

The Price of Intelligence

What happens to the rate card
when the work gets cheap

GAURAV TRIPATHI
MD, Accenture

“

The dirty secret
the industry never
said out loud:
**Getting the team
together was the
project. Delivery
came after.**



Cerebraix Capability Clouds

AI-Powered Managed Capability as a Service

Stop waiting out the hiring cycle. Cerebraix Capability Clouds give you validated AI, Data, Cloud and Engineering capability on demand — found, validated and kept in place for as long as your programme needs it.

Our AI-powered capability platform draws on 30,000+ curated digital and data science professionals.


Managed Capability as a Service (MCaaS)

You request a capability, not a headcount. We find it, validate it and deploy it — validated options within 24 to 48 hours of a brief, working inside one to two weeks. No bench to carry. No hiring cycle to wait out.

Validated Before You See It

Every specialist passes five levels of assessment — ML matching, AI scoring, expert technical vetting, rate optimisation and executive interlock — before a profile ever reaches you. Your team assesses fit, not fundamental competence.

Capability, Built to Your Brief!





Editor's Note

Rishi Bagga

Chief Executive Officer, Cerebraix Technologies
Editor, Cerebraix Talent Tech

Dear Readers,

Each quarter we try to find the question our industry has not yet answered. This time it turned out not to be a question about technology at all.

The firms under pressure in 2026 are not the ones behind on AI. They are the ones still invoicing for a unit of value their clients have quietly stopped buying. Ten Nifty IT constituents shed a combined \$73 billion in market value this year, while mid-tier providers grew several times faster than the giants. That gap is commercial, not technical, and it is the subject of this edition.

Across these pages we look at what clients are actually asking for, the contracts already being signed in place of the billable hour, and the harder internal question of what it costs a firm to make the change. Much of it is unglamorous: baselines, cost-to-serve, incentive design. That is where the work sits.

We are grateful to Gaurav Tripathi, Managing Director at Accenture, for the guest essay in this issue. Writing in a personal capacity, he argues that the real product of IT consulting was never the software but the mobilisation, and that assembling the team was, in truth, the project. Our conclusions meet at the same place. The scarcity has not disappeared. It has moved.


Happy Reading!

Rishi Bagga



COVER STORY

06 The Price of Intelligence:
Why the Rate Card Is
the Real Disruption




13
The same work for 30% less
Anatomy of the client ask

17
Five contracts that
broke the hour

22
Designing the unit
of value

26
What exactly are you
selling now?





39

Who owns the number?
Pricing changes what you
hire for

44


After the reset
Indian it services on the
other side of the rate card

30

Gain-share without giving
away the firm

34

When the contract moves
faster than the cost base



“

It is no longer "how do we resource this?" but
"should we resource this at all?" Underneath lies
something more uncomfortable: were we ever
getting value from the resourcing model, or just
paying for the difficulty of the alternative?



Gaurav Tripathi,
Managing Director, Accenture
Strategy & Consulting, AI & Data.

The Price of Intelligence

Why the Rate Card Is the Real Disruption

The Old World Ran on One Asset: People

For decades, IT consulting was a people business dressed up as a technology business.

The real product wasn't software. It wasn't strategy. It wasn't even transformation. It was mobilisation — hunting the right skills across shrinking talent pools, clearing legal and procurement hurdles, onboarding engineers across geographies and time zones, and assembling the right team at the right moment. That process alone could take weeks. Sometimes months. And the client paid for every hour of it.

Here's the dirty secret the industry never said out loud: **getting the team together was the project.** Delivery came after.

The entire rate card model — bill by resource, scale by headcount, grow by hiring — was built on a single assumption: that skilled human capacity was scarce, hard to assemble, and worth paying a premium for.

The process of mobilisation alone could take weeks. Sometimes months. And the client paid for every hour of it.



What clients bought, without knowing it, they were paying logistics costs and calling them strategy fees. The premium wasn't for the insight. It was for the difficulty of finding the people who had it.

That assumption just broke.

The Hidden Arbitrage, Now Exposed

Advisory value was never really about the advice. It was about access — to credentialed specialists, to frameworks encoding hard-won experience, to institutional knowledge that came only from having done this before.

Because that reserve was hard to replicate, the pricing held. Rate cards inflated. Bench strength became a competitive moat. The size of your talent pool was the size of your business.



What AI has done — quietly, then all at once — is make that reserve available on demand, at a fraction of the cost. The margin between what the insight costs and what the logistics of getting it cost is collapsing. Most firms haven't noticed, because revenue hasn't dropped. Their clients have already started doing the math.

What collapsed, though, was the bundle, not the need. Raw intelligence is now trivially available. Capability an enterprise can trust and deploy is not.

The Client Knows More Than You Think



The New World — AI Moved the Focus to Outcomes

Clients stopped asking "how many people can you put on this?" and started asking **"how quickly can you show results?"** The credibility of the team and the prestige of the firm now matter less than one question: what did you actually deliver?

CFOs are doing the math, and doing it specifically. What did the engagement cost, what did it produce, and what would the same outcome cost today under a different model? The numbers don't flatter the traditional approach.

CIOs face a different discomfort. For many the threat isn't vendor disruption but disintermediation from within. Marketing runs its own analytics; finance automates workflows that once needed IT sign-off. The gateway to technology is discovering the gates are open.

The Client Knows More Than You Think

For years the information asymmetry ran one way. The client had the business problem; the provider had the technical knowledge. That gap justified the engagement model, the day rates, the steering committees.

It is eroding fast. The same AI tools providers deploy internally now sit with the CFO's team and the in-house transformation office. What once required a room full of experienced consultants can now be initiated by a senior analyst with a well-structured problem.

What clients need external help for has changed. They no longer need providers to think on their behalf. They need providers who can think better than they can themselves, and they are increasingly able to judge the difference.



The Question Has Changed

For two decades the standard engagement opened with the same question: "What resources do you need, and for how long?" That made sense when the constraint was human capacity, and success meant execution fidelity.

The question clients ask now is different:

"What changed because of you?"

That is a question about causation, not execution, and it exposes the gap between effort and impact

the old model was good at obscuring. It also closes the old model's favourite exit: when the outcome doesn't materialise, the provider points to scope changes, client dependencies, data quality. The invoices were accurate. The hours were delivered. The outcome was someone else's problem.

