PRACTITIONER REPORT · AUGUST 2026

The Perception Gap

Why most IT leaders overestimate their ITSM maturity, and what it costs them.

A report based on twelve years of independent, consultant-led ITSM assessment experience across UK enterprise IT.
Not a survey. Not a published benchmark study.

THE ITSM PEOPLE · itsmpeople.co.uk

CONTENTS

What's in this report

THE CENTRAL FINDING

Leadership sees the framework. Practitioners live the reality.

Almost every organisation we assess has a maturity gap between how leadership rates its service management and how the people doing the work experience it. It is not a failure of leadership. It is a structural consequence of distance from operations, and it has become considerably more expensive to ignore since organisations began committing serious budget to AI-enabled service management.

This report sets out the five patterns behind that gap, why they persist, what they cost when AI sits on top of them, and what to do about it.

ILLUSTRATIVE PROFILE

What we typically see

Based on what we consistently see across UK enterprise IT, most organisations sit between Initial and Defined. This is a representative profile, not a measured average across our own client base.

1.8 / 4.0

Typical overall score

Initial → Defined

Where most enterprise IT functions land

TYPICAL SCORE BY DISCIPLINE

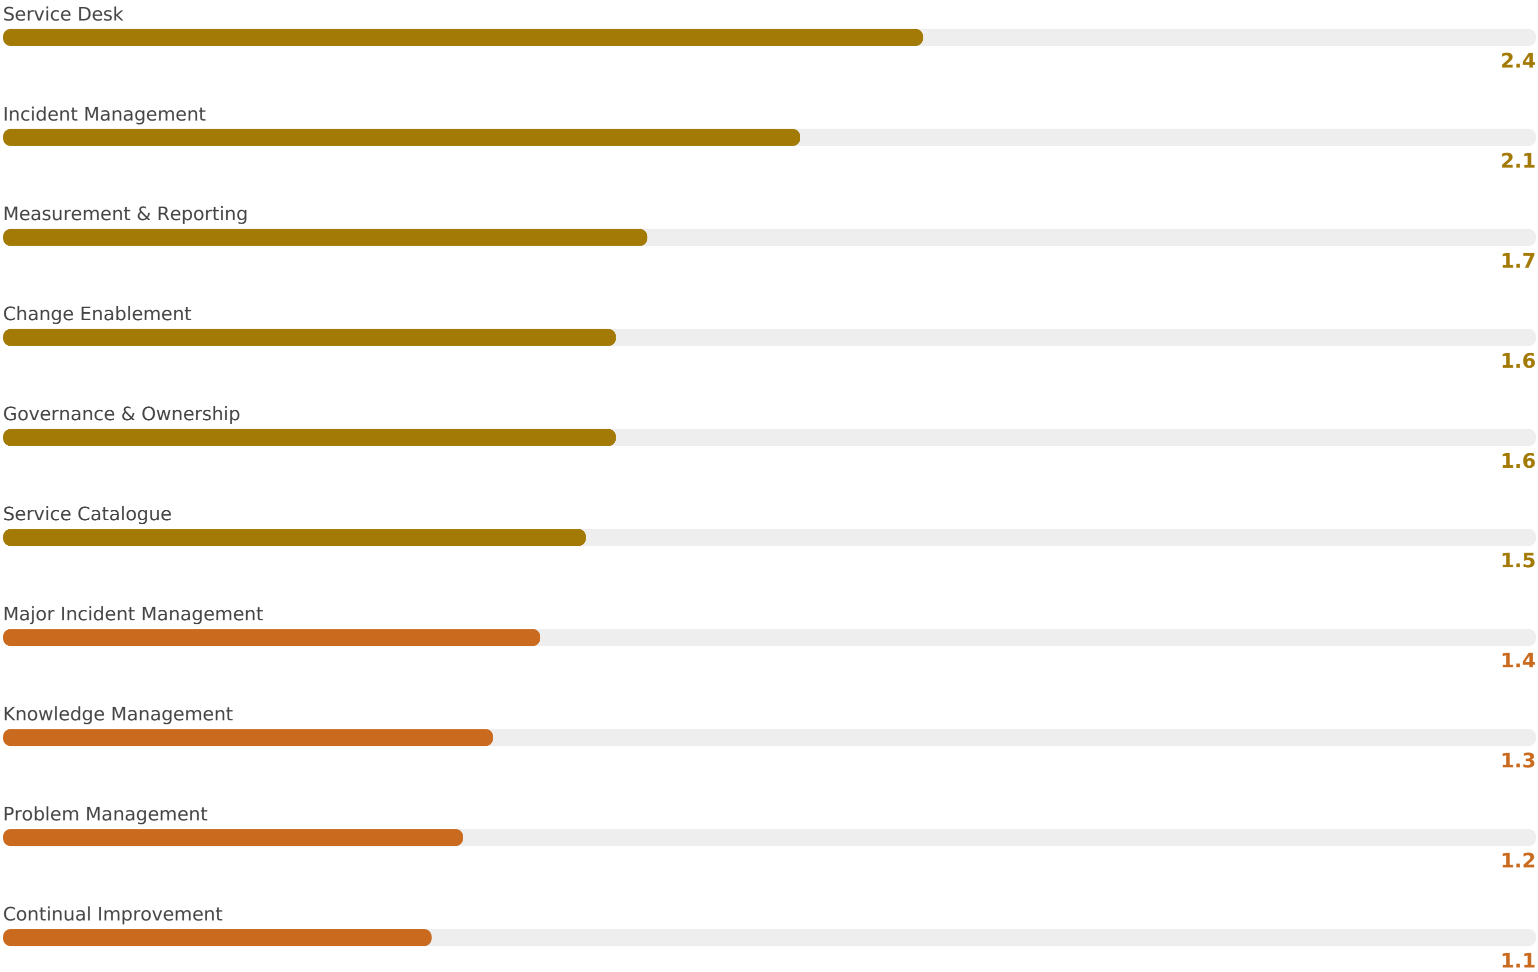

A representative profile based on our consulting experience across UK enterprise IT, not a published benchmark study. Individual organisations vary significantly. Scores use a 0-4 scale: 0 Absent, 1 Initial, 2 Defined, 3 Managed, 4 Optimised.

This is not a survey and it is not a benchmark study. It is a practitioner report. Since 2014 we have run independent, consultant-led ITSM assessments across UK enterprise IT organisations. What follows is what we consistently see. Where we describe a pattern as common, that is our professional judgement from twelve years of doing this work, not a statistic. Individual organisations vary significantly, and no client is identified. We are independent of every tool vendor and every framework body.

WHAT THE GAP TYPICALLY LOOKS LIKE

Leadership's self-rating **Managed**

Practitioner-reported reality **Initial – Defined**

A representative pattern from our consulting experience, not a measured average. The gap itself, not the exact positions, is what we consistently see.

1. The question we ask first

There is a question we ask at the start of almost every assessment.

We ask IT leadership to estimate how mature they believe their service management is. Then we ask the same question of the people doing the work.

The gap between those two answers is, without exception, one of the most revealing things we see.

