

The Resource Model Is Dying.

Outcome-Based
Model Is Already Here.

Created by Techwave's Cloud Competency Team



Why the shift is happening, what is driving it, and how
enterprise leaders can get ahead of it.



Most Enterprises Are Optimizing a Model That No Longer Fits the Market


When you look at your cloud and infrastructure managed services spend, are you buying compute, storage, and engineer hours or resilience, speed, and measurable outcomes? For most enterprises, the answer is the first, and the gap between what you pay for and what your business needs is becoming a significant strategic liability.

This point of view argues that the shift from resource-based to outcome-based managed services is not a future consideration but a current reality. Across Techwave's four pillars (cloud migration, managed operations, cost optimization, and platform modernization), the pattern is clear: enterprises that restructure commercial models around outcomes outperform those that buy by the hour.

The pages that follow explain why this shift is happening, what is driving it, where the traps are, and how to prepare for it. The data is real, the patterns are observable, and the window to act is narrowing.

Powered by NALORA, Techwave brings cloud strategy, intelligent delivery, and business outcomes together through a unified, governed AI-enabled operating model.

The managed infrastructure services market is valued at \$129.14 billion in 2025 and is forecast to reach \$210 billion by 2030. The growth is not coming from more of the same. It is coming from enterprises demanding a fundamentally different value proposition. (Mordor Intelligence)





How the Resource Model Was Built, and Why It Is Showing Its Age

The resource model was not a bad idea. It was a rational response to its time. Enterprises needed predictability. Providers needed stable margins. So they built a commercial framework around inputs: engineers, infrastructure, and hours of managed support.

The familiar delivery pattern followed: fixed compute and storage across multi-year cycles, monthly billing for a defined team regardless of deliverables, migrations scoped by effort, and value reported through uptime percentages and ticket resolution rates.

The model rewarded one thing above all else: resource deployment. More infrastructure meant a larger contract, which meant more personnel, with progress measured by renewal, not transformation.

None of this was dishonest. It was calibrated for a world where technology was a cost center, not a competitive differentiator. That world has changed. Cloud is now the engine of product delivery, customer experience, and organizational agility. And a commercial model built around resource consumption is structurally misaligned with those priorities.

One-third of enterprise technology leaders now say cost control is their most critical cloud focus, up 8 percentage points from the year prior. The resource model's inefficiencies are no longer invisible. (Flexera 2026 State of the Cloud Report)





Three Forces Are Making the Old Commercial Model Structurally Untenable

The transition away from resource-based managed services is not driven by a single trend. Three distinct forces are converging simultaneously, and their combined effect is accelerating the pace of change well beyond what most enterprise planning cycles anticipate.

AI Is Changing the Economics of Cloud Delivery

AI is changing the fundamental cost structure of running, monitoring, and optimizing cloud infrastructure. It predicts failures before they occur, resolves incidents without human escalation, and surfaces waste at a granularity no manual review can match.

The commercial implication is direct. The same outcomes that previously required a team of ten can now be delivered by a team of three, with better consistency and faster resolution times. A managed services partner still pricing for ten people is not a strategic partner; they are a structural cost.

30% boost in operational efficiency for enterprises using AI-enabled cloud services (Forrester, 2025)

Cloud Spend Is Growing Faster Than the Value It Produces

Enterprise cloud investment is accelerating at a rate that should concern every CTO and CFO in the room. Global cloud infrastructure services spending reached \$171.8 billion in 2024, up 22.5% year-over-year. IDC projects a further 33.3% increase in 2025, reaching \$271.5 billion. Yet a meaningful share of that spend is producing no measurable return.

29% of enterprise cloud spend is wasted in 2026, the highest level in recent years, driven by AI workload complexity and IaaS/PaaS proliferation (Flexera 2026)

Over 20% of organizations report having little to no visibility into what individual aspects of their business cost in the cloud. The resource model has no structural mechanism to address this, because its commercial logic is built around consumption, not optimization.

The Contract Structure of Managed Services Is Already Changing

Enterprises are not waiting for outcome-based models to mature before adopting them. They are already signing agreements that look fundamentally different from the retainer-based contracts of the previous decade. The shift is contractual and measurable.

41% of managed service deals signed in North America in H1 2025 incorporated variable billing tied to cloud resource usage or delivered outcomes (Future Market Insights)

MSPs are actively phasing out fixed-capacity retainers in favor of consumption-linked and outcome-indexed models. If your current managed services partner has not raised this conversation with you yet, that absence is itself informative.



