coordination knowledge: the senior engineer who knows that Hall B cooling behaves oddly under mixed training and inference load in summer. When that engineer leaves, that pattern has to be rediscovered the hard way.

What the literature does not yet show: there is no peer-reviewed quantification of knowledge loss cost specifically in data center operations. No MTTR degradation study from staff turnover. No incident recurrence rates from lost institutional memory. No operational cost of relearning system behavior after team changes. The data center industry has not studied itself through this lens.

### The Gap That Remains

*No peer-reviewed study quantifies MTTR degradation from data center staff turnover, incident recurrence rates from lost institutional memory, or relearning costs after team transitions. The anecdotal evidence is overwhelming. The academic framework exists. The data center-specific empirical study does not. First to publish owns the narrative.*

*Key Citations: Hong, J.H., JOPART, Oxford Academic, June 2025. Eurasia Review, The Silent Collapse: Workforce Turnover, Organizational Memory, and the Fragmentation of Knowledge Networks.*

## Multi-System Root Cause Analysis

Evidence Strength: STRONG

Validates: Core Coordination Architecture

### Why It Matters

Today, SCADA sees power. BMS sees cooling. DCIM sees infrastructure. Network monitoring sees packets. No system sees Power to Cooling to Compute to Workload to Economics as a unified causal chain. If most incidents span multiple domains and root cause is frequently upstream from visible symptoms, then siloed monitoring systems are architecturally incapable of correct diagnosis. Coordination is not optional. It is the only path to correct root cause.

### What the Research Shows

IoT Now (March 2026) names the problem precisely: the system zoo. Specialized tools, each authoritative in its own territory, none speaking to the others. Under stress, this architecture breaks down. When a circuit breaker trips, downstream effects hit engineering, servers, and network simultaneously. Each monitoring system generates its own alert stream. Response time comes down to how long one engineer needs to piece together a timeline across four or five consoles.

The article states directly: power consumption, thermal load, and server workload are tightly coupled. Shifts in one show up in the others, often within seconds. Without a view that ties power draw, inlet temperature, and server utilization together, answering a root cause question takes minutes. In a degrading situation, those minutes matter. This is a practitioner-facing technical publication describing the exact problem Synestra solves, published March 2026.

LEMMA-RCA (Zheng et al., UIUC / NEC Laboratories / Rutgers, 2024) provides the academic foundation: a large-scale, multi-modal, multi-domain dataset for root cause analysis across IT and OT operational systems. The paper's framing is directly applicable: recent advancements have led to increased complexity and interdependence in modern systems. This complexity heightens their vulnerability to faults arising from interactions among modular services.

### The Gap That Remains

*No study has quantified, at data center scale, the percentage of incidents that are multi-domain in origin, the average MTTR penalty from siloed diagnosis, or the revenue impact of extended resolution time on production AI workloads. These data points would make the Synestra RCA argument decisive.*

*Key Citations: IoT Now, Why DCIM Still Fails When Data Centres Need It Most, March 2026. Zheng et al., LEMMA-RCA, arXiv:2406.05375, UIUC / NEC Laboratories, 2024.*

## AI Factory Economics

Evidence Strength: STRONG (revenue metrics) / MODERATE (GPU throttle cost)

Validates: [The Synestra Value Reframe](#)

### Why It Matters

Synestra's value proposition is not save cooling energy. It is increase revenue output per megawatt. To make this case, the research must establish the revenue density of AI factories: what a megawatt of AI compute is worth, what a GPU-hour costs when idle or throttled, and what the economic consequence of coordination failure looks like in production terms.

### What the Research Shows

NVIDIA (Developer Blog, March 2026) introduces the token factory model: AI data centers as revenue-per-watt engines, where performance per watt is the defining metric. NVIDIA documents a 1,000,000x improvement in inference throughput per megawatt across six GPU generations, framing MW productivity as the primary competitive variable. This establishes the framework in which Synestra's value is measured: not PUE, but revenue per megawatt.

Market data on GPU economics is clear. H100 GPUs command \$25,000 to \$40,000 per unit. GB200 NVL72 rack-scale systems command \$1.8 million per unit. NVIDIA data center revenue reached \$47.5 billion in Q4 FY2025: 461 percent year-over-year growth. At \$12.5 million per megawatt per year in AI factory revenue density, every percentage point of Economic Availability gap on a 2.5 GW campus costs \$3.1 billion annually in unrealized output.

