

— SERVICENOW ECOSYSTEM

A Paradigm **Shift** in the ServiceNow Consulting Business?

While you're still hiring consultants hour by hour, some are paying half the price — and getting results twice as fast.

BJØRN JARL

CEO, The Cloud People Group

— INTRODUCTION

A Market Frozen in 2015

Something quiet is happening in the ServiceNow market. Not on conference stages. Not in analyst reports. It's happening in deliveries — at a handful of partners who have started working in a fundamentally different way. Their clients pay a fixed price, receive solutions faster, and don't have to worry about the hourly rate of the next consultant.

The question is no longer whether AI is changing the consulting role. It already has. The question is whether clients are aware of it — and whether they allow themselves to be served accordingly.

The model is old and simple: the client needs something built on ServiceNow, the partner provides a consultant, the consultant invoices hours. The more complex the project, the more hours. The more hours, the higher the bill. The client pays for time — not for value. This model has persisted for nearly ten years. It is familiar, controllable, and safe. But "safe" is not the same as "optimal."

"The client pays for time. Not for value. That distinction is worth dwelling on."

ServiceNow in 2026 is not an unfamiliar system. It is the world's most widely used enterprise workflow platform, used by 85% of the Fortune 500. Thousands of projects have been completed. Thousands of problems solved. Thousands of solutions documented, tested, and reused — somewhere. The only question is: by whom?

85%

OF THE FORTUNE 500 RUN ON SERVICENOW IN 2026

NTT DATA / INSEAD · MARCH 2026

Rising AI-driven productivity is challenging time-based pricing and accelerating the transition toward differentiated, productized, and platform-based models. Clients are already pre-pricing AI efficiency gains into their expectations — making traditional FTE-based models economically harder to sustain.

THE OLD MODEL

- Client requests a ServiceNow solution
- Partner provides a consultant
- Consultant invoices hours
- More complexity = more hours = higher bill
- Client pays for time, not for value

— THE QUIET REVOLUTION

You Do Not Need to Start from a Blank Page Anymore

Over recent years, The Cloud People has built something most partners don't have — a structured, systematised competence base in The Cloud People's Resource Management tool Cloudia — competence that's drawn from hundreds of completed ServiceNow projects across industries and countries.

When a new client requests a solution, the team doesn't start from scratch. They start from prior accumulated knowledge: pre-built solution components, a structured code bank, and specialized AI agents that are under constant training in TCP's own projects.

The result is not just faster production. It is a fundamentally different delivery model — one where the volume of prior work is a direct competitive advantage. Not on paper. In practice.

WHAT TCP STARTS WITH

- Pre-built components — core functionality already built, tested, and documented from comparable projects
- Structured code bank — verified code from real implementations, categorised and searchable
- Cloudia — AI resource management matching expertise to tasks
- TCP AI Agents — trained in ServiceNow projects, building up their experiences

10x: What Does It Actually Mean?

"10x" is not a marketing promise. It is an operational reality describing the relationship between what an experienced consultant can produce in a day, and what that same team augmented with AI, pre-built components, and structured knowledge can deliver in the same timeframe.

It is not about replacing the consultant. It is about eliminating the unnecessary work they currently spend time on. When these frictions are removed, quality increases, flexibility for change increases, and cost falls markedly. We are not talking about 10–20% efficiency gains. We are talking about an order of magnitude.

NTT DATA / INSEAD · MARCH 2026

GenAI significantly reduces task completion time across cognitive and technical domains, with duration often decreasing by more than half — with similar gains in programming, systems evaluation, and critical thinking. AI acts as a force multiplier for human capability.

THE PRODUCTIVITY SPIRAL

"The more we have delivered, the better and more cost-effectively we deliver the next time. It is a spiral moving in the right direction — just not for those paying by the hour."

— PLATFORM SHIFT

ServiceNow Goes AI-Native

ServiceNow has repositioned and is currently moving faster than any of their platform competitors in the market. What they are enabling is not about new features — it is a strategic repositioning around three foundational moves.

The Context Engine is an institutional memory connecting identity, resource, and decision data, enabling AI agents to act on genuine enterprise context rather than educated guesswork. Build Agent Skills is an SDK that lets developers build ServiceNow agents directly from tools they already use — including Claude Code — with automatic governance and audit trails from day one.

The new pricing model replaces Standard/Pro/Enterprise with Foundation/Advanced/Prime, AI included at every tier. Prime enables fully autonomous front-line role replacement. Critically: Anthropic/Claude has been selected as the default engine for Build Agent. ServiceNow — like The Cloud People — uses Claude internally for all 29,000 employees.

