

## Unlocking 40–50% More Capacity in Existing Data Centers

This is a bold claim. These are big numbers. But we've done the work, we have the receipts, and we believe it.

When you combine four things — a thorough [data center assessment](#), optimized [power infrastructure](#), modern [liquid cooling](#), and disciplined [operational efficiency](#) practices — many legacy environments can realistically unlock **40–50% more usable power and floor space** for AI workloads.

Not 40–50% better efficiency in one narrow slice. Not just a reduction in power cost. But 40–50% more actual, deployable compute capacity inside the four walls you already own.

That is a very big deal for organizations hitting hard limits on power delivery or facing expensive utility upgrades.

### Important Reality Check

It's going to cost money. You will have to open your wallet for CapEx to get these OpEx savings and capacity gains. There is no free lunch here.

We should also be transparent: there isn't one clean case study where a single large customer did all of this at once — assessment, power optimization, liquid cooling, and deep operational changes in one coordinated effort. Too bad, because that would have made our job easier. Instead, we're pulling together results and observations from many different environments and efforts. The 40–50% range represents what becomes possible when these pieces are attacked systematically rather than piecemeal.

### How the Gains Compound

- **Power + Cooling** upgrades alone often deliver a 30–35% reduction in total facility power draw for the same AI workload.
- A proper **data center assessment** surfaces ghost systems, stranded capacity, and low-utilization assets that quietly consume power and cooling.

- Strong **operational efficiency** practices (workload profiling, better scheduling, rightsizing, and decommissioning idle gear) squeeze even more waste out of the system.

In facilities with significant legacy air cooling and inefficient power chains, these layers reinforce each other. The whole becomes meaningfully larger than the sum of the parts.

We're not saying every site will hit 50%. Many won't. But a meaningful number of facilities with high legacy inefficiency can realistically land in the 40%+ range if they commit to doing the work properly.

### **Why This Matters Strategically**

For some organizations, this difference is the line between “we can handle our AI roadmap in our current facilities” and “we need to build something new or pay someone else to host it.”

That's not a trivial distinction when capital is expensive and utility upgrades can take years — if they're possible at all.

### **Final Caveats**

Results vary dramatically depending on the starting condition of your environment. Half-measures will get you half the gains. This is not magic — it's just systematic attention to waste that most enterprise data centers have historically ignored.

We'll continue watching real deployments closely. Early indications are encouraging.