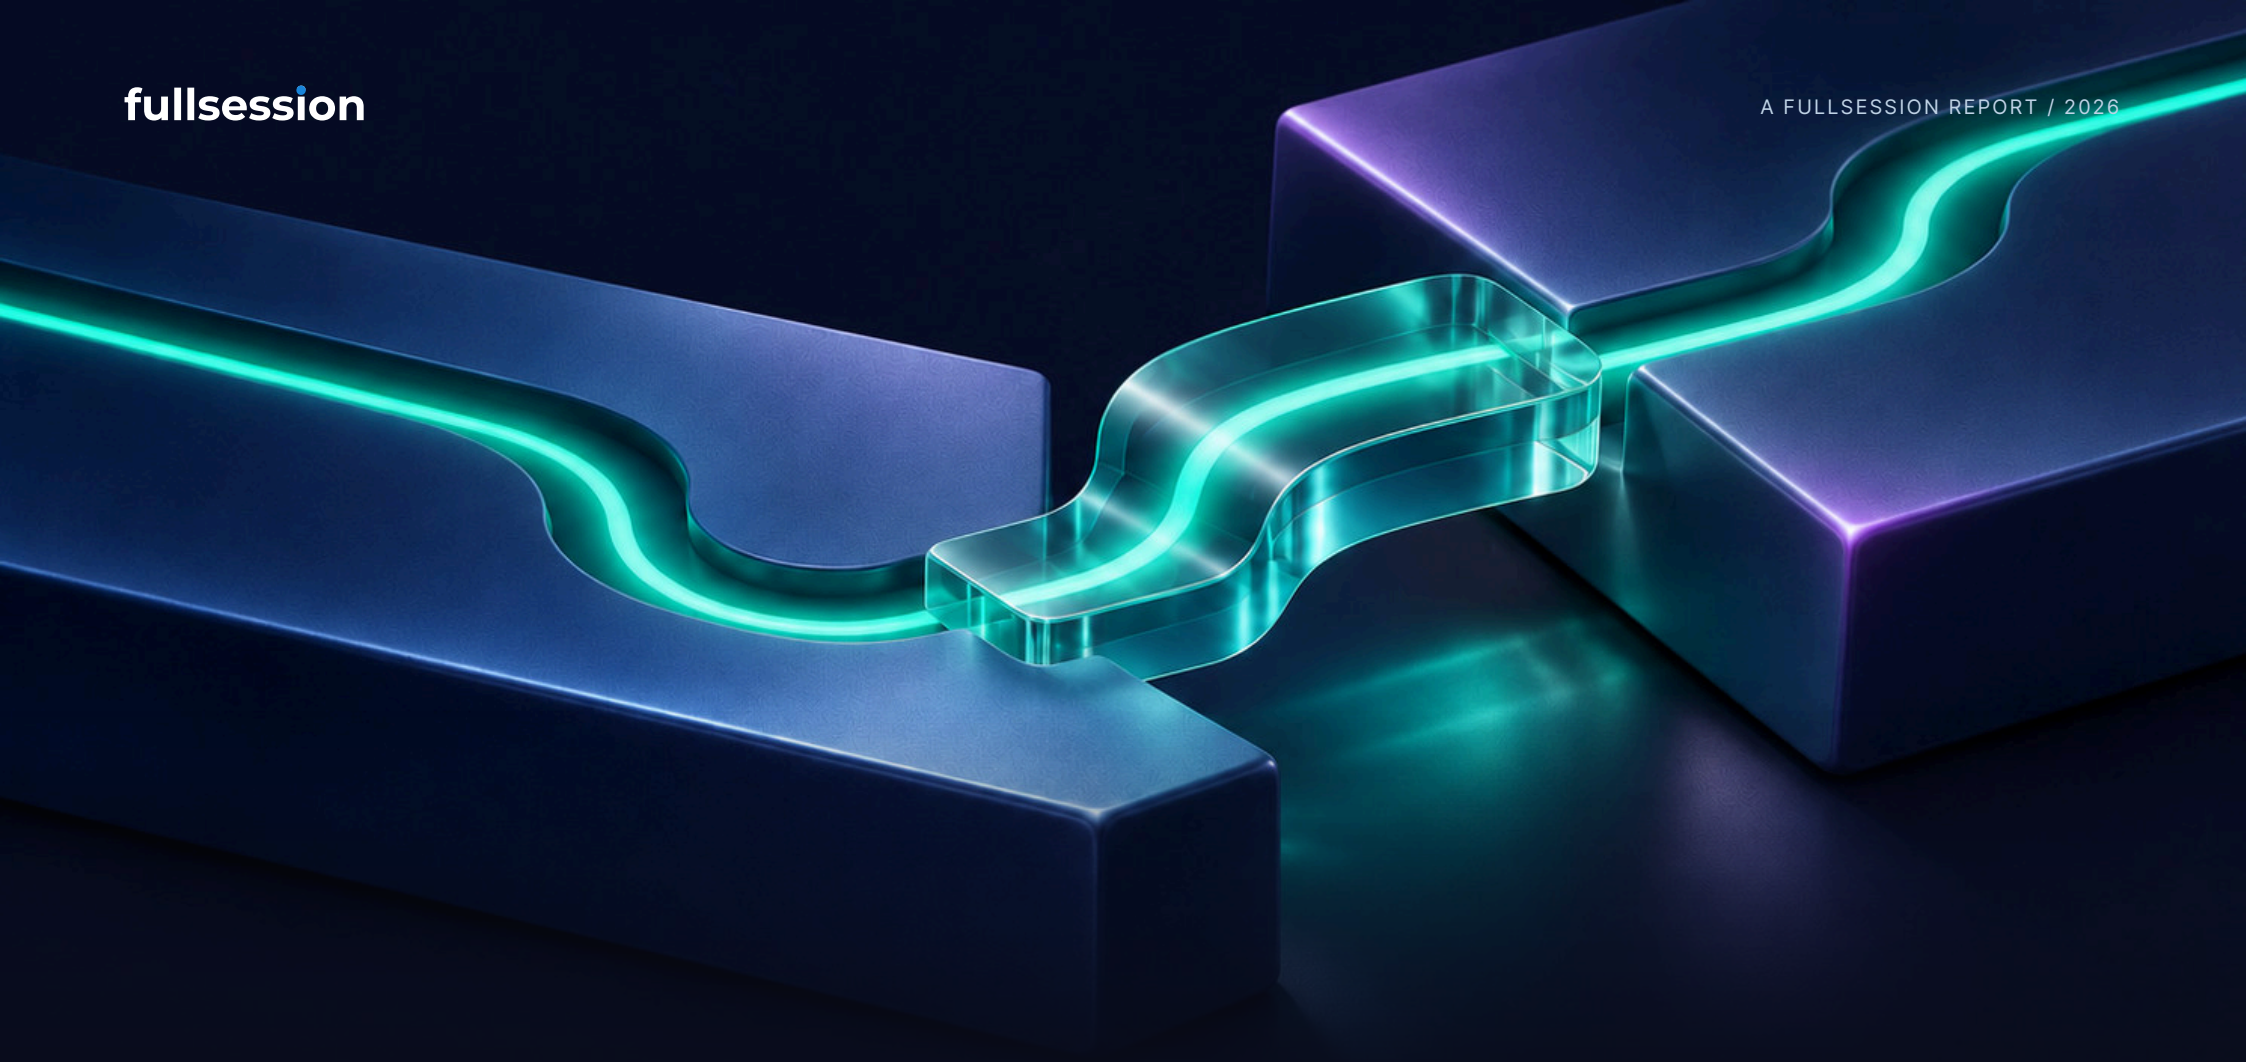
fullsession

A FULLSESSION REPORT / 2026



The Context Gap

Why AI inside the application is the next digital adoption decision.

CONTENTS

What is in this report

Introduction 3

How to read this report 4

Five findings 5

Act 1 Use is up, finished work is harder to see 6

Useful to people, harder to see in results 7

The work arrives in pieces 8

Three places the work stalls 9

Use thins out near the work 10

Leadership support goes with use 11

Help arrives before the work, or after it 12

One case, two records 13

Act 2 Where the thread breaks 14

Two journeys, one pattern 15

Signals worth measuring 16

Checking needs to be part of the task 17

People bring their own AI 18

Connected applications are the exception 19

Count every case, time only the work 20

Measure the first useful result 21

What an assistant needs to finish work 22

Act 3 Two ways to close it 23

Adoption is completed, correct work 24

One firm, help inside the task 25

Two routes, decided by ownership 26

Max, inside your application 27

Daleel, the shared layer 28

Detect and Coach, in the apps your team uses 29

More connected context, the same authority 30

From help to decision support 31

Measure, help, measure again 32

What counts as evidence 33

Privacy and scope 34

Worksheet 35

The thread, carried 36

About FullSession 37

Sources 38

Definitions and credits 39



Roman Mohren
CEO, FullSession

INTRODUCTION

AI becomes useful when it understands the task

Digital adoption used to mean getting people to use software their organization had already paid for. The measure was whether they logged in and whether they finished the training.