Outcome-based models don't allow that exit; they require real skin in the game. Clients pay more for certainty than for effort, and the margin on outcomes is structurally higher than on headcount — precisely because the judgment it demands is still scarce.



Outcome-based models don't allow that exit; they require real skin in the game.

Clients pay more for certainty than for effort

The margin on outcomes is structurally higher than on headcount



The Construct Is Collapsing

When you can stress-test a strategy across decades of data in minutes, the fundamental question changes. It is no longer "how do we resource this?" but "should we resource this at all?" Underneath lies something more uncomfortable: were we ever getting value from the resourcing model, or just paying for the difficulty of the alternative?

The whole construct — rate cards, bench strength, resource augmentation, onboarding timelines — was built for a world where human intelligence was the bottleneck. That world is over. What's left isn't a transition. It's a renegotiation of the entire value proposition, and when a construct this load-bearing goes, something has to take its place.

For the First Time, Business Is Driving IT

Here's what's historic about this moment: for the first time in the history of the IT industry, business users are driving technology decisions.

Not IT teams picking familiar vendors. Not architects designing in silos. People who understand outcomes, not infrastructure.

AI bridged the translation gap that twenty years of digital transformation never closed, making technology legible to people who don't speak infrastructure and business demands legible to systems that don't speak ambiguity. IT is finally enabling, the way it was always meant to.

The New Premium: Outcomes, Not Headcount

Clients aren't buying people anymore. They're buying results. No bench. No obligations. No onboarding debt. The model that made IT consulting a trillion-dollar industry is being renegotiated — not at renewal, but right now, inside active engagements.

The contracts that win are structured around outcomes rather than activities, carry enough risk to demonstrate conviction, and price for value delivered rather than time spent. They don't require the client to trust that work is happening: the result is the evidence. That is uncomfortable for firms whose delivery machinery — the PMO layer, the QA checkpoints, the stakeholder cadence — was built to show effort, not impact.

The value shifts from assembling effort to activating capabilities — reusable, governed, business-ready forms of intelligence



The Layer That Has to Exist

Abundance always creates a new scarcity. When intelligence was expensive, the constraint was access to it. Now that analysis is cheap and instant, the constraint has moved downstream to application — turning general capability into something specific, governed and dependable enough to run a business on.

This is the gap most enterprises are sitting in. They can generate a credible strategy in an afternoon and still cannot stand up a working capability in a quarter. The models are general; the problems are specific, regulated and entangled with legacy. The people who close that distance — who can specify AI-assisted work, vet what it produces and operate it under governance — are scarcer than the engineers the old model was built to supply.



Enterprises cannot build that capability fast enough, and they can no longer buy it by the hour, because the hour has stopped being the unit. As intelligence becomes abundant, they do not simply need more tools or cheaper access to analysis; they need a trusted human capability layer that turns AI potential into repeatable business outcomes. The value shifts from assembling effort to activating capabilities — reusable, governed, business-ready forms of intelligence that can be applied where outcomes matter most.

The unit of access becomes capability rather than headcount. And the boundary matters: this layer does not own the enterprise's outcome. The enterprise does. What the layer owes is that the capability is there, proven and dependable, at the moment the outcome depends on it.

That layer will exist. The open question is who operates it — the incumbent firms that already hold the client relationships, or the platforms being built specifically to do it.



Intelligence that
can be applied
where outcomes
matter most.




The Rate Card Is Falling. The Opportunity Is Rising.

As the cost of intelligence drops, something counterintuitive happens. The question stops being "how do we do this faster?" and becomes "what could we now attempt that we never dared before?"

Cheaper intelligence doesn't just compress timelines. It expands ambition. It opens doors that were always locked — not by capability, but by cost. A strategy that needed a six-month diagnostic can be pressure-tested in a week.

The firms that grasp this will grow faster than the old model ever allowed, because the market for "help me figure out what to do with this capability" is vastly larger than the market for "help me build what I already know I need." The old model scaled by hiring. The new one scales by thinking — and by having the capability to hand when the thinking is done.



Who Moves First?

As the cost of intelligence drops, something counterintuitive happens. The question stops being "how do we do this faster?" and becomes "what could we now attempt that we never dared before?"

Cheaper intelligence doesn't just compress timelines. It expands ambition. It opens doors that were always locked — not by capability, but by cost. A strategy that needed a six-month diagnostic can be pressure-tested in a week.

The firms that grasp this will grow faster than the old model ever allowed, because the market for "help me figure out what to do with this capability" is vastly larger than the market for "help me build what I already know I need."

The old model scaled by hiring. The new one scales by thinking — and by having the capability to hand when the thinking is done.

THE SAME WORK FOR 30% LESS

ANATOMY OF THE CLIENT ASK

Clients are asking IT services providers for the same scope of work at 25% to 30% below previous rates, on the reasoning that AI has lowered the provider's cost of delivering it. The request is usually reasonable on its own terms. What determines the outcome is whether the provider arrives with an alternative commercial construct or only with a defence of the existing one.

The ask, in the client's words

Persistent Systems chief executive Sandeep Kalra summarised the pattern in August 2026: clients want "the same work for 25% to 30% less," delivered faster and with higher productivity. Everest Group's Jimit Arora described the resulting market as one where "the odds are very much in favour of clients."

When a provider's own website claims a 30% uplift, procurement treats that as an admission, and prices accordingly.

Horizon Solutions Strategic Partnership Proposal

Business proposal is filed and corporate...
company (names or an excited company...
• Professional partnership-action and adaptive
• Project lifecycle management
• Project lifecycle and business proposal
• Evelop partnership business proposal

5. Project Budget Breakdown

The random list of budget line items are to the project subained the budget itave no eei-evoo-is becoms tbe not the budget line items:

Project Management	₹3,45,000
Resource Allocation	₹8,15,000
AI-assisted delivery	₹12,70,000
Data Analytics	₹4,90,000
R&D	₹6,25,000
Infrastructure	₹9,50,000
Marketing	₹5,80,000
Contingency	₹2,10,000
Total Proposed Budget (INR):	₹52,85,000

-30%
Revised: ₹37,00,000?

It helps to understand why the client believes this is fair rather than opportunistic. Enterprise buyers have watched their own engineering teams adopt AI tooling and have seen internal throughput change. They read the same vendor marketing about productivity gains that providers circulate.

When a provider's own website claims a 30% uplift, procurement treats that as an admission, and prices accordingly.

The provider is therefore negotiating against its own sales material.

