



QHappySignals

Why CSAT is not enough

and what to measure when your score looks fine, but
your people are feeling frustrated.

What's inside

Introduction: CSAT is a comforting number

Where CSAT came from and what it was never built to do

“Resolved” doesn’t mean resolved

The number is easy to move, but the experience isn’t

The people you never hear from

A score with no context can’t be acted on

What “enough” actually looks like

From a number to a habit

Checklist to get started

Keep the smile, add the substance

Page 3

Page 4

Page 5

Page 6

Page 7

Page 8

Page 9

Page 10

Page 11

Page 11



Introduction: CSAT is a comforting number

Picture the monthly service review where the SLA dashboard is a wall of green, tickets are closing on time, and CSAT is 4.5 out of 5.

Now, picture the same week from a different desk. What if an employee's VPN dropped right before a client call? An access request that should have taken minutes took three days? Or a ticket was marked "resolved," but nobody actually fixed the issue.

The dashboard is green, and the experience is red, simultaneously. We call it the watermelon effect, and CSAT is one of the reasons it hides so well.

CSAT is incomplete.

A single, easily-moved number that tells you someone clicked something positive, but almost nothing about what happened, why it happened, or what you should fix next.

This e-book is about closing that gap: keeping the parts of CSAT that work, and adding the parts that have been missing.



4.5



Where CSAT came from and what it was never built to do

CSAT became the default for a reasonable reason. For years, it was the one number that tried to represent the human side of IT, sitting alongside a portfolio of operational metrics that otherwise measured only how many and how fast.

In a world of ticket volumes and response times, a satisfaction score felt like a welcome dose of the customer's voice, offering insight into how users are actually feeling.

It earned its place on the slide. It's a number, it trends nicely, and leadership understands it instantly. That instant legibility is exactly why it stuck, and also, quietly, why it became a problem.

CSAT was designed to summarize a reaction, not to explain a reality.

It can tell you that a person felt okay (or didn't) about a single interaction. It was never built to tell you which

service failed, where the delay happened, how much working time the employee lost, or what to change so it doesn't happen again.

We've spent two decades asking one metric to answer questions it was never designed to hear.

Everything that follows is really just one idea, looked at from different angles: a score is not the same as understanding.

“Resolved” doesn’t mean resolved

Experience doesn’t live in a dashboard. It happens in very real, very specific moments, and most of those moments never make it into the metric. From an ITSM perspective, every interaction in which the ticket is resolved can appear to be a success.

The ticket was logged, handled, and closed inside SLA. From the employee’s point of view the issue is not resolved. Closing a ticket on time does not automatically mean someone can get back to work, and that gap is precisely where trust is won or lost.

The practitioners feel this more sharply than anyone. When Invesco’s experience leaders, Donna Xanthidis and Sharon Aggarwal, discussed the toughest feedback they’ve received, it wasn’t about broken hardware.

It was comments that boil down to “I wasn’t listened to” or “I was ignored.” The technical fault gets fixed; the experience still collapses at the last hurdle, because nobody kept the person informed.

“‘Resolved’ isn’t the same as ‘felt resolved.’ Even when the issue is fixed, the experience can collapse if people aren’t kept in the loop.”

- Donna Xanthidis & Sharon Aggarwal: Beyond CSAT.
[Read more here.](#)

A closed ticket measures the process. It says nothing about the moment an employee quietly decides “IT doesn’t care,” or “I won’t bother raising this next time,” or “I’ll just find a workaround.”

Those decisions are the real output of IT support and CSAT, and as most teams run it, they can’t see it.



“

“‘Resolved’ isn’t the same as ‘felt resolved.’ Even when the issue is fixed, the experience can collapse if people aren’t kept in the loop.”

- Sharon Aggarwal, Invesco

The number is easy to move, but the experience isn't

What should we improve to make you happier?

You may select more than one option

- ☐ Service personnel's skills
- ☐ I wasn't informed at
- ☐ Service personnel's attitude
- ☐ I had to explain my
- ☒ It was not clear where to get what I needed
- ☐ Instructions were h
- ☐ I didn't get what i requested
- ☐ Service was slow

Estimate the work time you lost

45 minutes



The uncomfortable secret of CSAT is that it's remarkably easy to improve the score without improving anything an employee would actually notice. You don't need to fix the service. You just need to make it easier to give a good rating.