That measure no longer answers the question leaders are asking. Assistants are now in browsers, inboxes and applications. The public research in this report shows people finding them useful for their own work, while far fewer organizations can point to a change in results. The work still crosses several systems, and the help still sits outside the step where things go wrong.

Our position is simple. An assistant becomes useful to an organization when it understands the task in front of the person, works within what that person is permitted to do in that application, and can establish whether the result came out right.

That is the context gap. Closing it is a decision about where AI sits: inside the application and the task, with the application's meaning, the person's goal and the relevant connected systems available to it, and with authority that stays with the organization.

This report combines cited public findings with FullSession's view. The research figures name their source, sample and method so you can check them. The worked examples are illustrative and say so. The method in Act 3 is one you can run yourself, whichever tools you use.

Who it is for depends on one question: do you own the application? Page 4 explains why that decides the route.

Roman Mohren

HOW TO READ THIS REPORT

Two readers, one question, three kinds of content

Which reader are you?

The problem in this report is the same for two groups of people. What they can do about it depends on whether they own the application where the work happens.

You own the application

Your product, portal or internal app. The people using it may be your customers or your employees. You can place help inside it.

FullSession Digital Adoption

You do not own it

Third-party web applications your teams work in every day: systems of record, portals, queues. You cannot change them.

FullSession Workforce

> Acts 1 and 2 apply to both. Act 3 shows the two routes and the measurement both share.

Three kinds of content

PUBLIC FINDING — A figure from a named, dated public source. The caption gives publisher, year, sample and method. Full references on pages 38 and 39.

FULLSESSION VIEW — Our interpretation and method. A position, not a measured result.

ILLUSTRATIVE — A worked example with invented values, used to show how something works. Not a benchmark, customer result or expected outcome.

What we did not do. Commission a survey, use customer data, pool figures across studies, or estimate the effect of any product.

Two terms we keep apart. An observed change is a measured difference from a baseline. An estimated effect compares that change with an uncoached group, under stated assumptions (page 32).

FIVE FINDINGS

AI reached the desk. The work still breaks.

01

Useful to people. Harder to see in results.

In McKinsey's 2026 survey, 80% say AI improved their individual productivity. A separate question finds 37% whose organizations attribute any EBIT impact to AI.

McKinsey, State of AI 2026, survey, n=1,719

02

The work arrives in pieces.

48% of employees in Microsoft's 2025 Work Trend Index say their work feels chaotic and fragmented. A case crosses several applications. In our view, an assistant that lives in one of them sees one step.

Microsoft WorkLab, June 2025, survey, n=31,000

03

Leadership support goes with use.

In BCG's 2025 survey, frontline employees with clear leadership support are twice as likely to use GenAI regularly. Only 36% of all respondents say their training was enough.

