



IT services run on 36% failure demand.

The State of Service Flow

The first published measure of failure demand and flow efficiency in IT services. 400 responses to six questions, and the two numbers no service report contains.

IN THIS REPORT

03	The findings in one page
04	Why nobody had the number
05	Read this before the numbers
06	Finding 1 · A third of the demand is self-inflicted
07	Finding 2 · The work waits nine days in ten
08	Finding 3 · "We do not know" was the most common answer
09	Finding 4 · Plan the team to 90 per cent and the work waits
10	Finding 5 · Providers count it, internal teams do not
11	Finding 6 · Fixing one does not fix the other
12	By sector and by size
13	How the numbers were made
14	Work out your own
15	Cite it, share it, add to it

WHAT THIS IS

The Service Flow Benchmark is an open, unsponsored attempt to measure two things IT services never have: how much of the demand exists because something already went wrong, and how much of the time a piece of work is actually being worked on. Six banded questions, anonymous, published free in aggregate only.

This report covers the first 400 responses. A flawed number, for reasons stated on page 5 before anyone else gets to state them. Also the only one.

CITE IT

Boyer, M. (2026) *The State of Service Flow: failure demand and flow efficiency in IT services*.
markboyer.co.uk/benchmark/findings/

Licensed CC BY 4.0. Reuse anything in it, including the charts, with attribution and a link.

400 responses · gathered 1 to 29 August 2026
Published September 2026 · free under CC BY 4.0
Published at
markboyer.co.uk/benchmark/findings/

TWO TERMS, PLAINLY

Failure demand

Demand that exists only because something already went wrong: chasing an update, reporting the same fault twice, asking for something that should never have needed asking. Not the fault itself, the contact it generated.

Flow efficiency

The share of the time you hold a piece of work in which it is actually being worked on: minutes of genuine effort divided by the days between raised and done. Utilisation tells you how busy your people are. Flow efficiency tells you how long their work waits.

FINDING ONE · FAILURE DEMAND

36%

of IT service demand exists because something already went wrong.

Band-midpoint estimate among the 311 respondents who could answer. Most common answer: 25 to 40%. Method, page 13.

Read together, the six findings say one thing. The profession measures how busy it is with great precision, and how much of that busyness it caused with none at all.

A third of the demand is self-inflicted. The work waits nine days in ten. Nearly two in five service organisations cannot state either number, and the ones that can are, mostly, the ones that invoice for it.

10%

Flow efficiency. For every 10 days a piece of work is open, it is being worked on for one. Estimate among the 251 who could answer.

43%

of those who measure flow efficiency put it **under 5 per cent**. 68 per cent put it under 10.

39%

could not state at least one of the two numbers. **21 per cent could state neither.** "We do not know" was the most common answer on flow efficiency.

64%

of teams planned to **over 90 per cent utilisation** report flow efficiency under 5 per cent. Planned under 70 per cent: 22 per cent.

12% v 27%

Providers versus **internal IT teams** who cannot state their failure demand. They have about the same amount. The providers count it.

10 pts

The entire spread in poor flow across every failure demand band, from under 10 per cent to over 60. **The two numbers do not move together.**

Every figure on this page is self-reported, banded and unverified. Page 5 says what that means, and page 13 says how each one was worked out.

WHY NOBODY HAD THE NUMBER

Thirty years of folklore and not one dataset.

Failure demand was named by John Seddon in the early nineties: demand that exists only because something already went wrong. Chasing an update. Reporting the same fault twice. Asking for something that should never have needed asking. Not the fault itself, the contact it generated.

The figure everyone quotes for it, forty to sixty per cent of demand in conventional service organisations, comes from consulting observation rather than published research. It has been doing service management duty ever since, and nobody has checked whether it transfers. The claim that up to eighty per cent of demand on public services is failure demand has the same problem, and it is a claim about a different population anyway. One policing study, in Gloucestershire, put it at thirty-two per cent. At least that one has a dataset.

Flow efficiency has a similar biography. The measure comes out of lean: the share of the time you hold a piece of work in which it is actually being worked on. The figures most often repeated for it, low single digits, trace back to practitioner experience across engagements rather than anything anyone can inspect.

For IT services there was nothing at all. No published failure demand figure for a service desk. No flow efficiency baseline for service management. The profession has spent three decades measuring how busy its people are and almost no time measuring how long its work waits.