THREE SERVICENOW MOVES**01 · Context Engine — Institutional Memory**

Connects identity, resource, and decision data. AI agents act on genuine enterprise context, not guesswork.

02 · Build Agent Skills — Developer SDK

Build agents from Claude Code, Cursor, and familiar tools — governance and audit trails built in from day one.

03 · New Pricing Model — Foundation / Advanced / Prime

AI embedded at every tier. Prime enables fully autonomous front-line role replacement. The platform is now AI-native.

Fixed Price: From Exception to Norm

In the traditional consulting market, fixed price is associated with risk. With 10x Augmented AI, this dynamic changes entirely. Pre-built solution components reduce estimation uncertainty dramatically. Changes mid-project are cheaper to implement. Quality assurance is built in, not a separate step consuming hours.

A fixed price is no longer a risky special offer. It is the natural result of having a delivery model designed for efficiency and reuse — not to maximise billable hours. A major difference is that the pre-project, the scoping and the statement of work will take a much larger part of the project than earlier — and the production less.

NTT DATA / INSEAD · MARCH 2026

Organisations evolving toward "Service as a Software" models — roughly 70% standardised, 30% client-specific — will be best positioned to capture AI-era value. Partners who translate strategic objectives into AI-enabled operating models will command a more defensible market position.

— A NEW WAY TO BUY

The TCP WOW Solutions

One named consultant orchestrating a certified team of AI specialists — accountable for every deliverable.

WOW is not a chatbot or a generic AI assistant. It is a method and a certified delivery configuration — TCP consultants with the right grading and competence, orchestrating a team of specialized AI agents, each qualified by the TCP Lab on real ServiceNow engagements.

AI AGENTS — CERTIFIED BY TCP LAB

- User Story Builder: Converts requirements into ServiceNow user stories in your team's format.
- Update Set Reviewer: Reviews every update set for best practices, naming conventions, and compliance.
- Estimation Analyst: Produces structured estimation breakdowns from a scoped requirement.
- Workshop Prep: Builds agendas, pre-reads, and output templates for requirements workshops.
- Instance Analyst: Indexes your ServiceNow instance for context-aware advice — not generic suggestions.
- Solution Designer: Drafts first-pass solution designs from approved requirements.

OPTION 01 — REPLACE

Same scope as today
 Lower bill than your current T&M run-rate
 No timesheet reconciliation

Best when: stable, established TCP engagement.

OPTION 02 — EXPAND

Same hourly team you have today
 Pod adds extra delivery capacity
 For backlogs, quality uplift, parallel projects

Best when: your team is at capacity.

OPTION 03 — CO-DESIGN

Tailored to your instance and priorities
 Joint case study at end of pilot
 Preferred pricing for first 12 months

Best when: you want to be a first reference customer.

Lower monthly bill than equivalent pre-AI T&M scope — visible on your invoice.

TCP carries the AI investment: tools, licences, the AI Lab. You see the saving directly. We share the productivity gain explicitly — rather than waiting for you to demand it. We invest. You save. We both win.

THE ORIGIN

Built for a Future That's Now Arriving

The Cloud People was founded in December 2019 — not as a conventional consulting firm, but as a deliberate architectural bet on how the industry would eventually have to work.

The founding idea was precise: competence should be documented, structured, and made machine-readable from day one. Not as an afterthought. Not as a reporting exercise. As the foundation of the entire business model.

Every consultant's experience, skills, and certifications are thoroughly logged. Every deployment is captured automatically inside TCP's own intellectual property: the ServiceNow Resource Management tool Cloudia. Not manually entered. Not periodically updated. Automatically recorded as work happens — project by project, skill by skill, outcome by outcome.

"Cloudia knows more about a consultant than what they know about themselves."

This was not built to impress clients. It was built because the founders understood that structured knowledge, not headcount, would become the primary competitive asset in this market. What seemed like an operational detail in 2019 is now the engine behind the 10x delivery model.

Today, AI agents shadow consultants across engagements — learning patterns, absorbing context, building capability. They don't just assist. They take positions inside TCP's smart resourcing team, matching the right combination of human expertise and trained AI capacity to each new engagement automatically.

TCP was built for a future that is seeing the light these days. The infrastructure was laid when almost no one in the ServiceNow ecosystem was thinking in these terms. That head start is not accidental. It is the company.

THE TIMELINE**DECEMBER 2019****The Cloud People founded**

With a founding thesis: structure competence as machine-readable capital from day one.

ONGOING · EVERY PROJECT**Cloudia logs every deployment**

Skills, certifications, and project outcomes recorded automatically — not manually — across every consultant.

NOW**AI agents shadow consultants**

Agents take active positions in TCP's smart resourcing team — combining human expertise with trained AI capacity on every engagement.

