Visual One Intelligence

AI Cost Governance
for the Post-Merger Enterprise

Two estates. One cost model. Every AI dollar owned.

Visual One Intelligence®, the FinOps Certified Platform from Clear Technologies, gives IT leaders one view of infrastructure cost across cloud, storage, and compute. Our Hybrid FinOps™ approach puts cloud and on-prem spend in the same reports, priced the same way, and pairs the platform with **Capacity Planning as a Service**: analyst-led reporting and consulting, so your team gets a working cost governance model rather than another dashboard to staff. As AI workloads move from pilots to production, the same discipline extends to the infrastructure they run on.

The Challenge

AI spend arrives faster than the operating model built to govern it. GPU purchases, cloud AI services, and token-metered APIs land on top of an IT estate that, after a merger, is really two estates: different cloud providers, different tagging conventions, and monitoring each team built for the environment it knew best. That tooling still does its job; what changed is everything now landing around it. Finance asks what AI is returning, and in most enterprises nobody can answer: 52% of organizations report no clear owner for AI cost (Harness, State of AI in FinOps, 2026). The gap is not effort or tooling. It is that ITFM discipline, chargeback, showback, budget-versus-actual, and defensible TCO, was never extended to the infrastructure AI actually runs on.

AI Cost Governance Use Cases

1 One cost model across two merging estates

Situation Post-merger integration means consolidating IT environments that grew up separately: one AWS-centered, one Azure-centered, each with its own on-prem footprint, tagging conventions, and reporting stack. Nobody's numbers match, and AI spend is scattered across all of it.

How VOI helps Vendor-neutral reporting puts both estates on one screen with one tag model across cloud and on-prem, independent of any provider's or MSP's own tooling. CapEx and OpEx in one language, priced the same way.

Outcome A Fortune 500 company's infrastructure partner saved over $1M and rebuilt trust across teams using VOI as an independent source of truth. The same neutrality is what lets two merged IT organizations agree on a single set of numbers.

2 Chargeback and showback for AI spend

Situation Data science, engineering, and business units all consume AI infrastructure, but the cost sits in one central IT budget. "Who pays for AI" has no answer, so IT absorbs spend it does not control.

How VOI helps Chargeback and showback reporting by business unit, owner, or cost center, using one tag model across cloud and on-prem, applied to AI infrastructure the same way it is applied to storage and compute today.

Outcome Each team sees what it consumes. AI cost becomes an allocated, owned line item instead of an unexplained total.

3 Defensible AI budgets

Situation Budget planning season asks a hard question: what did AI cost this year, and what will it cost next year? Homegrown monitoring can show usage in one cloud, but it cannot produce a budget-versus-actual view finance will accept.

How VOI helps Budget-versus-actual modeling, spend trends, and forecasting across the whole hybrid environment, with amortization and asset-level financial reporting ready to export for reconciliation with the general ledger.

Outcome One airline found 730TB of hidden capacity and avoided $1.5M in planned purchases. Budgets defended with data behind every number, and purchase requests answered before they become emergencies.

> **"Our manager speaks to the business about what Visual One would report, and they trust what he says while he trusts what Visual One says."**
>
> Stuart Way, McLane

4 An operating model, not another dashboard

Situation The tooling question is downstream of a harder one: who owns AI cost, on what cadence, with what reports? Most organizations have monitoring. Few have a governance operating model.

How VOI helps Capacity Planning as a Service pairs the platform with quarterly analyst-led reviews and monthly consulting calls. Our analysts work alongside your FinOps and infrastructure teams to stand up the allocation model, the review cadence, and the reporting finance actually uses, complementing the monitoring you have already built rather than replacing it.

Outcome A long-term AI cost governance model with named owners and a running cadence, operational in weeks and staffed by analysts rather than by another headcount request.

On the Horizon

VOI is extending the platform to measure AI infrastructure at the source. A GPU telemetry collector built on NVIDIA Data Center GPU Manager (DCGM) captures utilization directly from the hardware, with GPU utilization and cost reporting, and AI token usage tracking across major AI providers, now in development. These capabilities are roadmap items and their scope and timing are still being finalized.

We are actively looking for early adopters to join beta testing for these functionalities. Beta participants work directly with the Clear Technologies and Visual One Intelligence® product and analyst teams and help shape how GPU and AI cost reporting takes form.

Why Visual One Intelligence®

- One view of cloud and on-prem infrastructure cost across storage, compute, and virtualization
- Fully loaded on-prem pricing, not just cloud bills
- Analyst-led capacity planning: quarterly reports and monthly consulting calls, not just software
- FinOps Certified Platform, SOC 2 Type II certified
- Built and delivered by Clear Technologies, so the platform arrives with the team that implements and runs it