Over time, surveys drift in exactly that direction. Stars become smiley faces. Five questions become one. Friction gets stripped out "to improve response rates."

Each tweak makes the data easier to explain but less useful for learning. At some point, CSAT stops being feedback and starts being reassurance.

When the metric improves, but nothing else does

A higher CSAT score implies progress, even when the underlying service hasn't changed at all.

That's how good teams get quietly misled; they celebrate a trend line while the lived experience holds perfectly still. And if the score went up, but the only thing that changed was the survey, you haven't learned a single thing.

A few clicks decide the story for everyone

CSAT usually represents a very small slice of respondents, and not a random slice. The people who answer tend to be the ones who had the energy to click, not the ones who were most blocked or most frustrated.

So, you get response bias baked in: the very happy and the very angry, with the exhausted middle staying silent. A few percent of users end up writing the narrative for everyone else.

Put those together, and you have a metric that's easy to game, easy to misread, and built for reporting rather than learning.

As Invesco's Sharon Aggarwal put it, CSAT is "not bad", it's "just incomplete". "Like a signal, but a faint one."

The job isn't to throw the signal away. It's to make it loud enough to act on!

The people you never hear from

Post-contact CSAT has a structural blind spot that no amount of survey-tuning can fix: it only ever reaches people who contacted the service desk. Everyone who hit a problem and didn't raise a ticket (because they've given up, found a workaround, or assumed nothing would change) is invisible to it. And that group is not small.

62% of end-users say "I avoid my service desk."

58% say they're "living with ongoing problems the service desk can't fix."

Those figures come from Forrester's ["State of the Service Desk"](#) report, which found that 67% of people who contacted the desk were satisfied.

Read those numbers together, and the picture is unmistakable: you can hit your satisfaction target with the people who show up, while a majority quietly route around you.


If CSAT were an accurate read on what employees think of IT support, there wouldn't be so many people avoiding it.

This is why a "good" score can be genuinely dangerous. It doesn't just miss unhappy users, it actively reassures you that they don't exist.

The friction is still there. It's just showing up somewhere your metric can't see it: in escalations, in shadow IT, in lost time, or slowly declining trust.

68%
of end-users avoid
their service desk

58%
Have ongoing problems
the service desk can't fix



A score with no context can't be acted on

Even when CSAT is honest, it runs into a final wall: a single number can't tell you what to do. A standalone score cannot explain what actually caused the frustration, where the delay happened, which service or vendor let people down, or why work couldn't continue.

"CSAT comes too late to the party, and it's too generic to drive meaningful change."

- Shaun Brown, Director of Global Service Delivery, GEA.

[Read more here.](#)

Experience only becomes meaningful when it's connected to what actually happened in your ITSM reality.

Which ticket was this about? Which service was involved? How long did the employee wait, and for what? Was this their first attempt, or their fourth?

Without that context, you're left guessing, and guessing tends to default to the loudest voice in the room rather than the clearest signal in the data.

That last point matters more than it sounds. When prioritization runs on assumptions and seniority, a single complaint that reaches an executive can derail a roadmap. At the same time, a widespread, quiet frustration gets ignored because nobody important has felt it.

Context lets you respond at the right scale.

What “enough” actually looks like

So if CSAT alone isn't enough, what is? The answer is to keep measuring how people feel and connect it to what actually happened, as well as what it cost them in time. That's the shift from a comforting number to a system you can act on.

Add the half of the story CSAT leaves out: lost time

MTTR tells you how long a ticket took. It says nothing about how much working time the employee actually lost.

Capturing perceived lost time alongside sentiment changes the whole conversation: now you can prioritize by productivity impact rather than ticket volume, and quantify value in language the business already speaks.

[Read more: Ticket reassignment is the biggest cause of productivity loss](#)

Connect every piece of feedback to operational reality

Instead of a free-floating score, tie each response to the actual incident, request, or interaction that prompted it, such as the service, category, channel, location, assignment group, vendor, and resolution path.