What Changes When You Move From Buying Resources to Buying Outcomes

The outcome-based model is not a rebranding of the resource model with different pricing language. It represents a structural realignment of how managed services are scoped, delivered, and measured. The differences show up in every dimension of the relationship between an enterprise and its cloud partner.

Dimension	Resource-Based Model	Outcome-Based Model
What You Buy	Compute capacity, engineer hours, ticket handling	Uptime guarantees, cost reduction targets, migration milestones, optimization commitments
How It Is Priced	Fixed monthly retainer based on team size and provisioned infrastructure	Consumption-based or milestone-linked; billing tied to usage levels and delivered results
How Value Is Measured	SLA compliance: uptime %, mean time to resolution	Business metrics: cost-to-serve, deployment frequency, cloud waste reduced, resilience score
Who Bears the Risk	The enterprise: large team, poor outcomes, still billed in full	Shared: the provider is commercially incentivized by the same outcomes you care about
How AI Is Integrated	Add-on tooling, often billed as a separate line item	Embedded in delivery: AIOps, FinOps, and automated remediation are core service components
Migration Scope	Defined by effort estimates and project timelines	Defined by outcome targets: cost reduction, performance improvement, business continuity metrics
Governance Model	Reactive: spend reports generated after cloud costs are already incurred	Proactive: FinOps practices advise on architecture and provisioning decisions before workloads run

Across Techwave's NALORA-powered cloud managed services engagements, the shift from the left column to the right column produces consistent, measurable results: 85% proactive issue detection, 60% reduction in alert noise, 40% improvement in mean time to resolution, and 3x faster incident resolution. These are production outcomes, not projected benefits.



The Numbers Behind the Shift

The market trajectory reflects this structural realignment at a meaningful scale. These are not projections for a distant future state. They represent the market that is actively forming around outcome-based managed services today.

\$129.14B

managed infrastructure services market in 2025, projected at \$210B by 2030 (Mordor Intelligence)

\$66.1B

global cloud managed services market in 2025, projected to reach \$172.1B by 2034 at an 11.22% CAGR (IMARC Group)

What This Looks Like Across Cloud and Infrastructure Specifically

Techwave's Cloud & Infrastructure practice operates across four core service areas: cloud migration, managed cloud operations, FinOps and cost optimization, and platform modernization. Across all four, the shift to outcome-based delivery produces a consistent pattern. Contracts anchored to business metrics rather than resource consumption reduce cloud waste by 20 to 35% in the first year of engagement while improving service resilience and deployment velocity concurrently.

The technical mechanisms that make this possible are AIOps-driven operations and embedded FinOps governance. AIOps replaces reactive incident management with predictive monitoring and automated resolution. FinOps moves cost governance from a quarterly reporting exercise to a continuous architectural discipline. Together, they create the operational foundation on which outcome-based SLAs can be credibly committed to and consistently delivered.

63% of enterprises now have a dedicated FinOps team influencing cloud spend decisions, up from 51% in 2024. Value delivered to business units rose 12 percentage points year-over-year as the primary measure of cloud success. (Flexera 2026 State of the Cloud Report)





A Direct Comparison: What the Resource Model Costs vs. What the Outcome Model Returns


The following comparison uses representative industry figures to illustrate the financial and operational difference between the two models on a mid-size enterprise hybrid cloud estate. The patterns are drawn from real engagement data and publicly available market benchmarks.

The Resource Model: A Typical Enterprise Spend Profile

Consider an enterprise running a hybrid cloud estate, with approximately 60% workloads on public cloud platforms such as Azure or AWS and 40% on co-located private infrastructure. Under a conventional resource-based managed services contract, the annual spend pattern typically looks like this:

Managed infrastructure retainer, fixed team of 12 full-time equivalents: approximately \$2.4M per year	Estimated total: \$4.75M annually, with limited traceability from spend to business outcome
Reactive incident management and scheduled patching: approximately \$400K per year	Cloud compute and storage, overprovisioned by an estimated 30 to 35%: approximately \$1.8M per year
Annual cloud governance reporting, point-in-time and limited in actionability: approximately \$150K	

The Outcome Model: Same Estate, Restructured



Representative saving on a \$4.75M estate: approximately \$2M per year, a 42% cost improvement achieved while improving both service resilience and delivery velocity. This reflects patterns observed across outcome-based managed services engagements in Techwave's Cloud & Infrastructure practice.