Four clause types now appearing at renewal

The ask rarely arrives as a blunt demand for a lower rate. It arrives structured, and the structure matters more than the headline number.

1 Productivity pass-through.

A clause committing the provider to reduce charges by a stated percentage each year on the assumption that AI-assisted delivery improves. It looks modest at 4–5% annually and compounds into a serious margin problem across a five-year term.

A REASONABLE ANSWER:

Accept a pass-through only where it is tied to a measured baseline the two parties agree, and pair it with a floor below which rates cannot fall regardless of measured gain.

THIS IS A
DEFENSIBLE
READING OF
AN UNCERTAIN
MARKET

2 AI-adjusted rate benchmarking.

The client benchmarks the provider's rates against a market reference that already reflects AI-era pricing, then requires alignment. The reference set is usually opaque.

A REASONABLE ANSWER:

Ask to see the comparator basis and insist that benchmarking compare like-for-like scope and accountability, not headline rates alone. A provider carrying delivery risk is not comparable to one supplying capacity.

THIS IS A DEFENSIBLE
READING OF AN
UNCERTAIN MARKET,
AND IT IS THE
POSITION MOST MID-
MARKET FIRMS
HAVE TAKEN.

FOUR CLAUSE



S Mehta

3 Savings-share language.

The client proposes that AI-generated savings be split, which sounds equitable and often is. The difficulty is definitional: savings against what baseline, measured by whom, over what period.

A REASONABLE ANSWER:

Engage willingly, but make the baseline a joint artefact agreed before signature. Cognizant's February 2026 agreement with Daimler Truck, which splits AI-driven savings between vendor and client, is a working example of the structure done deliberately rather than retrofitted.

4 Term compression.

Rather than cutting rates, the client shortens the contract from five years to two, preserving the option to reprice sooner.

A REASONABLE ANSWER:

Price the shorter term accordingly. A two-year contract carries different amortisation of onboarding and transition cost than a five-year one, and that difference is legitimate to reflect.



Why holding the line usually loses

The instinctive response is to defend the rate card and wait, on the view that clients have not yet asked in most accounts and quality will tell. This is a defensible reading of an uncertain market, and it is the position most mid-market firms have taken.

It carries a specific risk. A provider who concedes nothing until the client insists ends up conceding at the client's number, at renewal, under time pressure, with no alternative construct on the table to trade against. The discount arrives anyway, and arrives as a pure margin reduction rather than as part of a restructured deal.

Firms that have moved earlier are trading the discount for something. ExlService reports outcome-based arrangements at roughly 36% of revenue with a target in the mid-40s over three to five years. Genpact, at around 24%, attributes about 300 basis points of gross-margin expansion to agentic solutions.

These are not firms that avoided the price conversation. They are firms that changed what was being priced.

What to have ready before the meeting

Preparation determines the outcome more than negotiation skill does. Four things are worth assembling before a renewal conversation opens.

What to bring

A measured baseline for the account

Two alternative constructs, priced

A clear statement of what you are accountable for

Your own cost-to-serve, honestly modelled

Why it changes the conversation

Turns "savings" from an assertion into arithmetic

Gives the client something to choose rather than something to refuse

Prevents scope drift into outcomes you cannot control

Tells you which concessions are survivable

THE LAST OF THESE IS WHERE MOST FIRMS ARE WEAKEST.

A provider who cannot state the cost of delivering a given outcome cannot know whether a proposed discount is painful or fatal, and will therefore either refuse reasonable terms or accept unreasonable ones.



FIVE CONTRACTS THAT BROKE THE HOUR

Five contract structures signed between mid-2025 and early 2026 show how IT services work is being priced once the billable hour stops being the unit. Each defines value differently, places risk differently, and demands a different internal capability from the provider. None is a template. Together they map the available design space.

Provider risk rises as the unit of value becomes more abstract and harder to measure.



1

Shared savings: Cognizant and Daimler Truck

In February 2026 **Cognizant** signed an AI and automation agreement with **Daimler Truck** structured so that the cost savings generated by AI implementation are split between vendor and client.

Unit of value

A rupee or euro of verified cost reduction.

Risk

Shared, roughly in proportion to benefit.

What it demands internally:

A baseline both parties accept, and instrumentation that survives audit. A shared-savings deal is arithmetic, and arithmetic requires an agreed starting number.

This is the most intuitive of the five structures and the hardest to retrofit.

Signed at the start of an engagement it works. Proposed halfway through, when the pre-AI cost base is no longer observable, it collapses into argument.



2

Deferred payment: HCLTech and E.ON

HCLTech's multiyear cloud management contract with E.ON, signed in June 2025, carries no compensation in the first year. Payments from year two are tied to efficiency gains and specific business outcomes.

Unit of value

Demonstrated efficiency, paid in arrears.

Risk

Carried almost entirely by the provider in the early period.

What it demands internally:

A balance sheet that can fund twelve months of delivery before revenue recognition, and enough confidence in the outcome to justify doing so.

This structure separates firms that can talk about outcome pricing from firms that can carry it.

A mid-sized provider might sustain one such contract. Three at once would be a financing decision rather than a commercial one.



3

Metric-linked managed services: Tech Mahindra in healthcare

Tech Mahindra won a healthcare engagement priced against a defined operational metric set: 40% fewer tickets, 20% lower resolution time, and a 30–35% reduction in technical debt.

Unit of value

Movement in named operational measures.

Risk

Provider, bounded by the specificity of the metrics.

What it demands internally:

The ability to forecast operational improvement with enough accuracy to price it.

The instructive feature here is restraint. Every metric is operational and within the provider's influence. None is a share of the client's business result. That boundary is what makes the contract deliverable, and it is the single most useful thing in this list for a firm writing its first outcome deal.

4

Credit-based value pricing: Grid Dynamics

Grid Dynamics introduced its GAIN framework in June 2025, using credit-based pricing tied to value delivered rather than effort expended. Clients draw down credits against outcomes instead of purchasing hours.

Unit of value

A defined credit, decoupled from headcount.


Risk

Provider, on the accuracy of credit definition.

What it demands internally:

CREDIT MODELS SOLVE A REAL COMMERCIAL PROBLEM.

They let a provider improve delivery efficiency without automatically surrendering the gain, because the client is buying credits rather than time. They fail when the credit is defined so loosely that clients cannot tell what they are buying.



5

Tokenised pods: Globant

Globant's AI Pods replace effort-based billing with a tokenised, outcome-driven model, packaging a mixed human-and-AI team as a consumable unit.