Suddenly, a number becomes a story about what happened and why, and you can drill straight to where to act next.

[Read more: From SLAs to XLAs](#)

Stop averages from lying

A single company-wide CSAT number hides pockets of genuinely poor experience.

Segmenting by role, location, channel, or support behavior surfaces the hidden pain fast.

This is the kind of insight that turns up when you look. End-users are consistently unhappier with their devices and core applications than with almost any other part of IT.

The most-used support channel (the self-service portal) often delivers the worst experience of all.

And a huge share of total lost time concentrates in a tiny minority of tickets, which means the right 13% of incidents matter far more than the comfortable average suggests.

[Read more: Support your IT support channels](#)

Don't reinvent the survey, and don't read 10,000 comments by hand

Validated, research-backed surveys give you a trustworthy baseline and a definition of “what good looks like” without a survey-design project.

[A global benchmark](#) tells you whether your trend is actually good or just better than last month.

And because the truth usually lives in open comments, AI that summarizes thousands of responses into clear themes turns “what's happening” and “what to do about it” into an actionable plan in minutes.

From a number to a habit

Time for a customer case!



At [Invesco](#), “You said, we did” became a mantra repeated across the technology group. The trick wasn’t the slogan; it was how they used it. Instead of announcing “we fixed the thing,” they anchored every update in human impact.

Things like what it did to someone’s day, how it affected productivity, and what friction it removed.

That single change turns an IT announcement into proof that the employee’s voice actually mattered. It closes the loop, and doing so builds trust.


Most feedback, once you actually collect and read it, is positive, and naming individuals in that positive feedback, especially on the service desk, becomes immediate recognition rather than a stick to beat anyone with.

People see their work land in the hands of a real human. That changes team energy fast.

When you can see where friction genuinely concentrates, from the business’s perspective, you stop debating perceptions and start fixing the things that matter most. A single loud complaint no longer outranks a thousand quiet ones!

That’s the real “beyond CSAT” story. Not replacing one score with another score, but building a way of working where IT can confidently answer the three questions that actually matter:

What are people really experiencing? Where is productivity being lost? What should we fix first, and why?



Checklist to get started

- **Start. Don't wait for perfect.** You can tweak surveys and wording as you learn.
- **Treat feedback like a loop, not a scoreboard.** If you ask, you have to act, or employees will stop believing it matters.
- **Make communication part of the definition of "resolved."** The last mile is where trust is won or lost.
- **Use experience to prioritize value, not noise.** Move away from internal politics and toward visible friction points.
- **Celebrate what works.** Most feedback is positive when you actually listen, and that can change team energy fast.

Keep the smile, add the substance

CSAT isn't dead, and it doesn't need a funeral. It needs a promotion. On its own, it's a faint signal: easy to move, easy to misread, blind to the people who never contact you, and unable to tell you what to fix. On its own, it's a faint signal: easy to move, easy to misread, blind to the people who never contact you, and unable to tell you what to fix.

But connected to operational reality, paired with lost time, segmented properly, and acted on out loud, that same signal becomes the engine of continual improvement!

[Read more: Continual Improvement](#)

If your score improves but the experience doesn't, you've learned nothing. If your score is grounded in real incidents, real requests, and real outcomes, you've got something you can finally trust and act on.

[Read more and take our 2-minute CSAT self-assessment to see where your organization stands on CSAT maturity.](#)



Ready to drive IT Experience Management in your organization?

If you're ready to move beyond SLAs and start managing IT by outcomes that actually matter to your business, HappySignals gives you the platform to do it.

Visit us at happysignals.com to find out more!

HappySignals is the leading SaaS company for IT Experience Management, empowering enterprises to change their culture to be more open, outcome-focused, and data-driven. The HappySignals Platform enables IT leaders to get a real-time understanding of the experiences they are delivering to end-users. Using HappySignals, customers can make informed decisions that increase employee happiness and productivity. HappySignals discovers the experiences of millions of employees in 130 countries. Customers have been able to make employees happier and increase productivity by 26%.