This report is the first attempt to put a number on both. A flawed number, for the reasons on the next page. Also, for now, the only one.

WHAT GETS MEASURED INSTEAD

Ticket volume. First contact resolution. CSAT. Cost per ticket. SLA attainment, to two decimal places.

Everything except whether the work needed doing, and how long it spent waiting to be done.

WHERE THE FOLKLORE COMES FROM

FIGURE	POPULATION	BASIS
40 to 60%	Conventional service organisations	Consulting observation
Up to 80%	Public services	Claimed, unsourced
32%	Policing, one force	A study with a dataset
None	IT services	Until this one

READ THIS BEFORE THE NUMBERS

Every vendor survey skips this page.

It is the reason to trust this report more than theirs, and the reason to trust it less than you would like.

Self-reported, not instrumented

Nobody counted anything. 400 people estimated, from memory, in ranges, on a phone, in about ninety seconds.

Self-selected

People who care about failure demand answered a survey about failure demand. The population that never thinks about it is under-represented here, and it is probably worse.

Unverified

No submission was checked against anything. The three questions about who you are were used only to cut the data, never to identify anyone.

Banded, so the aggregate is a range

The single figures in this report are band-midpoint estimates. Moving the top band's assumed midpoint by five points moves the headline from 35.2 to 36.6. Read 36 as "the mid thirties", not as a decimal.

Under-reported

People scoring their own failure demand almost certainly score it low, because the number is unflattering. 36 per cent is more likely a floor than a ceiling.

The don't-know rate is a finding, not a gap

Nobody was forced to guess. The 22 per cent who could not state their failure demand and the 37 per cent who could not state their flow efficiency are the most reliable numbers in the report, because they are the only ones nobody had to estimate.

THE SAMPLE

400 responses, gathered through markboyer.co.uk/benchmark between 1 to 29 August 2026, after a keynote and a LinkedIn post. 41 per cent serve customers under contract, 46 per cent serve colleagues inside one organisation, 13 per cent do both. Nine sectors and five size bands, none smaller than 23 responses. Aggregate only, throughout.

WHAT IS NOT PUBLISHED


No cut with fewer than 20 responses. That removes the 90 respondents who do not plan to a utilisation figure from the utilisation analysis, because only 12 of them could state a flow efficiency. Sector and size cuts are published with their n, and should be read as indicative.

FINDING ONE · FAILURE DEMAND

A third of the demand is self-inflicted.

Among the 311 respondents who could put a figure on it, the estimate is 36 per cent. The most common answer is 25 to 40%, and two in three put it at a quarter or more.

How much of your demand exists because something already went wrong? Share of all 400 responses



Of those who know, 67 per cent put failure demand at a quarter or more of everything coming in. 39 per cent put it at 40 per cent or higher, which is where the consulting folklore starts. 10 per cent say it is under 10 per cent. Either they are very good, or they have not looked.

Read it as a floor. The people answering were scoring their own operation, and it is not a number anybody wants to be high. Read it as a range too: the honest statement is that the typical IT service organisation sits somewhere in the 25 to 40% band, with a long tail above it.

At 36 per cent, roughly one contact in three exists to chase, repeat or correct an earlier one. If your model bills per ticket, that is revenue. If it does not, it is pure cost with nothing on the other side. Finding 5 is about which of those you are.

ANSWER	N	ALL	WHO KNOW
Under 10%	30	7.5%	9.6%
10 to 25%	72	18.0%	23.2%
25 to 40%	89	22.2%	28.6%
40 to 60%	77	19.2%	24.8%
Over 60%	43	10.8%	13.8%
We do not know	89	22.2%	


n = 400. "Who know" is the share among the 311 who did not answer "we do not know".

FINDING TWO · FLOW EFFICIENCY

The work waits nine days in ten.

Among the 251 who could answer, the estimate is 10 per cent. For every 10 days a piece of work is open, it is being worked on for one. 43 per cent put it under 5.

Of the time you hold a piece of work, how much of it is actually being worked on? Share of all 400 responses



Flow efficiency is the share of the time you hold a piece of work in which it is actually being worked on: minutes of genuine effort divided by the days between raised and done. Nobody plans to it, and it shows. Of the 251 who could answer, 43 per cent put it under 5 per cent and 68 per cent under 10. 15 per cent put it under 1 per cent, which is a request open for a full working week and touched for about 24 minutes of it.