BCG, AI at Work 2025, survey, n=10,635; frontline n=3,537

04

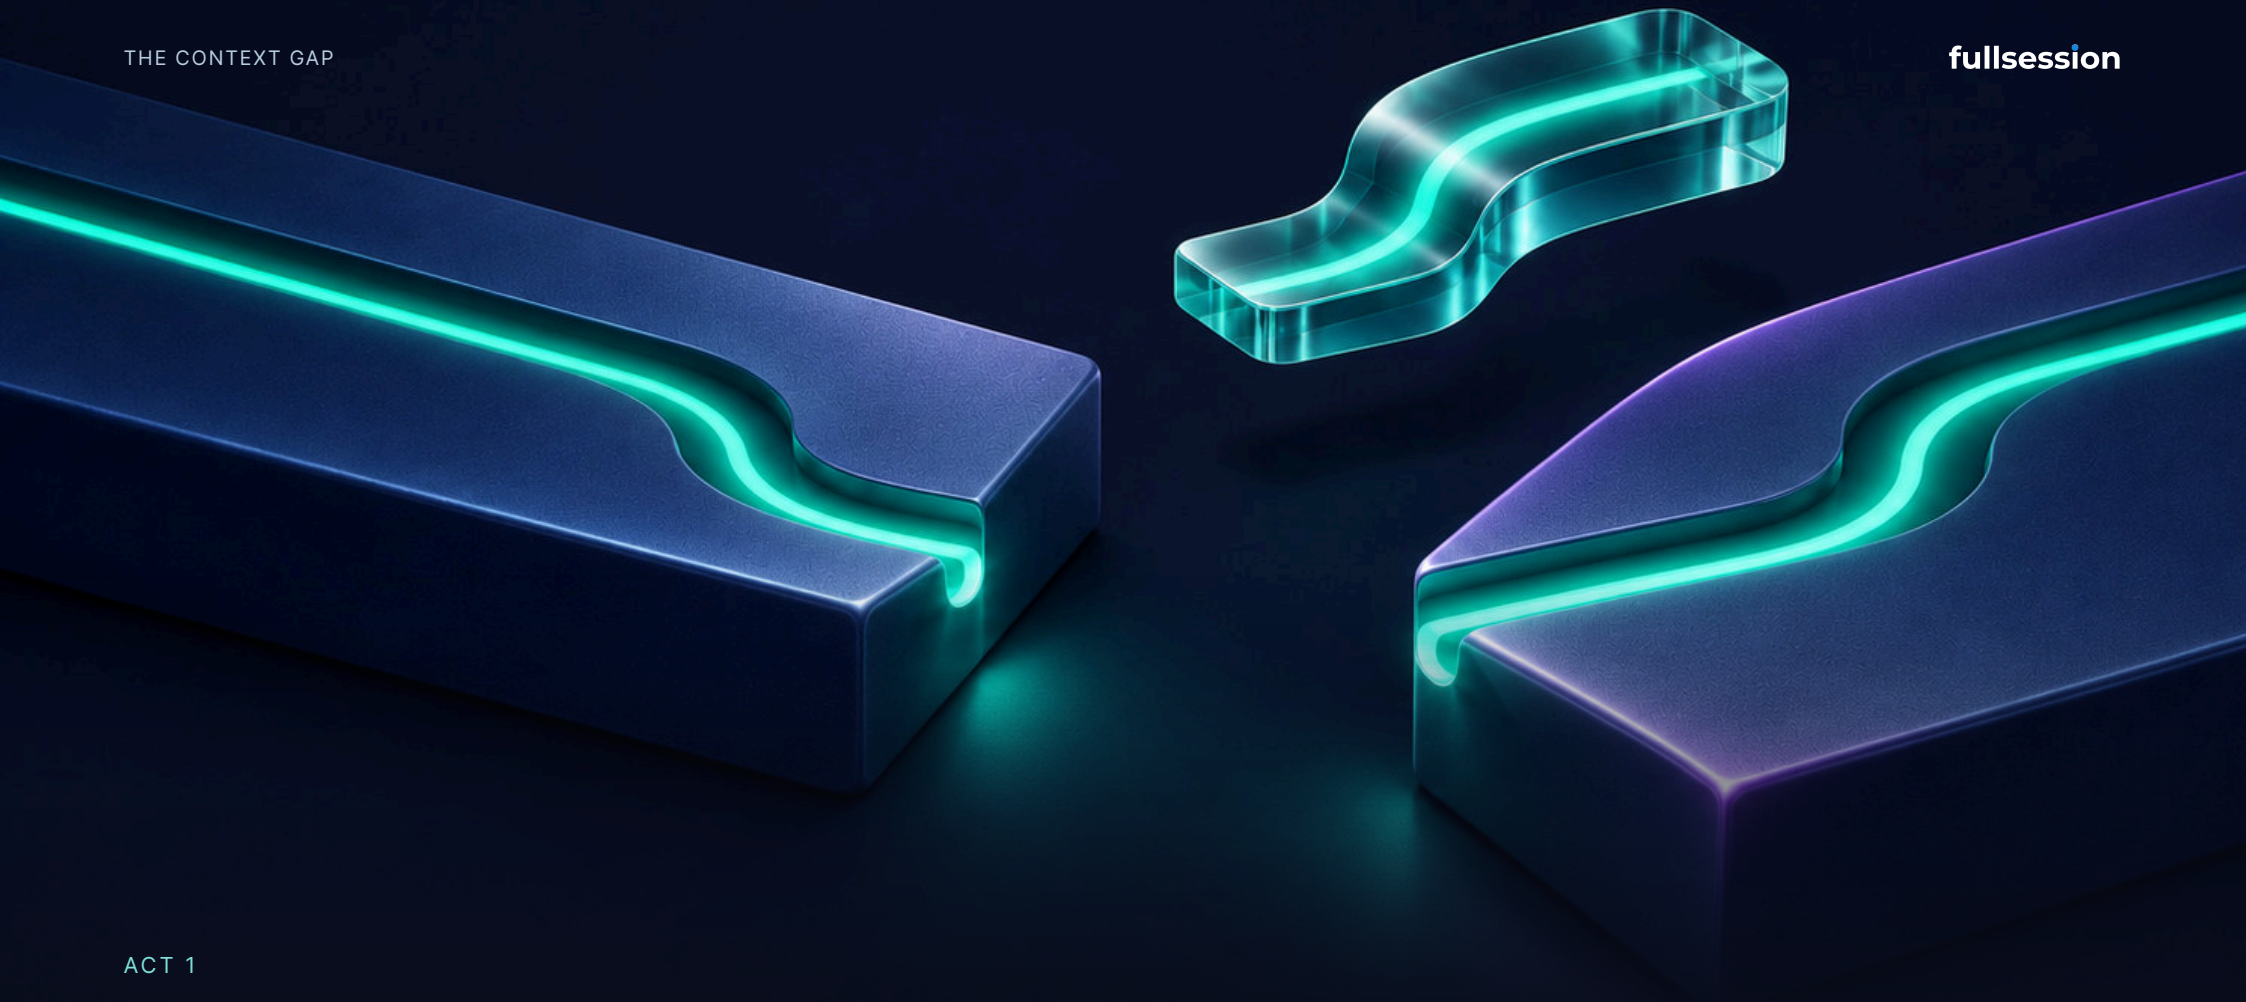
Context limits what an assistant can finish.

