EVIDENCE MAP · JUNE 2026

Research Gaps and Evidence Map

Where the evidence is strong, where it is emerging, and where the territory is open.

ABSTRACT

The Synestria thesis is built on published evidence. Where that evidence is strong, Synestria stands on it directly. Where it is emerging, Synestria extends it with a clearly labeled thesis. Where it is absent, Synestria identifies the gap and builds toward it. This paper maps the full evidence landscape across eight research areas: Economic Availability as a metric, operational knowledge loss, the 24-month data moat, coordination layer value, cross-domain consequence tracing, AI-native management platform design, benchmark methodology for EA measurement, and cross-campus knowledge inheritance. Transparency about where the evidence ends and the thesis begins is not a weakness. It is how durable positions are built.

PUBLICATION	Synestria Research Series | White Paper No. 4	**DATE**	June 2026
STATUS	Public — For Distribution	**CONTACT**	research@synestria.com

The strongest positions know exactly where their evidence stops.

Investors, operators, and technical advisors who engage with an early-stage infrastructure thesis should expect a clear map of what is documented and what is not. This paper provides that map. It is not a defense of positions under pressure. It is an honest account of where the evidence is strong, where it is emerging, and where the open territory represents both risk and opportunity.

The evidence categories used in this map carry specific meanings. Strong evidence means independently replicated findings at production scale. Emerging evidence means documented findings that require extension to apply to the Synestria context. Open territory means an area where the mechanism is plausible but the data does not yet exist in published form.

> *Most of the core Synestria thesis is built on strong evidence. The emerging and open categories are not gaps in the argument. They are the territory where Synestria is actively building. Naming them is not a risk disclosure. It is a research agenda.*

Evidence categories used in this map.

Strong: Multiple independent studies at production scale. Findings replicated across organizations and geographies. Mechanism well understood.

Emerging: Documented findings at meaningful scale with reasonable generalizability. One or two authoritative sources. Mechanism clear but not exhaustively replicated.

Open Territory: Mechanism is plausible and consistent with strong adjacent evidence. No direct production data published in peer reviewed form. Synestria is building toward this data.

Eight research areas. Honest strength ratings.

RESEARCH AREA	STRENGTH	KEY SOURCES AND NOTES
Stranded Capacity and Utilization Gaps	Strong	LBNL 2024 confirms 30%+ comatose servers in production. IDC, Gartner, and multiple national lab studies corroborate. Mechanism well documented across facility types and generations.
Thermal Inefficiency and Cooling Optimization	Strong	Google DeepMind 2016 production deployment. 40% cooling energy reduction. Multiple subsequent hyperscale deployments. Mechanism replicated at scale. PUE measurement limitations documented independently.
Software Coordination Value	Strong	Google ASPLOS 2020: 25%+ oversubscription at tens of MW. BNL/LBNL 2025: 4x power swing from single parameter change. Results independently documented at production scale.
Physical Fabric and Networking Loss	Strong	Meta LLaMA 3: 419 interruptions in 54 days. ByteDance: 38,000+ failures in 3 months. Borrill 2026: ghost event every 48 seconds at 2025 cluster scale. Checkpoint tax 12 to 43% of training time. Multiple authoritative production sources.
Cross-Campus Knowledge Inheritance	Emerging	McKinsey 2025: $250B design cost recovery from operational data. Mechanism plausible and consistent with adjacent strong evidence. Direct measurement of cross-campus learning curves not yet published in peer reviewed form.
Economic Availability as a Formal Metric	Emerging	Component losses documented independently and strongly. EA as an integrating metric is a Synestria extension. The component evidence is strong. The aggregation methodology is a Synestria contribution to the field.
24-Month Operational Memory Moat	Emerging	Google DeepMind training data timeline implies 12-month minimum for predictive thermal models. No published study has directly measured the moat timeline for a multi-vector operational memory system. This is open research territory.
AI-Native Infrastructure Management Platform	Open Territory	No published peer reviewed evidence for a cross-domain AI-native platform at production hyperscale. All adjacent evidence is strong. The platform category itself is a new territory. Synestria is the first entrant.

> *The four strong evidence areas — stranded capacity, thermal inefficiency, software coordination, and physical fabric loss — represent the foundation of the Synestria thesis. They are not disputed. They are documented at the largest computing operations in the world. The four emerging and open areas represent the territory where Synestria is building the next layer of evidence.*

Open territory is not empty. It is where the work happens.

The open research territory in this map is not a weakness in the Synestria thesis. It is a description of the frontier. Every strong evidence area in this map was open territory at some point. It became strong because someone built the data. Synestria is building the data for the open areas.

Economic Availability methodology

Synestria is developing the EA measurement methodology against production data from deployed campuses. The first published EA baseline established in under 30 days will be the foundational data point for validating EA as a standard metric. That publication is a near-term research output.

The 24-month data moat

The 12-month threshold for predictive thermal model maturity is inferred from the Google DeepMind deployment timeline. Synestria's deployment data will allow direct measurement of the threshold across all four learning vectors. This is the most important open research question in the Synestria program.

Cross-domain consequence tracing

No published study has traced a consequence chain from the physical fabric layer through the network layer, the compute layer, and to an economic output measurement in a single continuous trace. Synestria's Consequence Intelligence module is designed to produce exactly this trace. The first documented cross-domain consequence trace from a production deployment will close this gap.

AI-native platform design

The design principles for an AI-native infrastructure management platform do not yet exist in published form. Synestria's architecture is informed by the strong evidence from adjacent domains. The platform design itself is the contribution. The deployment results will produce the evidence base for this research area.

Synestria did not invent these problems. The industry has already documented them. Synestria is designed to measure, attribute, and recover their economic impact across the AI factory chain. The open territory above is not the foundation. The strong evidence is the foundation. The open territory is where the next five years of compounding advantage gets built.