10 per cent report over 25 per cent. There is no way of knowing from this data whether they measured it or are optimists, and the two are not mutually exclusive.

Utilisation tells you how busy the people are. This number tells you how long their work waits. Finding 4 is about what one does to the other.

ANSWER	N	ALL	WHO KNOW
Under 1%	37	9.2%	14.7%
1 to 5%	71	17.8%	28.3%
5 to 10%	63	15.8%	25.1%
10 to 25%	54	13.5%	21.5%
Over 25%	26	6.5%	10.4%
We do not know	149	37.2%	

n = 400. The don't-know rate is highest here: 37 per cent, against 22 per cent for failure demand. The SLA clock measures elapsed time. Timesheets, where they exist, measure effort. Nobody divides one by the other.

FINDING THREE • WHAT NOBODY MEASURES

"We do not know" was the most common answer.

On flow efficiency it was the most common answer by a distance.

On failure demand it tied for first. And not knowing one number very nearly guarantees not knowing the other.

22%

cannot state their
failure demand

37%

cannot state their flow
efficiency


21%

can state neither

15%

have none of the three
numbers, planned
utilisation included

Cannot state flow efficiency...



If you cannot state your failure demand, there is a 93 per cent chance you cannot state your flow efficiency either. If you can, that falls to 21 per cent.

Not knowing is not incompetence. Nothing on a standard dashboard would reveal either number. A ticket count does not separate the demand you exist to handle from the demand you created earlier. An SLA clock measures elapsed time, not touch time. The dashboard says the desk is busy. It does not say why, and it does not say how long anything waited.

What the data adds is that measurement is a habit rather than a metric. Organisations either look at the work or they do not, and the ones that do not tend not to plan to a utilisation figure either: 60 respondents, 15 per cent, had none of the three numbers.


Which makes the don't-know rate the most reliable finding in this report. It is the one answer nobody had to estimate.

FINDING FOUR • UTILISATION AND FLOW

Plan a team to 90% and the work waits.

Among teams planned to over 90 per cent utilisation, 64 per cent report flow efficiency under 5 per cent. Planned under 70, it is 22. Kingman said so in 1961.

Share reporting flow efficiency under 5 per cent, by planned utilisation



Kingman's formula says the time work spends waiting rises sharply, not steadily, as utilisation approaches 100 per cent. The last few points of utilisation cost more waiting than all the rest put together. It is sixty-five years old and it is still not on the capacity planning slide.

The benchmark reproduces it in self-reported data. The share reporting flow efficiency under 5 per cent roughly triples from the coolest band to the hottest, and the estimate falls from 14 to 6 per cent along the way. Utilisation is the number capacity planning is proudest of. It is also the one that makes the queue.

One group is missing from the chart. 90 respondents do not plan to a utilisation figure at all, and only 12 of them could state a flow efficiency, which is below the publishing floor. Not planning to a number and not measuring the flow travel together.



PLANNED TO	N	KNOW FLOW	UNDER 5%	ESTIMATE
Under 70%	39	32	21.9%	13.6%
70 to 80%	87	68	26.5%	13.6%
80 to 90%	93	72	43.1%	8.7%
Over 90%	91	67	64.2%	5.7%
No figure planned	90	12	n/p	n/p

n/p: not published, fewer than 20 could answer. "Under 5%" and the estimate are among those who could state a flow efficiency.

FINDING FIVE · WHO YOU SERVE

Providers count it. Internal teams do not.

12 per cent of providers cannot state their failure demand. For internal IT teams it is 27 per cent, and for those who serve both, 39. The estimates barely differ. The counting does.



Providers are more than twice as likely as internal teams to know their failure demand. It is not that they have less of it: the provider estimate is 37 per cent against 34 for internal teams, and 42 per cent of providers put it at 40 per cent or higher. They have roughly the same amount. They categorise it.

The commercial reading is the uncomfortable one. If your model bills per ticket, failure demand is revenue, and nobody leaves revenue uncounted. For an internal team it is pure cost with nothing on the other side, and nobody counts a cost that never appears as a line. The incentive to measure runs exactly opposite to the incentive to eliminate.

The gap closes on flow efficiency, where 34 to 43 per cent of every group cannot answer. Nobody invoices for waiting, so nobody measures it. The 51 who serve both customers and colleagues know least of all. Two masters, measured for neither.