Help that understands one screen is useful. Finishing a task also needs the application's meaning, the person's goal and permissions, and the systems that hold the rest of the case.

05

Two routes close it, and both can be measured.

Where you own the application, help goes inside it. Where you do not, it goes into the browser your team already uses. Either way, measure completed, correct work against a comparison group.



ACT 1

Use is up.
Finished work is harder to see.

Assistants are deployed and people use them. The research shows where that use stops short of finished, correct work, and who is left without help.

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

Useful to people. Harder to see in results.

Most respondents to McKinsey's 2026 survey say AI makes them personally more productive. A separate question finds far fewer organizations attributing any effect on earnings to it.

The two figures answer different questions about different levels, so the distance between them is not a shortfall to subtract. What they show together is that individual usefulness does not add up to organizational results on its own. Something has to connect the faster step to the finished case, the correct submission, the resolved customer.

McKinsey also reports that nearly three quarters of the respondents it classes as AI high performers have fundamentally redesigned workflows because of AI. The work changed, not only the tools.

> High performers is McKinsey's defined subgroup: respondents attributing at least 5% of EBIT to AI and reporting significant value.

PUBLIC FINDING · TWO SEPARATE QUESTIONS

Asked of individuals

80%

say AI has improved their individual productivity

Asked about the organization

37%

say their organization attributes at least some EBIT impact to AI

McKinsey & Company, The state of AI in 2026: On the road to ROI, 25 August 2026. Online survey of 1,719 participants in 97 nations, 4 May to 8 June 2026. Self-reported. Separate questions; not a measured gap. mckinsey.com

PUBLIC FINDING

48%

of employees say their work feels chaotic and fragmented. Among leaders it is 52%.

Microsoft WorkLab, Breaking down the infinite workday, 17 June 2025. Survey of 31,000 full-time knowledge workers in 31 markets, 6 February to 24 March 2025, by Edelman Data x Intelligence. Self-reported. microsoft.com

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

The work arrives in pieces

A case rarely lives in one application, and a new user rarely gets value on one screen.

An operations case starts in a queue, moves to a system of record, detours through a portal or a worksheet, and comes back for review. A new customer in a product signs up, connects something, invites someone, and only then sees the result they came for. Every step is a place where context gets lost: what was agreed, what was already checked, what the next screen expects.

An assistant that sees one step can help with that step. Fragmentation tells you where to look. It does not tell you which step costs the most; Act 2 shows how to find out.

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

Three places the work stalls

The same three gaps appear whether the person is a customer in your product or an employee in someone else's.