What the research does not yet contain: no published study quantifies the economic cost of GPU throttling specifically caused by infrastructure coordination failure. The throttling data exists. The infrastructure causation has not been formally connected to revenue loss in the literature.

### The Gap That Remains

*The single highest-value research investment Synestra can make is a production study quantifying the economic cost of GPU throttling events caused by infrastructure coordination failure. This study does not exist anywhere in the literature. It would simultaneously validate Gap 01, Gap 04, and Gap 05 in a single publishable result. It is the research that makes the investment thesis undeniable.*

Key Citations: NVIDIA, *Scaling Token Factory Revenue and AI Efficiency by Maximizing Performance per Watt*, Developer Blog, March 2026. Luminary Signal, *GB200 NVL72 analysis*, April 2026.

## Landlord Economics

Evidence Strength: STRONG

Validates: Asset Performance Platform

### Why It Matters

Data center real estate has become the highest-performing asset class in commercial real estate. The question for Synestra is whether operational coordination software can be positioned as an asset performance tool that landlords pay for independently of tenants. The research here is strong.

### What the Research Shows

SitusAMC's 2025 analysis documents the asset class performance: 21 percent year-over-year FFO growth and 7.2 percent NOI gains. Data center REITs command a 5.0 percent cap rate: the lowest of any commercial real estate sector. In concrete terms, a 1 percent NOI improvement on a 60 megawatt colocation facility generating \$30 million per year in NOI translates directly to \$6 million in asset value at a 5.0 percent cap rate.

Transaction volume confirms the scale of the market: \$52 billion in data center transactions in 2024, plus \$13 billion in the first eight months of 2025 alone. The asset class is being actively priced. Any performance delta attributable to coordination software is now economically material at transaction scale.

Nareit data center REIT performance data for Q1 2025 confirms: 21 percent FFO growth, 7.2 percent NOI gain, 5.0 percent cap rate. These are not projections. They are reported results from the public REIT market.

### The Gap That Remains

*No published study has isolated the asset value premium attributable to operational coordination software, or modeled the NOI uplift from stranded capacity monetization at the REIT portfolio level. This is the analysis Synestra should produce with a real estate investment bank or REIT-focused advisory firm.*

*Key Citations: SitusAMC, Data Centers in an AI-Driven Era: Key Trends Reshaping Valuation, 2025. Nareit, Data Center REIT Performance Data, Q1 2025. Preqin, Data Center Transaction Data, 2024 to 2025.*

## OEM Benefits

Evidence Strength: THIN — Confirmed Research Gap

Validates: OEM Partnership Thesis

### Why It Matters

Why would Schneider Electric, Vertiv, Eaton, Cummins, GE Vernova, or NVIDIA care about Synestra? The intuitive answer is that real-world operational telemetry from Synestra's coordination layer provides OEMs with failure pattern data, product feedback loops, and warranty cost reduction intelligence they cannot otherwise access. This is compelling. The research basis is thin.

### What the Research Shows

The OEM market context is clear. Vertiv commands 23 percent global market share in precision cooling with \$6.9 billion revenue. Schneider Electric generated \$34.2 billion in revenue with data center infrastructure as a significant segment. The liquid cooling market is projected to grow from \$5.52 billion to \$15.75 billion by 2030. These are companies for whom even a 1 to 2 percent reduction in warranty and field service costs represents hundreds of millions of dollars.

However, no published peer-reviewed research quantifies data center OEM warranty expense rates, the cost of field failure root cause investigation, or the economic value of operational feedback loops. Industry sources confirm this is considered proprietary and competitive-sensitive by all major OEMs. The academic framework exists in manufacturing literature. It has never been applied to data center critical infrastructure at scale.