Unit of value
Pod capacity, consumed as tokens.

Risk

Shared, with consumption variability sitting mostly with the client.

What it demands internally:

A delivery model genuinely designed around mixed human and machine capacity, not a traditional team with tooling added.

This is the structure closest to how software is sold, and the furthest from how services have traditionally been staffed. It only works if the underlying delivery model actually is what the pricing claims.

Two patterns are worth noticing.

First, provider risk rises as the unit becomes more abstract. A ticket is easy to count and hard to argue about. A credit is flexible and contestable. The commercial upside runs in the same direction as the measurement difficulty.

Second, each structure fails on an internal capability rather than on a commercial idea. The constraint on adoption across the industry has been consistently reported as provider capacity to absorb financial risk, not client appetite. Reporting on Indian IT in August 2026 identified exactly this: relatively few providers have the operational depth and skills to take on the risk that outcome pricing implies.

That is a solvable problem, but it is solved in delivery and finance, not in the sales conversation.

Where to start

A firm signing its first non-hourly contract should look hardest at the Tech Mahindra shape. Bounded operational metrics, within the provider's influence, measured by instruments both sides can read. It carries real risk without carrying unbounded risk, and it teaches the measurement discipline that every other structure on this list assumes.

The deferred and credit-based models are for later, once the firm knows its own cost of achieving a result rather than its cost of supplying a person.



DESIGNING THE UNIT OF VALUE: WHAT EXACTLY ARE YOU SELLING NOW?

The unit of value is the thing a client actually buys and the thing an invoice counts. In a time-and-materials contract it is a person-month. Once that unit stops making sense, choosing its replacement becomes the hardest commercial decision a services firm makes, because every other term in the contract is derived from it.


Why the choice determines everything downstream

Two firms can agree on the same strategy and reach entirely different outcomes because they picked different units. Cognizant reports roughly 45% of BPO contracts on outcome-based models.

Coforge puts outcome-based work at 6–7% of total revenue. Both are moving the same way. Much of the distance between them sits in how tractable a unit they chose and whether their delivery model could carry it.

A well-chosen unit makes pricing, measurement, dispute resolution and capacity planning follow naturally.

A poorly chosen one turns every month-end into a negotiation.



The candidates

Unit	Works well when	Fails when
Ticket or incident	Volume is high and definitions are stable	Clients change what counts as a ticket
Resolution or case closure	Quality is measurable and disputes are rare	Complexity varies widely between cases
Transaction processed	The process is standardised and audited	Volume is outside the provider's influence
Model or feature in production	Deployment is a clear, dated event	"In production" is loosely defined
Capability-month	The client needs assured access to scarce skill	The client expects an outcome, not access
Business KPI	The provider genuinely controls the lever	Anyone else's decisions move the number

The last row is where most first-time outcome contracts go wrong, and the reason is understandable. A business KPI is the easiest unit to sell, because it is the thing the client cares about. It is also the unit most contaminated by decisions the provider does not make.



Four tests for any candidate unit

Before a unit reaches a proposal, run it against four questions. A unit that fails any one of them will cause trouble; a unit that fails two should not be offered.

Explanation:

1 Does the client recognise it as value?

Your answer here....

If the buyer has to be persuaded that the unit matters, the unit is wrong. Credits and tokens fail here more often than their designers expect. Grid Dynamics' GAIN framework and Globant's AI Pods both work as pricing constructs precisely because the vendor invested in making the unit legible to a buyer.

2 Can it be measured without argument?

Your answer here....

The measurement has to come from an instrument both parties can read, on a cadence both parties accept. If the number arrives in a monthly provider report that the client cannot independently verify, the unit will eventually be disputed.

3 Can either side game it?

Your answer here....

A ticket-based unit rewards ticket creation. A resolution-based unit rewards closing cases fast rather than well. Every unit creates an incentive; the question is whether that incentive points in a direction both parties can live with. Pairing a volume unit with a quality gate usually resolves this.

4 Do you know your cost of delivering one?

Your answer here....

This is the test most firms skip and the one that decides profitability. If a provider cannot state the cost of resolving one case or deploying one model, it cannot know whether a price is generous or ruinous, and it will discover which only after the contract is signed.

The baseline problem

Most outcome contracts are lost at signature rather than in delivery, and the usual cause is a missing baseline.

If the pre-engagement position was never measured, every subsequent improvement

becomes contested. The provider says ticket volume fell 40%; the client says volume was already falling, or that the definition changed, or that seasonality explains it. Neither party is acting in bad faith. They simply have no shared starting number.

The discipline that avoids this is unglamorous:

- 1 Instrument the account before proposing, not after signing.
- 2 Agree the measurement method, the data source and the reporting cadence in writing, separately from the commercial terms.
- 3 Record the baseline as a joint artefact, signed by both parties, with the definitions attached.
- 4 Write the dispute mechanism while both sides are still optimistic. It is the only time it can be written fairly.

Tech Mahindra's healthcare engagement, priced against 40% fewer tickets, 20% lower resolution time and a 30–35% reduction in technical debt, only functions because each of those measures has a defined starting point. Without the baseline, the same contract is unenforceable in either direction.

A note on vendor productivity claims

Firms designing their first unit often anchor on published productivity figures. A 30% uplift here, six-times-faster adaptation there, three-times-less production time elsewhere. These numbers come from companies selling the frameworks that produce them.

Treat them as positioning, not as planning inputs. A pricing model built on someone else's

marketing claim will price the work at a cost the firm cannot actually achieve.

The alternative is slower and more reliable: run two or three engagements with the new delivery model, measure the real cost of a unit in your own environment, and price from that. The first contracts should be priced to learn rather than to win.



GAIN-SHARE WITHOUT GIVING AWAY THE FIRM

Gain-share pricing pays a provider a share of the value its work creates rather than a fee for the effort it expends. It is the most commercially attractive of the new structures and the one most capable of damaging a mid-sized firm, because the upside is capped by the deal while the downside is capped only by the balance sheet.

Why the industry keeps stalling here

Reporting on Indian IT through 2026 has been consistent about what limits adoption of outcome and gain-share models, and it is not client appetite. Relatively few providers possess the operational depth and financial capacity to assume the risk these contracts transfer.

Deals also tend to be smaller and shorter than the contracts they replace, which shortens the period over which a provider can recover an investment.