THE RESULT**The 10x model becomes possible**

The infrastructure laid in 2019 is now the engine behind outcome-based delivery at scale.

WHY IT MATTERS NOW

The founding decision to make competence machine-readable — not as a later upgrade but as the company's core architecture — is precisely what makes the transition to AI-augmented delivery real at TCP, rather than a promise on a slide.

THE QUESTION CLIENTS ARE STARTING TO ASK**QUESTION**

“Are we going to stop buying consulting hours altogether?”

Bjørn Jarl · CEO, The Cloud People Group

Not overnight. But directionally, yes — and the direction is clearer than most people in this market want to admit.

The shift will happen step by step, in line with how our code banks grow and how our AI agents become more trained and capable. As that intellectual capital deepens, deliveries will become progressively more outcome-based and less hour-based. That is not a prediction. It is already happening in our most advanced engagements.

“The more projects we complete, the better our agents understand ServiceNow delivery. And the better they understand delivery, the less sense it makes to sell the hours it takes to run them.”

But here is the dynamic that surprises most people: as the shift to outcome-based delivery accelerates, hourly rates will actually go up — not down.

Why? Because a consultant who has access to TCP's full intellectual capital — the code banks, the certified agents, the accumulated knowledge from hundreds of projects — can now produce a multiple of what a standalone consultant could produce even two years ago. What used to require a team of four can now be delivered by one person with the right tools. That person's effective productivity is dramatically higher. And productivity commands a price.

The important caveat: the productivity gain only materialises if the partner takes genuine responsibility for production — not just for time spent. A consultant billing hours has no incentive to use the AI toolchain at full capacity. A partner delivering against a fixed outcome does. The model must change for the gain to become real.

QUESTION

“Most larger customers build internal teams to run their projects. They know their own business, their processes, their needs — and those needs often shift mid-project. That won't change. So how can you hand off production to an external partner?”

Bjørn Jarl · CEO, The Cloud People Group

This is exactly the right question — and the honest answer is: you don't hand off the knowledge. You never did. What changes is where the production happens.

What you will see is that the pre-project and planning phase demands a more accurate and thorough approach than before. Typically a "waterfall method" on the planning but with the ability to have an agile project as of the way we now are working with agents following the project. The requirements need to be sharper, the architecture more deliberate, the scope more carefully defined going in. That investment at the front end is real.

But the payoff is equally real: when adjustments come — and they always do — they become dramatically easier and faster to implement. Because the production engine already understands your instance, your patterns, your history. It doesn't start over. It adapts from a position of accumulated knowledge.

Internal teams remain exactly where they belong: owning the business logic, setting priorities, making decisions. TCP takes care of the production — not the judgement. That said, The Cloud People also acts as a trusted advisor in the outcome discussion before the scope is defined — helping clients ask the right questions before the work begins.

THE TRANSITION IN THREE POINTS

01

Step by step

The shift to outcome-based pricing will be gradual — paced by code bank depth and agent maturity, not by executive mandate.

02

Rates go up, not down

One consultant with TCP's intellectual capital produces what four did before. Higher productivity commands a higher rate.

03

Accountability unlocks the gain

The gain only materialises if the partner takes responsibility for the outcome. The model must change for the productivity to become real.

“If the partner is not willing to put their name on the outcome, they are probably not using the tools at the level that justifies a conversation about 10x.”

Bjørn Jarl · CEO, The Cloud People Group

NTT DATA / INSEAD · MARCH 2026

Traditional FTE-based pricing is at risk. Clients already expect efficiency gains from AI and demand lower prices — confirming that the pressure to move away from T&M; is not just a supplier-side choice but a client-side demand.

QUESTION

“Is fixed price a required condition for this model to work?”

Bjørn Jarl · CEO, The Cloud People Group

No. Fixed price is a signal of maturity, not a prerequisite. You can receive an estimated price and run the entire project on Time and Material — that is absolutely an option, but as there are no timesheets involved there are then milestones related instead.

What is not optional is this: TCP needs to steer the production. That means putting our intellectual capital, our AI agents, our local expertise, and our architects to work — and taking overall responsibility for what comes out the other end.

“The commercial model is flexible. The accountability model is not.”

The distinction matters because the productivity gain — the 10x — only materialises when TCP can orchestrate the full stack: code banks, certified agents, domain knowledge, and senior architecture. A T&M; arrangement where TCP simply supplies hours defeats the purpose. A T&M; arrangement where TCP owns the delivery engine is another matter entirely.

Fixed price or time and material: we can work in either model. What we cannot do is work in a way that prevents us from bringing everything we have built to bear on your project.

