SYNESTRIA

RESEARCH SERIES · WHITE PAPER NO. 1

WHITE PAPER · JUNE 2026

Operational Memory and Economic Availability

Why AI factories need infrastructure that learns.

ABSTRACT

Technical availability tells you the facility is running. Economic Availability tells you whether it is producing. Operational memory determines how much it improves. A Tier IV facility operating at 99.995% technical availability may simultaneously operate at 65 to 78% Economic Availability. That gap represents up to $6.6 billion in annual economic output lost to coordination failures that technical availability was never designed to detect. This paper establishes Economic Availability as the formal measurement framework, documents its five loss components using peer reviewed research and production data, and demonstrates that EA improvement is the primary evidence that operational memory creates compounding business value.

PUBLICATION	Synestria Research Series | White Paper No. 1	DATE	June 2026
STATUS	Public — For Distribution	CONTACT	research@synestria.com

The gap that monitoring cannot close.

AI factories are becoming too large for experience to remain trapped inside people. The accumulated knowledge of experienced engineers. How systems interact. How thermal sequences develop. Which interventions succeed. That knowledge disappears when engineers leave, transfer, or retire. The infrastructure forgets.

The problem is not that AI factories fail. Most do not. The problem is that they forget. Infrastructure that forgets cannot improve.

Operational memory is the systematic capture of what infrastructure learns during operation. Every intervention, every consequence, every outcome, every pattern. When infrastructure remembers, it improves. When it does not, it merely operates.

> *EA is not the thesis. EA is the evidence. The thesis is that infrastructure which remembers becomes infrastructure that improves. And that improvement compounds. Economic Availability is the metric that makes that improvement legible.*

Economic Availability: a formal definition.

Economic Availability is the ratio of a facility's actual annual economic output to its maximum possible annual economic output given its provisioned physical infrastructure. Technical availability measures whether systems are running. Economic Availability measures what they are producing.

99.995%	65–78%	$1.6B
TECHNICAL AVAILABILITY — TIER IV STANDARD	ACTUAL ECONOMIC AVAILABILITY AT THE SAME FACILITY	ANNUAL RECOVERABLE ECONOMIC OUTPUT AT 1.4GW FRONTIER SCALE — 3-POINT EA GAP, BASE SCENARIO

These numbers coexist without contradiction. A facility can satisfy every technical availability SLA and simultaneously convert only two thirds of its provisioned capacity into economic output. The other third disappears into coordination failures that no availability dashboard captures.

> *These coordination failures are not random. They are the predictable consequences of infrastructure that has no memory. They are reduced — measurably and progressively — as operational memory accumulates.*

Five components that close the EA gap.

Economic Availability loss is not uniformly distributed. It concentrates in five structural categories, each documented in peer reviewed research and each addressable through operational memory.

LOSS COMPONENT	SOURCE	EA IMPACT	MODULE
Stranded Capacity	LBNL 2024: 30%+ servers in comatose state	2–5 pts	Resource Intelligence
Thermal Inefficiency	Google DeepMind: 40% cooling reduction	1–3 pts	Thermal Intelligence
Coordination Loss	Google ASPLOS: 25%+ additional capacity	3–6 pts	Compute Intelligence
Knowledge Loss	McKinsey: $250B design recovery	2–4 pts	Lifecycle Analytics
Physical Fabric	Borrill 2026: fiber loss every 48 seconds	1–3 pts	Edge Data Intelligence

Why the losses compound

Each loss category does not operate independently. Stranded capacity increases the thermal burden on the active portion of the facility. Higher thermal load reduces power delivery efficiency. Reduced power efficiency increases scheduling uncertainty. Scheduling uncertainty degrades compute utilization. All five categories reduce EA simultaneously. Operational memory addresses all five.

> *A facility at 72% EA with Synestria deployed is not a facility with an 8-point problem. It is a facility with five concurrent, interacting losses. The value of resolving them compounds because the categories interact. Fixing thermal fixes coordination. Fixing stranded capacity fixes power. The second improvement is always larger than the first.*

What the literature establishes

The five loss components are documented independently across peer reviewed research and production deployments at Google, Meta, Lawrence Berkeley National Laboratory, ByteDance, and McKinsey. None of the source papers uses the term Economic Availability. Each measures a different loss category in isolation. Synestria is the first framework to aggregate them into a single metric and trace them to a common root cause: infrastructure without operational memory.

The compounding case for infrastructure that learns.

Operational memory does not produce a one-time improvement. It produces a compounding one. Every event captured makes the next event faster. Every campus observed makes the next campus smarter. The advantage is structural and it cannot be purchased by a competitor.

Year 1: Baseline and discovery

EA baseline established within 30 days of deployment. Five loss categories quantified individually. Recovery roadmap built from the facility's own history. The infrastructure sees itself clearly for the first time.

Year 2: Compounding returns

Operational patterns recognized from 12 months of history. Interventions validated against prior outcomes. Thermal pre-cooling deployed ahead of workload arrivals. The infrastructure acts on what it has learned.

Year 5: Durable moat

Five years of operational history. Facility specific consequence patterns that cannot be transferred to a competitor. Every engineering decision informed by real outcomes from this specific campus. The advantage is built into the facility itself.

The non-transferable asset

A competitor can buy the same hardware. They can hire the same engineers. They cannot purchase five years of operational memory built on your specific infrastructure, your specific workloads, and your specific failure history.

What this means at 1GW scale.

$456M	$543M	$630M
ANNUAL RECOVERY — CONSERVATIVE (3 PTS EA, 5.6 DAY PAYBACK)	ANNUAL RECOVERY — BASE CASE (4 PTS EA, 4.7 DAY PAYBACK)	ANNUAL RECOVERY — OPTIMISTIC (5 PTS EA, 4.1 DAY PAYBACK)

At $20 to $38 million per megawatt per year in recoverable value intensity — the conservative to base scenario in the Synestria commercial model — every percentage point of EA recovered at a 1GW facility is worth $200 to $380 million annually. The infrastructure investment to close that gap is zero. The data is already being generated. The only missing element is a system designed to read it.

> *Economic Availability is not a new metric that requires new infrastructure. It is a new lens on the data that your facility is already generating. The gap between what your infrastructure could produce and what it actually produces is visible the moment you know where to look.*