Infosys chief executive Salil Parekh has described client interest in outcome-based work as real but not yet a large part of activity. The gap between stated direction and booked revenue across the industry is mostly a risk-carrying gap, and the numbers show its width. Cognizant reports roughly 45% of BPO contracts on outcome models; Coforge puts outcome-based work at 6–7% of total revenue. Both firms want the same destination.

One has built the capacity to carry the risk and the other is building it.



That is worth stating plainly because it reframes the decision. Declining a gain-share deal is not a failure of ambition. Signing one the firm cannot absorb is a failure of arithmetic.

The four mechanisms that make risk survivable

CAPS:

A ceiling on provider upside. Clients ask for these and providers resist, but caps are the price of floors and generally worth paying. A capped, floored contract has a known range. An uncapped, unfloored one has a known range too, and the bottom of it is unpleasant.

RATCHET PROTECTION:

The subtlest of the four. In a multi-year gain-share, this year's achieved outcome becomes next year's baseline, so the provider must keep improving against its own improvements to earn the same money. Without an explicit reset mechanism, a successful contract becomes progressively harder to deliver profitably until it turns loss-making.

Ratchet protection usually takes one of three forms: a baseline reset at defined intervals, a declining share percentage that acknowledges diminishing returns, or a fixed floor that rises with demonstrated performance.

DEPENDENCY CLAUSES:

Named conditions outside the provider's control that suspend or adjust the measurement. Client staffing changes, scope additions, data quality below an agreed threshold, adoption rates below a stated level. These are not escape hatches; they are the difference between a contract the provider can deliver and one where the provider is exposed to the client's own decisions.

FLOORS:

A minimum payment regardless of measured outcome, typically covering direct delivery cost. A floor converts a potentially catastrophic contract into a low-margin one. Most first-time gain-share deals should carry a floor at or near cost-to-serve, and a provider unwilling to discuss floors is usually a provider who has not modelled the downside.

Caps

Ratchet
protection

Dependency
clauses

Floors





The downside test

Before signing, the question is not what the contract earns if it works. It is what happens if it does not.

A workable discipline is to model three cases and one decision:

Case	Assumption	What to check
Expected	Outcome achieved as forecast	Margin versus the time-and-materials alternative
Plausible miss	Outcome achieved at 60% of forecast	Whether the contract is still cash-positive
Bad	Outcome not achieved; floor only	Whether the firm can carry it for the full term
Concurrent	Three such contracts in the bad case	Whether the firm survives that year

The concurrent case is the one firms skip and the one that matters. A single contract in difficulty is a management problem. Three at once, in a firm of a few hundred people, is a solvency problem.

HCLTech's cloud management agreement with E.ON is instructive at the extreme end. Signed in June 2025, it carries no compensation in the first year, with payment from year two tied to efficiency gains and business outcomes.

HCLTech can fund twelve months of delivery from its balance sheet without strain.

A provider one-hundredth the size structurally cannot, however good the engagement looks.





Some gain-share proposals should be refused politely and early.

Deals worth declining

1

No agreed baseline.

If the starting position cannot be measured and jointly signed, the contract has no arithmetic and will end in dispute.

2

Outcomes driven by client behaviour.

Adoption rates, internal process compliance and business volumes are not provider-controllable. Sharing risk on them is gambling on someone else's management.

3

Unbounded exposure.

Any structure where the provider's downside has no defined floor should be declined regardless of upside.

4

Short term, long payback.

A two-year contract requiring eighteen months of investment leaves almost no earning period, and the client holds the option to reprice at the moment it becomes profitable.

5

First deal, largest account.

Learning should not happen on the account the firm cannot afford to lose.

The upside is real

None of this argues against gain-share. The evidence on the other side is substantial. ExlService reports outcome-based arrangements at roughly 36% of revenue, targeting the mid-40s within three to five years. Genpact, at around 24%, attributes roughly 300 basis points of gross-margin expansion to agentic solutions. Firms that have built the capability are earning more per unit of delivery, not less.

The pattern among them is consistent. They started on bounded contracts, built measurement

discipline, learned their real cost of achieving a result, and scaled from there. None of them began by putting a flagship account on an uncapped share of business outcomes.

Risk-sharing is a capability, and capabilities are built in a sequence.


The firms that will be taking meaningful commercial risk in 2028 are the ones taking small, well-structured, slightly uncomfortable risk this year.

WHEN THE CONTRACT MOVES FASTER THAN THE COST BASE

Outcome-based pricing makes revenue variable while a permanent delivery bench keeps cost fixed, and the gap between those two facts is where margin disappears. This is the structural problem underneath the entire pricing transition, and it is more consequential for a mid-sized firm than any individual contract term.

Time and materials

Revenue and cost move together




Person off project > gap shows in days

— Revenue - - - Delivery cost

Outcome-based pricing

Variable revenue, fixed bench



Bench cost stays, margin disappears

■ Margin ■ Margin lost

The arithmetic

Under time and materials, revenue and cost move together. A person on a project generates revenue; a person off one stops, and utilisation reporting makes the gap visible within days.

Under outcome pricing that link breaks, and it breaks in both directions. Revenue attaches to a result delivered on the client's timetable. Cost attaches to people employed on the provider's payroll. A contract that pays on measured improvement in month nine still requires a fully staffed team in months one through eight.

Three further characteristics of Al-era deals sharpen the problem. Contracts are reported to be smaller and shorter than those they replace, which compresses the period over which fixed capacity can be amortised.

Payment is often back-loaded, as in the multiyear cloud management agreement HCLTech signed with E.ON in June 2025, where no compensation flows in the first 12 months. And the skills required change faster than a permanent workforce can be retrained, so a team optimised for last year's delivery model is not automatically the team that hits this year's outcome.

The scale of the squeeze is visible in the 2026 numbers. The Nifty IT index fell 20% across the year, with 10 constituents shedding a combined \$73 billion in market value, while Persistent Systems reported clients asking for the same scope at 25% to 30% less.


A firm absorbing a 25% price cut on a cost base that has not moved has a structural problem, not a margin one.

The options, honestly costed

No arrangement removes the mismatch. Four manage it, each with a real price.