SERVE	N	FD EST.	FD 40%+	FLOW EST.
Providers	165	36.9%	41.8%	8.9%
Internal IT	184	34.1%	35.1%	10.1%
Both	51	38.7%	38.7%	10.9%


Estimates are band-midpoint means among those who could answer. "FD 40%+" is the share of those who know who put failure demand at 40 per cent or more.

FINDING SIX • FAILURE DEMAND AND FLOW

Fixing one does not fix the other.

Across every failure demand band, 36 to 46 per cent report flow efficiency under 5 per cent. A 10-point spread. The two numbers do not move together.

Share reporting flow efficiency under 5 per cent, by failure demand band. The 245 who could state both



The intuition is that they travel together: an organisation drowning in failure demand should wait longer, because the repeat contacts clog the queue. In this data they do not. Organisations reporting under 10 per cent failure demand wait about as long as those reporting over 60. A flat line is the least exciting chart in the report and possibly the most useful.

Failure demand is about what comes in. Flow efficiency is about what happens to it once it is in. One is a design problem: are we causing our own demand? The other is a queue problem: how many hands, how many handoffs, how hot we run. Programmes routinely attack one and report success on the other: the desk halves its repeat contacts and every remaining request waits exactly as long as before.

If you can only measure one, measure the one you are worse at. You will need both to find out which that is.

FAILURE DEMAND	N	FLOW UNDER 5%	FLOW EST.
Under 10%	22	36.4%	9.2%
10 to 25%	58	39.7%	9.6%
25 to 40%	69	46.4%	9.9%
40 to 60%	60	43.3%	9.6%
Over 60%	36	41.7%	10.2%

Among the 245 who could state both numbers.

THE CUTS • SECTOR AND SIZE

Everyone has about a third. The waiting varies.

Failure demand estimates sit between 30 and 38 per cent in every sector and every size band. Flow efficiency spreads from 6 to 12.

BY SECTOR

SECTOR	N	FD EST.	FLOW EST.	DK FD	DK FLOW
Technology	77	36.3%	11.0%	16.9%	36.4%
Professional services	68	36.9%	11.4%	26.5%	30.9%
Financial services	55	37.3%	11.6%	23.6%	36.4%
Public sector	54	34.5%	6.6%	14.8%	37.0%
Something else	33	37.8%	9.3%	33.3%	54.5%
Retail and hospitality	31	30.1%	9.9%	35.5%	54.8%
Healthcare	30	37.8%	8.5%	16.7%	30.0%
Manufacturing	29	37.2%	5.8%	24.1%	31.0%
Education	23	32.2%	9.3%	13.0%	30.4%

BY SIZE OF THE ORGANISATION SERVED

PEOPLE SERVED	N	FD EST.	FLOW EST.	DK FD	DK FLOW
Under 250	55	34.0%	12.2%	21.8%	40.0%
250 to 1,000	81	37.1%	9.9%	32.1%	34.6%
1,000 to 5,000	124	34.3%	8.4%	24.2%	39.5%
5,000 to 20,000	88	37.2%	8.5%	13.6%	36.4%
Over 20,000	52	37.5%	11.7%	17.3%	34.6%

Estimates are band-midpoint means among those who could answer. DK: the share answering "we do not know".

SECTOR

The sectors sit within 8 points of each other on failure demand, which is either a real regularity or the limit of what self-report can resolve. Flow is where they separate. Manufacturing (6 per cent) and public sector (7) report the longest waits; financial services, professional services and technology sit at 11 to 12.

Retail and hospitality reports the lowest failure demand (30 per cent) and one of the highest don't-know rates (36 per cent), which is one way of getting a low number.

SIZE

Failure demand is flat across size. Flow efficiency is worst in the middle: organisations serving 1,000 to 20,000 people report about 8 per cent, against about 12 at either end. Enough process to wait, not enough scale to fix it.

READ THEM AS INDICATIVE

The smallest sector cell is 23 responses and the largest 77. Every cell clears the publishing floor. None is large enough to settle an argument, and a difference of a point or two between two rows is noise.

METHOD

How the numbers were made.