The AIOps and FinOps Dividend

The two capabilities that most directly underpin outcome-based delivery in cloud and infrastructure managed services are AIOps and FinOps. Their impact on managed services economics is quantifiable.

18 to 24%

reduction in service delivery cost through automated workflow execution in managed services engagements (Future Market Insights, 2025)

12 pts

year-over-year increase in enterprises measuring cloud value by business outcomes rather than cost efficiency alone (Flexera 2026)

The measurement shift is telling. Enterprise leaders are no longer satisfied with a cost reduction report from their cloud partner. They are asking what the infrastructure enabled the business to do: how much faster did releases go out, how much more resilient are critical systems, and what did optimization free up for reinvestment. Those questions cannot be answered by a resource-based contract. They are precisely the questions the outcome-based model is designed to address.





Why Spending More and Moving Faster Does Not Automatically Create Better Outcomes

More cloud spend does not guarantee more business value. A larger team does not guarantee better delivery. And faster infrastructure deployment does not automatically translate into competitive advantage. Three specific traps emerge when enterprises modernize without changing the underlying commercial model.

Productivity Gains That Do Not Reach the Business

Your teams are delivering more than they were two years ago, but has that translated into measurable outcomes, such as lower cost-to-serve, faster time-to-market, or higher platform resilience? For most enterprises, the answer is only partial. Gains exist at the team level but are absorbed into existing complexity rather than surfaced as a competitive advantage. This is the productivity trap.


More Infrastructure, More Complexity, Not a Better Platform

AI-assisted tooling accelerates the scale and complexity of enterprise systems, but more workloads migrated at speed does not mean a better-governed platform. Governance gaps, security debt, and architectural fragmentation are the shadow costs of moving fast without an outcome framework. If workloads are not right-sized and governed, you are accumulating technical and financial debt at the pace of your automation.

A Contract That Does Not Reflect What You Actually Need

If your managed services agreement still prices by resource consumption or headcount, and does not tie your partner's incentives to your business outcomes, you have a commercial misalignment at the foundation of your delivery model. A resource-based contract incentivizes your partner to maximize resources deployed to your account, not to automate low-value tasks, right-size your estate, or surface overprovisioning. The incentives point in the wrong direction.

A straightforward diagnostic: ask your managed services partner what they are commercially incentivized to deliver. The clarity, or absence of it, in that answer will tell you a great deal about whether the relationship is structured for the outcome-based era.





A Practical Framework for Moving From Resource Contracts to Outcome-Based Delivery

The transition to outcome-based cloud and infrastructure managed services does not require replacing everything at once. The enterprises making this shift successfully are progressing through three stages, each building the operational and commercial foundation for the next.

Stage 1: Migrate With Intent, Not Just Speed

Migration is a business decision, not an IT project. Every workload moved to cloud or to a managed infrastructure environment should carry a defined business case: what outcome does this enable, what does steady-state operation cost, and how will success be measured. This requires a managed services partner that brings a point of view on architecture, cost optimization, and business continuity from the outset of the engagement.

Set outcome targets before migration scope: cost reduction, resilience improvement, and time-to-market acceleration

Insist on right-sizing assessments before provisioning cloud resources, not after initial deployment

Treat migration milestones as business outcome checkpoints, not technical delivery gates

Stage 2: Operate With Intelligence, Not Just Coverage

Steady-state management of your cloud and infrastructure estate is where the outcome-based model produces its most differentiated value. AIOps-driven operations, including predictive monitoring, automated remediation, and intelligent capacity management, reduce the human effort required for routine management and redirect that capacity toward higher-value work.

Shift your managed services SLAs from uptime-only to outcome-composite: uptime, cost-to-serve, mean time to resolution, and deployment throughput

Embed FinOps as a continuous architectural discipline, not a quarterly reporting exercise

Measure your partner's automation maturity directly: the ratio of automated to human-resolved incidents is a meaningful leading indicator





Stage 3: Optimize Continuously, With Committed Targets

The outcome-based model is not oriented toward a single transformation milestone. It is designed around continuous improvement, with your managed services partner held to measurable gains against agreed targets on a rolling basis. This is the commercial structure that aligns provider incentives with enterprise priorities.

Structure contracts around committed outcomes: waste reduction, deployment frequency, and critical incident rate

Review commercial terms against delivered outcomes annually, not only at contract renewal

Build internal governance frameworks that make outcome measurement systematic and continuous

When a global IT distributor across North America restructured its 700-server estate under this model, Techwave's NALORA-powered operations delivered 99.9% or higher system availability, sub-one-hour critical incident response, and 3x faster resolution through AI-driven anomaly detection. The outcome-based SLA was not aspirational. It was the contract.