APPROACH	What it does	What it costs
HOLD THE BENCH	Retains capability and culture	Full fixed cost through revenue troughs
CORE-AND-FLEX	Permanent core, variable surround	Coordination overhead; weaker cohesion at the edges
PARTNER CAPACITY	Contracted access to skills on demand	Margin share; dependency on partner quality
SHRINK AND BUY LATE	Minimal standing capacity	Slow mobilisation; the exact weakness AI made expensive

The fourth option deserves attention because it is the instinctive response to margin pressure and the most dangerous. A firm that strips its bench to protect this quarter's margin recreates the mobilisation delay that used to be the industry's differentiator, except that clients now expect speed and will not pay a premium for it.



Core-and-flex, in practice

Most firms that have made this work have landed on some version of a core-and-flex structure, and the design questions it raises are unusually specific.

CORE


What belongs in the core?

Client relationship depth, domain knowledge, delivery leadership, architecture, and the capability the firm competes on. These are the things that take years to build and cannot be acquired at short notice.

FLEX

What belongs in the core?

Specialist skills needed in bursts, capacity for volume peaks, and capability the firm needs occasionally but cannot keep utilised. A machine-learning engineer needed for eleven weeks a year is a flex requirement regardless of how strategically important those eleven weeks are.



What is the ratio?

There is no industry answer, but the question can be derived. Take the portfolio, identify revenue that is genuinely committed for more than twelve months, and treat that as the maximum defensible permanent capacity. Everything above it is a bet on renewal.

Most mid-sized firms discover their permanent headcount implies confidence in revenue they do not actually have contracted.

What does the flex cost?

Contracted access to capability carries a margin share that looks expensive against an employed rate card. The comparison is misleading, because the employed rate card excludes the cost of carrying that person through non-billable periods. The correct comparison is flex cost against fully-loaded annual cost divided by realistic utilisation,

and on that basis the gap narrows sharply.

This is the practical case for treating capability as something a firm accesses rather than owns — the argument behind managed capability models generally, and one worth evaluating on the arithmetic rather than the positioning.



What to change first

Three moves, taken in order, do most of the work.




MEASURE COST-TO-SERVE AN OUTCOME, NOT A PERSON.

Most firms know their cost per person-month precisely and their cost per resolved case not at all, and outcome pricing requires the second number. Until it exists, every outcome contract is priced on hope rather than on arithmetic.



SEPARATE COMMITTED REVENUE FROM EXPECTED REVENUE IN CAPACITY PLANNING.

Permanent headcount should be sized against contracted revenue, not forecast revenue. Firms routinely plan against the latter and then describe the consequence as a bench problem when renewals slip.



BUILD THE FLEX RELATIONSHIP BEFORE IT IS NEEDED.

A partner arrangement negotiated under pressure, mid-contract, costs more and delivers worse than one established in advance. What is being purchased in a standing arrangement is lead time on access to scarce capability, which is precisely the thing that cannot be bought quickly.

The measure worth watching

Utilisation was the right metric for a time-and-materials business. It answers a question that matters less each quarter: what proportion of the people we employ are billing.

The useful replacement is a cost-to-outcome ratio — what it costs to deliver one unit of the thing clients now buy, tracked over time. A firm improving that number is genuinely getting better at the new model. A firm with excellent utilisation and a flat cost-to-outcome ratio is running the old business efficiently while the market moves.

Persistent grew revenue 16% and Coforge 33% in the second quarter of 2026 while TCS, Infosys, Wipro and HCLTech managed 1–3%. Size was not the differentiator between those two groups; the ability to reconfigure a delivery model at speed was, and that ability is largely a function of how much of the cost base can actually move.

WHO OWNS THE NUMBER?

Once a services contract pays against a business metric, the provider becomes exposed to everything that moves that metric, including decisions made entirely inside the client organisation. Defining which part of the number the provider owns is not a legal formality; it is what separates a contract that can be delivered from one that cannot.

The exposure matters more each quarter because the contracts carrying it are multiplying.



The exposure nobody prices

Consider a provider paid on a 25% reduction in customer service handling time. It redesigns the workflow, deploys automation and retrains the support team's tooling. Handling time does not move, because the client's product team shipped a release generating a new class of complex queries and the client's own agents were never given time to adopt the new process. The provider did the work; the number did not move; under a poorly drafted contract the provider is not paid.

This is the most common way outcome contracts sour, and it damages the relationship rather than merely the commercials, because both parties believe they behaved correctly.

The exposure matters more each quarter because the contracts carrying it are multiplying. Cognizant signs roughly 45% of its BPO contracts on outcome-based models, and around 80% of TCS contracts in its finance, HR and business-services segment now carry outcome performance measures, double the proportion of late 2023. Reporting on Indian IT in August 2026 identified provider capacity to absorb financial risk as the main brake on adoption, and much of that capacity question is not balance sheet at all. It is whether the provider drew the accountability line precisely enough to know what it accepted.

The line that holds

There is a defensible boundary, and it is worth stating precisely because commercial pressure erodes it.

A provider can reasonably commit to four things:

1 **Capability**

the skills and capacity required, available when required.

2 **Availability**

service levels, response times, coverage windows.

3 **Throughput**

volume processed, cases handled, features deployed.

4 **Quality**

defect rates, rework, accuracy, compliance with defined standards.

A provider cannot reasonably commit to four others:

1 ADOPTION: whether the client's people use what was built.

2 SCOPE DISCIPLINE: whether requirements stay stable through the term. service levels, response times, coverage windows.

3 DATA QUALITY: the state of inputs the provider does not generate.

4 BUSINESS RESULTS: revenue, market share and customer satisfaction, which depend on pricing, product, competition and a dozen factors outside the engagement.

The second list is not a list of excuses. Each item describes a lever held by the client. A provider who accepts accountability for them has not been bold; it has been imprecise, and imprecision in a contract is a liability that arrives later.

Why Tech Mahindra’s metric set is worth studying

The healthcare engagement Tech Mahindra won, priced against a 40% reduction in tickets, a 20% reduction in resolution time and a 30–35% reduction in technical debt, is instructive for what it does not contain.

	Operational, not business	Within provider influence	Measured from shared systems
40% fewer tickets Ticket volume	✓	✓	✓
20% faster resolution Resolution time	✓	✓	✓
30–35% less technical debt Code and platform health	✓	✓	✓

NOT IN THE CONTRACT: BUSINESS RESULTS

e.g. Patient outcomes

e.g. Revenue growth

e.g. Client cost savings

None of those three measures is a business result. Each is operational, each sits substantially within the provider’s influence, and each can be measured from systems both parties can read. The contract carries genuine risk — missing any of them has consequences — but the risk is bounded by the provider’s own performance rather than by the client’s decisions.