THE SIX QUESTIONS, AS ASKED

- 01 Who do you serve?**

Customers, under contract · Colleagues, inside one organisation · Both
- 02 How big is the organisation you serve?**

Under 250 · 250 to 1,000 · 1,000 to 5,000 · 5,000 to 20,000 · Over 20,000
- 03 Which sector?**

Public sector · Financial services · Healthcare · Manufacturing · Retail and hospitality · Technology · Professional services · Education · Something else
- 04 How much of your demand exists because something already went wrong?**

Under 10% · 10 to 25% · 25 to 40% · 40 to 60% · Over 60% · We do not know
- 05 Of the time you hold a piece of work, how much of it is actually being worked on?**

Under 1% · 1 to 5% · 5 to 10% · 10 to 25% · Over 25% · We do not know
- 06 What utilisation do you plan your teams to?**

Under 70% · 70 to 80% · 80 to 90% · Over 90% · We do not plan to a figure

THE SAMPLE

400 responses through markboyer.co.uk/benchmark, 1 to 29 August 2026, collected with a plain web form and stored by the site's host. No weighting. No verification. An email address was optional, was left blank by most people, and played no part in any analysis. Nothing that identifies a respondent or an employer was asked for, and nothing of the kind is held in the published data.

THE ESTIMATES

Answers are bands, so every single figure is the mean of band midpoints among respondents who could answer. "We do not know" is reported as its own category and is never imputed.

FAILURE DEMAND	MIDPOINT	FLOW EFFICIENCY	MIDPOINT
Under 10%	5	Under 1%	0.5
10 to 25%	17.5	1 to 5%	3
25 to 40%	32.5	5 to 10%	7.5
40 to 60%	50	10 to 25%	17.5
Over 60%	70	Over 25%	30

The open top bands are the judgement calls. Putting "over 60%" at 65 or 75 instead of 70 moves the failure demand estimate to 35.2 or 36.6. Putting "over 25%" at 35 moves flow efficiency to 10.2.

THE FLOOR

No cut with fewer than 20 responses is published, anywhere. Where a cell falls below it the table says so rather than showing a number.

LICENCE

Published free at markboyer.co.uk/benchmark/findings/ under CC BY 4.0. Cite it, quote it, reuse the charts; the only condition is attribution and a link.

DO IT PROPERLY

Work out your own numbers. It takes an afternoon.

Neither of these needs a tool, a project or a licence. They need one week of contacts, twenty closed requests and the willingness to divide one number by another.

FAILURE DEMAND

Take one week of contacts. All of them. For each one ask a single question: would this contact exist if we had done the earlier thing right? Chasing an update is a yes. A second report of the same fault is a yes. A password reset is a real request, badly designed, and you can argue about it later.

Count the yeses and divide by the total. That is your failure demand. It will be higher than the number you would have guessed, which is the point of not guessing.

FLOW EFFICIENCY

Take twenty recently completed requests, picked at random rather than the ones you are proud of. For each, add up the minutes anyone actually spent working on it, then divide by the hours between it being raised and being closed. Average the twenty.

If the answer is above 10 per cent, check the arithmetic. If it is above 25, check who chose the twenty.

THEN WHAT

Failure demand is a design problem. Go to the earliest point in the journey that generated the contact and fix that, not the contact. Flow efficiency is a queue problem: fewer handoffs, smaller batches, and a utilisation target the queue can survive. Neither shows up on a dashboard built to count tickets, which is why neither was ever on one.

If you would rather work them out properly than estimate them, that is what Shift Right is for: markboyer.co.uk/shift-right/



Add your numbers to the
benchmark

Six questions · ninety seconds
markboyer.co.uk/benchmark

The benchmark stays open. Add your numbers.

Every response makes the number less of a guess. Six questions, ninety seconds, anonymous, and published only in aggregate.



markboyer.co.uk/benchmark

36%

of demand exists because something already went wrong

10%

flow efficiency: the work is being worked on one day in 10

39%

cannot state at least one of the two numbers

CITE IT

Boyer, M. (2026) *The State of Service Flow: failure demand and flow efficiency in IT services*.
markboyer.co.uk/benchmark/findings/

Licensed CC BY 4.0. Reuse the findings and the charts with attribution and a link. 400 responses · published September 2026.

ABOUT THE AUTHOR

Mark Boyer is Global Services Director at Softcat plc and the creator of the Human Operating Model. He speaks and writes about failure demand, flow, and the redesign of IT services around the people they serve.

markboyer.co.uk ·
markboyer.co.uk/speaking/
linkedin.com/in/mark-boyer