### The Gap That Remains

*There is no publicly available research quantifying the cost of data center OEM warranty claims, failure analysis investigation expense, or the economic value to OEMs of production operational telemetry access. First publisher controls the narrative for this segment of the Synestra thesis. Synestra can produce and publish the first structured analysis of OEM operational data value, potentially in partnership with one anchor OEM willing to pilot a feedback loop arrangement.*

*Key Citations: Introl Blog, Vertiv vs Schneider vs Eaton: Cooling Solutions Comparison for AI Data Centers, March 2026. Liquid cooling market: \$5.52B to \$15.75B by 2030 (industry consensus).*

## What Happens at 2.5 GW?

Evidence Strength: MODERATE

Validates: Synestra's Category-Creation Argument

### Why It Matters

The most powerful argument Synestra can make is not efficiency. It is that existing operating models were not designed for GW-scale AI campuses, cannot scale to them, and will fail structurally as the industry makes this transition. This is a category-creation argument. It positions Synestra not as a better DCIM, but as a new infrastructure category required by the new scale of AI compute.

### What the Research Shows

McKinsey (August 2025) explicitly acknowledges the scale transition: data center campuses must expand from providing tens of megawatts to hundreds, even expanding to accommodate gigawatt scale. Global power demand for data centers is projected to reach 1,400 TWh by 2030: 4 percent of total global power demand, with US capacity growing from 30 GW in 2025 to 90+ GW by 2030, a 22 percent CAGR. The research identifies this transition as requiring a reevaluation of design and construction methods, but stops short of analyzing operational management complexity at GW scale.

This is the gap. The entire research literature was developed for data centers measured in tens of megawatts. At 2.5 GW, representing 40 to 60 individual data center halls, SCADA, BMS, DCIM, and ITSM tooling generates alert volumes, coordination requirements, and interdependency graphs that fundamentally exceed their design parameters.

Consider what 100 racks actually means operationally. At 800VDC bus systems, 100 racks generates thousands of telemetry data points per second across per-rack power, busway load, coolant temperatures, CDU health, GPU thermals, and NVLink fabric status. None of these systems were designed to correlate with each other. A NOC team can see events. It cannot see the causal chain forming before an alarm fires. At 1,500 racks per building approaching one gigawatt of power draw, this is operationally dangerous. The industry is building at this scale right now.

### The Gap That Remains

*No study exists modeling operational complexity, staffing requirements, or tool failure modes at GW scale. The absence of published research is itself evidence: the industry has not yet studied a problem it is actively creating. First to publish creates the category framework that all subsequent analysis, including competitor and investor research, will reference.*

*Key Citations: McKinsey, Building Data Centers Bigger, Faster, August 2025. McKinsey, The Next Big Shifts in AI Workloads and Hyperscaler Strategies, December 2025. Data Center Frontier, February 2026.*

# SYNTHESIS

## Three gaps to own. Five gaps to cite.

Of the eight gaps, three represent genuine first-mover research opportunities where no peer-reviewed literature currently exists: Operational Knowledge Loss in data centers (Gap 03), OEM Operational Data Value (Gap 07), and GW-Scale Operating Model Failure (Gap 08). These are not weaknesses in the thesis. They are intellectual territory that Synestra can claim through primary research, white paper publication, or academic partnership. First to publish owns the framing.

Five gaps have strong to moderate existing evidence that Synestra can cite immediately. Gaps 01, 02, 04, 05, and 06 are supported by peer-reviewed papers, national laboratory reports, tier-one consulting analysis, and production operator data. The evidence is sufficient to build investor materials, customer-facing briefings, and product roadmap justifications today.

Gap	Topic	Evidence	Research Priority
01	Workload to Infrastructure Coupling	STRONG	Highest
02	Commissioning Economics	STRONG	High
03	Knowledge Loss	MODERATE	High
04	Multi-System RCA	STRONG	High
05	AI Factory Economics	STRONG / MODERATE	Highest
06	Landlord Economics	STRONG	Medium
07	OEM Benefits	THIN	Medium
08	What Happens at 2.5 GW	MODERATE	Highest

### THE HIGHEST-VALUE RESEARCH INVESTMENT

The single most important study Synestra can commission is a production quantification of the economic cost of GPU throttling events caused by infrastructure coordination failure: connecting workload, power, cooling, and revenue into a single measured data point at campus scale.

**This study does not exist anywhere in the literature. It would simultaneously validate Gap 01, Gap 04, and Gap 05 in a single publishable result. It is the research that makes the investment thesis undeniable.**

*Part of a three-brief series. See also: [The Hidden Cost of Operational Hyperscale](#) and [The Economic Recovery Case](#). | [synestra.ai](#) | [john.chavner@synestra.ai](mailto:john.chavner@synestra.ai)*