Compare that to a contract paying on patient satisfaction or cost per episode of care. Those are the outcomes the client ultimately wants, and they are also shaped by clinical staffing, payer behaviour and regulation. A provider carrying that risk is being paid to influence a system it does not control.

The discipline is to price against the operational measures that drive the business result, not against the business result itself.

Writing it down

Four mechanisms make the boundary enforceable rather than aspirational.



01

Dependency clauses.

Named conditions that must hold for the measurement to apply: client staffing at agreed levels, data quality above a defined threshold, adoption above a stated percentage, scope frozen or formally repriced. Each should specify the consequence if it fails, whether that is suspension of the measure, an adjusted target, or a fallback to a time-based fee.



02

Change control with teeth

Outcome contracts fail through accumulated small scope additions more often than through single large ones, so the mechanism has to be light enough that people use it and firm enough that it is not bypassed.



One large change
Less common cause



Many small additions
The usual failure



03

A monthly joint governance forum.


Both parties' operational leads reviewing the metric and its dependencies together, to surface drift while it is still correctable rather than at the annual true-up.



04

A named owner on each side.

Every measure should have one accountable person at the provider and one at the client. Metrics without named owners become nobody's problem until they become everybody's dispute.



The conversation to have out loud

The most useful thing a provider can do in a negotiation is state the boundary explicitly rather than leaving it to drafting. Something close to:

We will commit to the capability being available, to the throughput, and to the quality standard. We will work alongside you on adoption and report on it. We cannot be paid or not paid on whether your teams adopt it, because we do not control that lever, and a contract that pretends otherwise will damage the relationship in month nine.

Clients respond better to that than providers expect, and a buyer who has run a failed outcome contract before recognises the honesty immediately. The alternative, accepting unbounded outcome accountability to win the deal, purchases a signature at the cost of a dispute that arrives later, costs more, and usually ends the account.



The governance point most firms miss

There is a version of this that goes wrong internally rather than contractually. When a firm signs its first metric-linked contract, accountability inside the delivery organisation often stays where it always was, with a project manager measured on schedule and effort. The commercial model changed; the internal accountability model did not, and the team optimises for the old measures while the contract pays on new ones.

Someone in delivery has to own the contracted number, with the authority to change how the work is done in order to move it. Without that, the firm has taken commercial risk it has no internal mechanism to manage.

PRICING CHANGES WHAT YOU HIRE FOR

Hiring built for time-and-materials optimises for utilisation: fill seats, keep them billable, grow headcount with revenue. Hiring built for outcome-based delivery optimises for hitting a number. Those two objectives produce different team shapes, different seniority mixes and a different tolerance for capability the firm cannot deploy quickly, and the second one is considerably harder to run.



The question changes

Under the old model the workforce planning question was whether the firm could fill the seats a signed contract required. It was a supply question, answered with a pipeline and a notice period.

Under outcome pricing the question is whether the firm can hit the contracted number, and how quickly it can assemble the people who will. That is a capability question, and it has a clock attached.

A contract priced on a twelve-month improvement does not allow three months of recruitment before work starts.



The practical consequence is that the firm's hiring problem moves earlier — from after the deal to before the proposal. A provider cannot responsibly price an outcome it has no route to staffing.

What changes in team shape



Smaller and more senior.

Outcome delivery rewards teams that get the approach right early, because rework is absorbed by the provider rather than billed. A team of six people who scope correctly beats a team of fourteen that iterates to the same place.

Deeper specialisation, held for less time.

The skills that move a specific outcome are often narrow and needed in bursts. A firm delivering an operations improvement may need a process-mining specialist for seven weeks and never again on that account.



Fewer generalist junior seats.

This is the uncomfortable one. Under utilisation-led staffing, junior engineers filled billable seats and learned on the job. Under outcome pricing, work that used to be a junior's apprenticeship is frequently the work AI now does, and the seat no longer exists in the same form.

The market data reflects this across the industry. SignalFire's State of Talent 2026 reports new-graduate and entry-level hiring down roughly 65% at the largest technology employers and around 76% at early-stage companies against a 2019 baseline. Research from the Observer Research Foundation on India's GCC sector reaches a similar conclusion, finding AI-driven displacement concentrated among entry-level roles and disrupting the traditional pathway for fresh engineering graduates.


The trap in cutting the bottom

Freezing fresher intake is individually rational and collectively expensive. The senior engineer of 2031 is a junior today, doing work that an agent now does faster. A firm that stops hiring juniors for three years buys its 2031 senior bench at market rates from firms that did not stop.

There is also a delivery argument. Outcome contracts require judgement — knowing which approach will work, spotting when a metric is drifting for the wrong reason. Judgement is built

by doing work and being corrected, and a firm with no junior intake has no mechanism for producing it internally.

The firms handling this well are redesigning the first year rather than deleting it. Juniors placed as supervisors of AI-assisted workflows, as evaluators checking machine output, as the people who build and maintain the context that automated delivery depends on. The work is different; the learning function survives.



It is worth being honest that this is a cost with a deferred return, which is exactly why the pressure runs the other way.


The scarcity moved

A common misreading of this period is that skilled people are becoming less scarce. The evidence points the other way.

What has become abundant is generic capacity — the ability to supply a competent engineer against a standard specification. What has become scarcer is the narrow, current, deployable capability that determines whether an outcome contract is met: the person who has actually built the thing being priced, available in the window when it is needed.

India's position in this is structurally strong and unevenly realised. The country accounts for roughly 28% of the global STEM workforce. More than 1,200 of India's 1,700-plus global capability centres now host AI and machine-learning capability, with around 250 running dedicated innovation centres.

Depth exists.
Availability of
specific, current skill
at short notice is the
constraint.



Three practical shifts

From	To	Why it matters
Hiring to fill signed seats	Confirming capability route before pricing	You cannot price what you cannot staff in the window
Utilisation as the primary metric	Time-to-assemble a capable team	Speed of assembly is now the commercial variable
Permanent headcount for all capability	Permanent for core, accessed for specialist	Burst skills cannot be kept utilised

The third row is where most of the financial consequence sits, and it connects directly to the cost base problem. A firm that employs every capability it might need carries the cost of idleness. A firm that employs none of them cannot mobilise. The design task is deciding which capabilities are core to what the firm competes on and which are genuinely episodic.