Leadership consistently rates it higher. Not because they are wrong about what exists. The processes exist, the tools are in place, the policies are written. They are right about all of that. They are simply insulated from how those things function on a Tuesday afternoon when someone is off sick and a change has gone wrong.

They see the framework. Practitioners live the reality.

This is not a failure of leadership. It is the natural consequence of distance from operations, and it is structural rather than personal. But it has consequences, and those consequences have become considerably more expensive in the last two years.

2. Five patterns we see everywhere

01
Problem management in name only

02
Knowledge in people's heads

03
Inconsistent categorisation

04
Change bypassed for speed

05
Metrics that measure activity

Pattern one: Problem management in name only

Almost every organisation we assess has a problem management process. Very few have problem management.

Problem records are raised after major incidents. Investigations begin. Then they stall. Root causes are never formally confirmed, workarounds are logged but never communicated to the service desk, and the same incidents recur three months later. The problem record eventually closes, not because the problem was solved, but because it got old enough that nobody remembered to chase it.

Leadership sees a problem management practice with an owner and a process document. The service desk sees the same incidents coming back.

The tell: ask to see the last five problem records that were closed with a confirmed root cause and a change that resulted from it. The answer tells you everything.

Pattern two: Knowledge that lives in people's heads

Critical knowledge about systems, workarounds and edge cases exists in the minds of specific individuals. Everybody knows who those individuals are. Frequently there are only two or three of them.

When those people are on holiday, resolution times spike. When they leave, capability leaves with them. The knowledge base exists, was populated during a tool implementation, and has not been meaningfully reviewed since.

This is the single most reliable predictor of fragility we encounter, and it is almost never visible in reporting, because it does not show up in any standard metric until the person is gone.

The tell: what happens to your service quality during the summer holidays? If you know the answer without looking it up, that is your answer.

Pattern three: Inconsistent incident categorisation

The same type of incident is categorised differently depending on who logged it, which team handled it, or what time of day it arrived.