Enterprises That Got Ahead of This Shift Are Already Seeing the Difference

The following patterns are drawn from real-world engagements. They show what outcome-based delivery produces in practice, and why the stakes for enterprises that delay this transition are higher than they appear from within a resource-based contract.

The Mid-Market Enterprise That Competed Above Its Weight Class

A regional financial services firm (\$500M in annual revenue) was competing against tier-one institutions many times its size. Instead of matching headcount, it restructured under an outcome-based model with SLAs tied to cost-to-serve, deployment frequency, and resilience. Within 18 months, cloud spend was down 31%, release cycles moved from quarterly to fortnightly, and the firm achieved 99.97% uptime. The advantage was not the budget; it was the commercial and operational model.

The Large Enterprise That Stayed With the Resource Model Too Long


A global manufacturer with a \$3B technology estate spent three years in proof-of-concept mode running pilots and renewing resource-based contracts while deferring structural decisions. During that period, two regional competitors captured meaningful market share with leaner, faster cloud-native models. The lesson is that the resource model creates inertia: more spend and resources deployed look like progress without changing the commercial logic. Renewal looks like forward motion. It is not.

When a European manufacturing spin-out needed full IT independence within 11 months, Techwave moved the organisation directly to outcome-based delivery: infrastructure right-sized from day one, AIOps embedded at launch, and commercial terms tied to uptime and cost targets rather than headcount. IT operating costs fell 60%. Alert noise was eliminated by 60% from day one, and 100% IT independence was achieved within the window with zero production disruption.

The Market Is Consolidating Around Outcome-Native Providers

The managed services landscape is reshaping around providers with outcome-based capabilities at their core. IBM's \$6.4B acquisition of HashiCorp to embed infrastructure-as-code automation, Accenture's absorption of Navisite's multi-cloud operations, and private equity consolidating regional MSPs for edge-to-cloud coverage, all signal the same directional shift. The provider ecosystem available to enterprises in three years will look materially different from today.

\$210B projected managed infrastructure services market by 2030 (Mordor Intelligence) – the largest share will go to outcome-native providers





The Shift Has Already Started. The Question Is Where You Are In It.

Every enterprise is somewhere on this transition. Some have already restructured their cloud and infrastructure managed services contracts around outcome-based delivery. Others are still in resource-based models, often without a clear view of what that is costing them relative to what it could be delivering.

Neither position is permanent. What matters is the direction you are moving, and the speed at which you are moving there. The commercial model for cloud managed services is changing structurally. The enterprises that recognize this and act on it now, rather than at contract renewal, are the ones that will be best positioned to capture the full value of their cloud investments in the years ahead.

The Techwave Cloud & Infrastructure practice works with organizations across cloud migration, managed operations, FinOps and cost optimization, and platform modernization. If you are evaluating your current managed services model or preparing for a contract restructure, our team is available to walk through what this transition looks like in the context of your specific estate and commercial arrangements.

Sources

- [1] Flexera. 2026 State of the Cloud Report. flexera.com, March 2026.
- [2] Mordor Intelligence. Managed Infrastructure Services Market 2025-2030. mordorintelligence.com, February 2026.
- [3] IMARC Group. Cloud Managed Services Market Size and Trends Report 2034. imarcgroup.com, 2025.
- [4] Future Market Insights. Managed Services Market Demand & Trends 2025 to 2035. futuremarketinsights.com, July 2025.
- [5] IDC. Cloud Infrastructure Spending Q4 2024 and 2025 Forecast. idc.com, March 2025.
- [6] Gartner / CIO Dive. Global IaaS Market 2024: Cloud Infrastructure Services Spend \$171.8B. ciodive.com, August 2025.
- [7] Forrester Research. AI-Enabled Cloud Services Operational Efficiency Prediction, 2025.
- [8] Globenewswire. Managed Cloud Services Market Overview 2025: Hybrid, Multi-Cloud and Edge Architectures. globenewswire.com, September 2025.

About Techwave

Headquartered in Houston, TX, Techwave provides Global Technology and Engineering Services and Solutions to more than 600 clients across 5 continents. Our end-to-end capabilities in Data & Analytics, Software Engineering, AI/ML, Cloud Engineering, and Strategic Sourcing empower clients to maximize the value of their digital investments.



13501 Katy Fwy Suite 3305,
Houston, TX 77079, USA.



info@techwave.com