What to measure instead

Two metrics tell a services leader more about outcome-readiness than headcount growth.



TIME TO ASSEMBLE.

From a signed outcome contract to a fully staffed, capable team at work. If this exceeds a month, the firm cannot credibly price short-horizon outcomes.



CAPABILITY COVERAGE AGAINST THE PIPELINE.

For the deals currently being pursued, what proportion have a confirmed route to the specific skills required. A pipeline priced on capability the firm has not secured is a forecast of future difficulty.

Neither appears on a standard services dashboard. Both determine whether the commercial model the firm has announced is one it can actually deliver.

AFTER THE RESET: INDIAN IT SERVICES ON THE OTHER SIDE OF THE RATE CARD


Five predictions about Indian IT services by the end of 2028, each stated specifically enough to be graded later, and each paired with the evidence that would prove it wrong. Predictions without a falsifier are press releases, so this edition closes by attaching one to each.




1 Outcome-linked work passes a quarter of industry revenue



EARLY MOVERS ARE ALREADY THERE



What would prove it wrong



Watch for: first significant public arbitration over an outcome contract

The call.

By the end of 2028, non-effort-based pricing — outcome-linked, gain-share, consumption and credit models combined — will exceed 25% of revenue across the listed Indian IT services sector, up from low single digits at most firms today.

The reasoning.

The trajectory already exists at firms that moved early. ExlService reports outcome-based arrangements at roughly 36% of revenue with a stated target in the mid-40s within three to five years. Genpact sits near 24%. Cognizant has around 45% of BPO contracts on outcome

models. Roughly 80% of TCS contracts in its finance, HR and business-services segment now carry outcome performance measures, double the level of late 2023. When the largest providers move a segment that far in two years, the industry average follows.


What would prove it wrong.

A sustained failure of measurement infrastructure. If disputes over baselines and attribution become common enough that large clients revert to fixed-price contracting for risk reasons, adoption stalls well below this level. Watch for the first significant public arbitration over an outcome contract.

2 The mid-market growth premium narrows but does not close




Q2 2026 growth



Why it narrows but does not close



What would prove it wrong



Precedent: Capgemini–WNS, Oct 2025, 11.2x EBITDA, agentic AI framing

The call.

Mid-tier Indian providers will continue to outgrow the largest firms through 2028, but the gap will compress from roughly ten-to-one to something closer to two-to-one.

The reasoning.

The 2026 gap was extreme: Persistent grew 16% and Coforge 33% in the second quarter against 1–3% at TCS, Infosys, Wipro and HCLTech. That advantage comes from agility in repricing and delivery redesign, and it is a transition advantage rather than a permanent one. Large firms are moving, slowly, and each segment they convert

removes part of the gap. The premium persists because structural agility persists; it narrows because the specific transition ends.

What would prove it wrong.


A wave of large-firm acquisitions of exactly these mid-tier assets, which would convert the growth into the acquirers' numbers and eliminate the comparison. The Capgemini–WNS transaction in October 2025, at 11.2 times EBITDA and framed around agentic AI-powered operations, is the shape of what that would look like.

3 Consolidation removes firms that could not reprice




Exit driver: commercial model, not technical capability

Why this size band is squeezed



What would prove it wrong



Demand cushioned previous downturns and could again

The call.

A meaningful number of Indian services firms in the 500-to-5,000 employee range will be acquired, merged or wound down by end-2028, and the distinguishing characteristic of those that exit will be commercial model rather than technical capability.

The reasoning.

The squeeze is specific to this size band. These firms are too large to operate as boutique specialists and too small to absorb the investment that outcome delivery requires: measurement

infrastructure, balance sheet capacity for deferred contracts, and commercial design capability. Ten Nifty IT constituents lost a combined \$73 billion in market value across 2026 while the index fell 20%, and the valuation reset changes the arithmetic of both buying and selling.

What would prove it wrong.

A broad recovery in discretionary technology spending that lifts revenue across the size band before the commercial transition is forced. Demand has been the cushion in previous downturns and could be again.



4 Utilisation stops being the primary delivery metric



Delivery capacity



Risk: time-and-materials stays the majority
Utilisation remains the primary metric

The call.

By 2028, most mid-sized Indian providers will report internally on a cost-to-outcome measure alongside or instead of utilisation, and permanent headcount as a share of delivery capacity will fall materially.

The reasoning.

The metric follows the revenue model. Utilisation answers what proportion of employed people are billing, which matters when billing is by person and matters progressively less when it is not.

Firms pricing on outcomes need to know the cost of delivering one unit of the contracted result, and that number cannot be derived from utilisation reporting. The shift in cost structure follows for the same reason: variable revenue against fixed cost is not sustainable at scale.

What would prove it wrong.

Persistence of hybrid portfolios in which time-and-materials remains the majority of revenue for long enough that firms reasonably keep running the old metric as the primary one.



5

The scarcity moves rather than disappearing

The call.

Demand for narrow, current, deployable specialist capability will remain tight through 2028 even as generic engineering capacity becomes abundant, and the entry-level pipeline will not recover to 2019 levels.

The reasoning.

Entry-level technology hiring is down roughly 65% at the largest employers and 76% at early-stage companies against 2019, according to SignalFire's 2026 talent research. Analysis of India's GCC sector finds AI displacement concentrated specifically among entry-level roles.

Meanwhile outcome contracts require people who have actually delivered the thing being priced, available inside a contractual window — a much narrower requirement than "a competent engineer."

What would prove it wrong.

A structural response at scale: government or industry apprenticeship programmes that rebuild the first rung, or a shift in AI capability that makes senior judgement itself cheaply substitutable. Neither looks likely inside three years, but both are observable if they begin.

What none of this settles

Three things remain genuinely open, and honest forecasting should say so.

Whether outcome pricing is structurally more profitable than time and materials at industry scale is unproven. Genpact attributes around 300 basis points of gross-margin expansion to agentic solutions, which is encouraging and is one firm.

Whether clients maintain appetite for shared risk through a downturn is untested. Risk-sharing is popular when buyers feel confident; procurement behaviour in a recession may look different.

And whether the mid-market's agility advantage survives contact with the capital requirements of outcome delivery is the central question for the readership of this magazine. Agility helps a firm reprice. It does not by itself fund a year of deferred revenue.

Next October's edition opens by grading these five.



YOUR PATH TO SUCCESS STARTS WITH CEREBRAIX

Where talent meets opportunity
and rewards follow