This looks like a housekeeping issue. It is not. It undermines every piece of trend data you produce, makes SLA reporting arguable, and renders problem management analysis close to worthless because you cannot reliably identify patterns in data that was inconsistently classified in the first place.

It has also become the single most consequential data quality issue in enterprise IT, for reasons we come to in section four.

The tell: pull a hundred recent incidents and look at how they were categorised. You will either find consistency or you will find catch-all categories absorbing everything ambiguous, and re-categorisation happening at volume after the fact.

Pattern four: Change management bypassed for speed

A formal change process exists and is presented as mature. Governance is in place. There is a CAB.

In practice, a significant proportion of changes bypass it entirely, particularly in development and infrastructure teams. Change-related incidents are then attributed to other causes, because attributing them accurately would require admitting the change was never registered.

The organisations where this is worst are rarely the ones with weak processes. They are the ones with processes so heavyweight that following them is genuinely slower than working around them. The people who comply are perceived as slower than the people who do not, so the bypassing continues and spreads.

Every unregistered change is a process failure. But when teams bypass a process at scale, the problem is usually the process, not the teams.

The tell: compare your change volume to your deployment volume. If they do not roughly reconcile, you have your answer.

Pattern five: Metrics that measure activity, not value

Dashboards show ticket volumes, response times and SLA compliance. They are produced monthly, reviewed in a governance forum, and filed.

Ask what any of it means for the business and the room goes quiet.

This matters more than it appears to, because it determines whether IT can defend its budget in language the business understands. An IT function that can only describe its performance in ticket volumes is an IT function that will be assessed on cost, because cost is the only dimension the business can evaluate.

The tell: could your CFO explain, from your reporting, whether IT service quality improved or declined last quarter, and what that was worth?

“

The problem record closes, not because the problem was solved, but because it got old enough that nobody remembered to chase it.

ON PATTERN ONE · PROBLEM MANAGEMENT IN NAME ONLY

3. Why these patterns persist

Discipline scores tell you what is broken. They do not tell you why it stays broken. Four cross-cutting themes appear in almost every organisation we assess, and addressing them is what separates an improvement programme that works from one that produces a report nobody implements.

Ownership without accountability. Somebody is named as the owner of Incident Management. That person has no dedicated time, no authority to enforce standards, and no quality targets. Having an owner and having effective ownership are not the same thing, and the gap between them is where most improvement programmes quietly fail.

Silos that predate the process. IT operates in functional towers. Incident, change and problem management are owned by different teams who rarely collaborate. Each team's process is locally sensible and collectively incoherent.

Culture that treats process as bureaucracy. Workarounds are celebrated as pragmatism. This is rational behaviour: practitioners who have watched three improvement initiatives come and go without lasting change have earned their scepticism. Resistance to ITSM improvement is almost never irrational, and treating it as something to be managed rather than understood is the fastest way to guarantee it continues.

The report that gathered dust. A previous assessment produced a detailed roadmap. It was presented, approved, filed and never implemented. Two years later the same gaps reappear. This is the single most common reason organisations hesitate before commissioning another one, and it is a reasonable hesitation.

4. Why the gap has become expensive

The perception gap has always existed. What has changed is the cost of getting it wrong.

Organisations are making significant investments in AI-enabled service management: intelligent automation, virtual agents and self-service. Almost every one of those investments quietly assumes foundations that are more capable than most organisations realise they need to be.

AI amplifies what is already there. It does not repair it.

WHAT AI DOES TO A WEAK FOUNDATION

Inconsistent process

Categorisation, knowledge, ownership all vary by team

→

Amplified at machine speed

AI learns the inconsistency and repeats it, faster and with more confidence

This is why the organisations getting genuine value from AI in service management fixed their foundations first. Automating a poor process does not fix it. It delivers the same poor outcome, faster.

- **Inconsistent incident categorisation** produces unreliable AI triage. The model learns patterns from data that was classified inconsistently, so its recommendations are frequently wrong, agents override it, and within weeks it becomes something people work around rather than with.
- **An unmaintained knowledge base** produces confident, incorrect answers at scale. An AI that surfaces a workaround that no longer works does not merely fail to help. It actively damages user confidence in the service desk.
- **An out-of-date service catalogue** means AI self-service fails at first contact. Users who encounter an assistant that cannot help with their actual request do not try again. They pick up the phone.
- **Absent problem management** means predictive analytics has nothing coherent to detect. Recurring incidents that were never investigated look identical to emerging new problems.