QUESTION

“How early should you involve the partner? And at what level — business goals, process transformation, or only once the scope is set?”

Bjørn Jarl · CEO, The Cloud People Group

As early as possible — and always before the scope is set. That is the honest answer, and it matters more in an AI-augmented delivery model than it ever did in a traditional one.

The conversation should start at the business level: what efficiency gains are you actually looking for? What transformation do you want to achieve — not just in the system, but in the way your organisation works? Those questions need to be on the table before anyone starts talking about features, modules, or user stories.

From there, the discussion moves down to the process level — mapping how work actually flows today, where the friction is, and what a better version looks like. This is where The Cloud People's role as an advisor is most valuable. We have seen how hundreds of organisations have solved similar problems. We know what works, what gets scoped incorrectly, and where the risks tend to hide before the project starts.

“The scope is not the starting point. The scope is the output of a good conversation about what you actually want to achieve.”

Involving a partner only at the point of scoping — after the business decisions have already been made — means the partner is optimising for the solution you have described, not necessarily the outcome you need. Those are sometimes the same thing. Often they are not.

Early involvement is not about selling more hours. In a fixed-price or outcome-based model, it is in TCP's direct interest to help clients scope correctly from the start. A well-scoped project is faster to deliver, easier to quality-assure, and produces a better result. The pre-project phase is not overhead — it is where the value is created.

— FOR CLIENTS

What You Should Be Asking For

Most ServiceNow clients evaluate partners on experience, certifications, and references. That's not wrong — but it is insufficient in 2026. The questions that actually determine what you get for your money are fundamentally different.

The next time you are in dialogue with a ServiceNow partner about a project, ask for a fixed price. Not as a provocation — but as a quality test. A partner with pre-built solution components, structured knowledge, and AI-augmented delivery will say yes.

A partner who always ends up explaining why fixed price isn't possible right here, on this project, with this scope — then you know something important about their delivery model. The market is ready for change. The only question is who is ready to demand it.

QUESTIONS TO ASK YOUR PARTNER

- Can you offer fixed prices as standard — not as an exception?
- What is the reuse ratio in a typical delivery?
- How is AI integrated in actual delivery, not just the sales presentation?
- What AI tools do your own consultants use every day?

RESEARCH PERSPECTIVE — INSEAD/NTT DATA

Traditional FTE-based pricing is at risk. Clients already expect efficiency gains from AI and demand lower prices — confirming that the pressure to move away from T&M; is not just a supplier choice but a client-side demand.

KEY TAKEAWAYS

What This Means For You

01

ASK FOR FIXED PRICE

The single most effective quality test for a partner. Those with AI-augmented workflows will say yes. Others will always find a reason why "this project is different."

02

INTERROGATE REUSE

What percentage comes from pre-built components? High reuse means lower cost, better quality, and less risk. A structured code bank is a hard asset — ask to see it.

03

DEMAND AI IN DELIVERY

Many partners feature AI in sales decks. Fewer embed it in actual delivery. Ask specifically: how does AI reduce the cost and time of our specific project?

04

THE SHIFT IS STRUCTURAL

INSEAD/NTT Data confirm FTE-based pricing is under pressure. Outcome-based models are not the future — they are the present.

05

SERVICENOW IS AI-NATIVE NOW

Foundation/Advanced/Prime tiers embed AI at every level. Anthropic/Claude is the default Build Agent engine. Partners not keeping up will fall behind.

06

CHOOSE TRANSFORMED PARTNERS

INSEAD: providers who have adopted AI in their own operations are significantly better positioned to guide client transformations. Ask what tools their consultants use daily.

07

THE TALENT MODEL IS CHANGING

The pyramid of junior staff executing repetitive work is shrinking. Orchestrators, validators, and domain experts are expanding. Better partners already deliver differently.

08

THE WINDOW IS SHORT

Partners who have invested in pre-built assets and AI agents have a head start. As the market normalises, this advantage will compress. Ask better questions now.

— CONCLUSION

“The question is not whether 10x delivery will become the norm. It will. The question is whether clients will demand it — or wait for someone to tell them they can.”

The market is ready for change. The transition is not a cliff. It is a curve. And The Cloud People are already on it.

“Enterprise value will be created by companies that challenge their workflow, orchestrate AI responsibly, and deliver measurable results. This is where CIOs should focus their work — we'll take care of the production.”

Bjørn Jarl · CEO, The Cloud People Group

Bjørn Jarl

CEO of The Cloud People Group

Pure-play ServiceNow partner · 400+ consultants · 9 countries

Part of NTT's global ServiceNow practice — 1,100 consultants

Reference: "Impact of AI on the IT industry: Talent transformation in AI era" — NTT Data & INSEAD · March 2026