Automating a poor process delivers a poor outcome faster. That warning is not new. What is new is the size of the budgets now being committed on the assumption that it does not apply.

The organisations getting genuine value from AI in service management share one characteristic: they fixed their foundations first, then deployed into an environment where the inputs were reliable. That sequence takes longer to show results. The results, when they come, are real.

“

An organisation that believes it is Managed and operates as Initial has a different, harder problem than an organisation that is honestly Initial and knows it.

ON WHY THE GAP IS MORE DANGEROUS THAN THE SCORE

5. Why we built a platform to close it

The perception gap is not a new observation. Every experienced consultant has seen it. What is new is what we decided to do about it.

Standard assessment methodology asks one set of questions and gives one group the microphone, usually leadership, because leadership commissions the assessment and leadership is easiest to access. The resulting report reflects what the organisation believes about itself. It rarely captures what the organisation actually does, because the people who could tell you that were never systematically asked.

We built the Clarity Assessment™ platform to close that gap by design, not by accident.

Every assessment we run uses stakeholder-specific question sets, scored independently against evidence rather than opinion. Where views diverge sharply, the platform surfaces that divergence explicitly, as a finding in its own right rather than noise to be averaged away.

LEADERSHIP	IT MANAGEMENT	PRACTITIONERS
Strategic view. What they believe is in place.	Operational view. How it is run day to day.	Ground truth. What actually happens.

↓ Every discipline scored independently against evidence, not opinion ↓

That divergence is frequently the most important line in the entire report. An organisation that believes it is Managed and operates as Initial has a different, harder problem than an organisation that is honestly Initial and knows it. The first organisation is not planning to fix anything, because nothing looks broken from where they sit. The second is already halfway to a solution, because they have stopped arguing about the diagnosis.

We are independent of every tool vendor and every framework body, and now, deliberately, we are independent of the frameworks themselves. Our question set was built from twelve years of running these assessments, not adapted from someone else's model, so it asks what we have learned actually matters rather than what a framework says should matter. Where ITIL, SIAM or ISO 20000 provide useful language, we use it. Where our own experience diverges from the textbook, we follow the experience.

The platform does not tell you what you want to hear. It tells you where your leadership team's view and your practitioners' reality stop agreeing, and by how much. In our experience, that single number, more than any individual discipline score, predicts whether an improvement programme will actually change anything.

The five patterns, at a glance

A quick reference for the diagnostic conversation. Each one has a test you can run in under an hour.

01 Problem management in name only

The tell: ask to see the last five problem records closed with a confirmed root cause and a resulting change.

02 Knowledge that lives in people's heads

The tell: what happens to service quality during the summer holidays? If you know without checking, that is the answer.

03 Inconsistent incident categorisation

The tell: pull a hundred recent incidents and look at how consistently they were categorised.

04 Change management bypassed for speed

The tell: compare change volume to deployment volume. If they do not roughly reconcile, you have your answer.

05 Metrics that measure activity, not value

The tell: could your CFO explain, from your reporting, whether service quality improved last quarter, and what that was worth?

---

6. What to do about it

1. Ask the two questions

Ask your leadership team to score your service management maturity. Ask your practitioners the same question. Do it independently and compare. You do not need a consultant to run this, and the gap itself is diagnostic before anyone scores anything formally.

2. Look at the tells, not the documentation

Every pattern in this report has a test that takes under an hour and cannot be answered from a process document. Run them.

3. Fix categorisation before you buy AI

It is unglamorous, it is the least interesting item on any roadmap, and it is the single highest-return preparation you can do. Consistent categorisation is a precondition for every AI capability worth having.

4. Be honest about the last roadmap

If a previous assessment produced recommendations that were never implemented, understand why before commissioning another. The reason is almost never that the recommendations were wrong.

5. Get an outside view, from someone with nothing to sell you

The perception gap exists because proximity makes objectivity difficult. That is not a character flaw, it is a structural feature of being close to the work, and the only reliable correction is an independent one.

WHERE DOES YOUR ORGANISATION ACTUALLY STAND?

Let's find out, honestly.

A free 30-minute strategy call, no obligation. We will tell you honestly whether an assessment would help, and what it would involve. No sales pitch: just a direct conversation with someone who has run this diagnostic hundreds of times.

itsmpeople.co.uk/contact
hello@itsmpeople.co.uk
+44 1564 340327

© The ITSM People 2026. The ITSM People® is a registered trade mark of The ITSM People Ltd, Company 09098884.
This report is based on consulting experience across UK enterprise IT and is not a survey or published benchmark study. Individual organisations vary significantly and no client is identified.