FULLSESSION VIEW

1

The help does not know the application

It can describe a feature in general. It cannot say that this field is required for this account type, or that the approval being waited on was already granted.

2

The help is not where the task is

Training came before the work and documentation lives in another tab. When a person hesitates they ask a colleague, guess, or export to a spreadsheet.

3

The systems do not share the case

The customer record, the ticket, the order and the usage data each hold part of the story. Help without them can still answer a question or support a bounded decision, but it may miss the rest of the case.

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

Use thins out near the work

The people closest to the work use GenAI least. In BCG's 2025 survey, regular use falls from 85% of leaders to 51% of frontline employees.

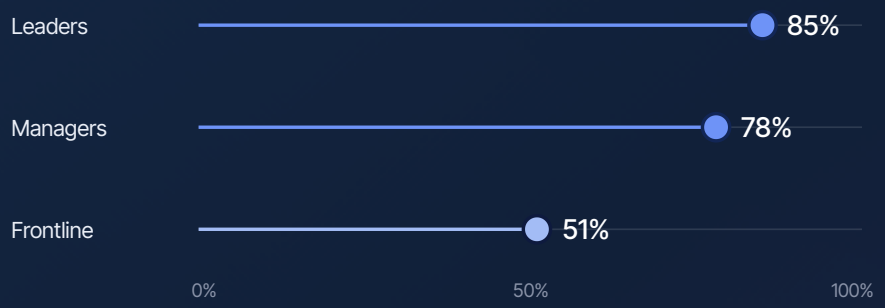
The frontline is where cases are completed and customers are served, and it is where help is thinnest. BCG names three things in short supply there, proper training, the right tools and leadership support, and treats them as complementary.

BCG surveyed employees. The parallel we draw with customers is ours, not a measured finding: for a product team, the people closest to the work are often the customers using your application; for an operations leader, the team working in software the organization bought.

PUBLIC FINDING

Regular GenAI use by level, 2025

Share of respondents at each level who use GenAI several times a week or daily.



BCG, AI at Work 2025: Momentum Builds, but Gaps Remain, June 2025, deck page 6. Survey of 10,635 leaders, managers and frontline employees in 11 countries and regions. Self-reported. bcg.com

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

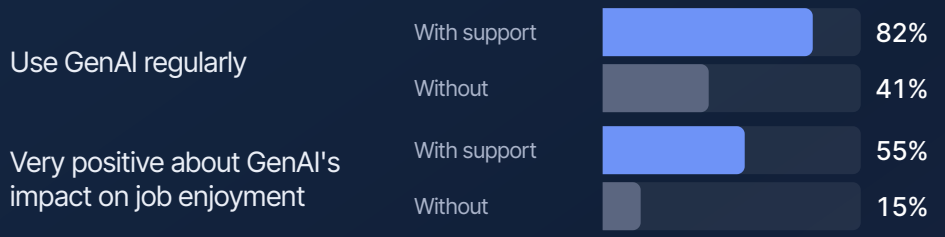
Leadership support goes with use

Frontline employees who report clear leadership support on GenAI use are twice as likely to use it regularly, and far more likely to say it makes their job more enjoyable.

This is an association in a survey, not a controlled result, and it concerns leadership support, not in-application guidance. We use it for one point only: what surrounds a tool goes with how much it is used. Whether help inside the task changes outcomes is a question each organization can test, and page 32 shows how.

PUBLIC FINDING · ASSOCIATION

Frontline employees, with and without clear leadership support on GenAI use



BCG, AI at Work 2025: Momentum Builds, but Gaps Remain, June 2025, deck page 15. Frontline subgroup, n=3,537. Self-reported survey; association, not a causal estimate. bcg.com

ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

Help arrives before the work, or after it

PUBLIC FINDING

36%

of respondents say the training they received on GenAI was enough.

BCG, AI at Work 2025: Momentum Builds, but Gaps Remain, June 2025, deck page 13. n=10,635. Self-reported. bcg.com

FULLSESSION VIEW

Training is necessary, and it often happens before the work. Support tickets and coaching often happen after it. The moment that decides whether a case comes out right, the hesitation on a step, is often covered by neither. That is true for a new customer in your product and for an analyst in a system your company licenses. It is the moment Act 3 is about.



ACT 1 · USE IS UP, FINISHED WORK IS HARDER TO SEE

One case, two records

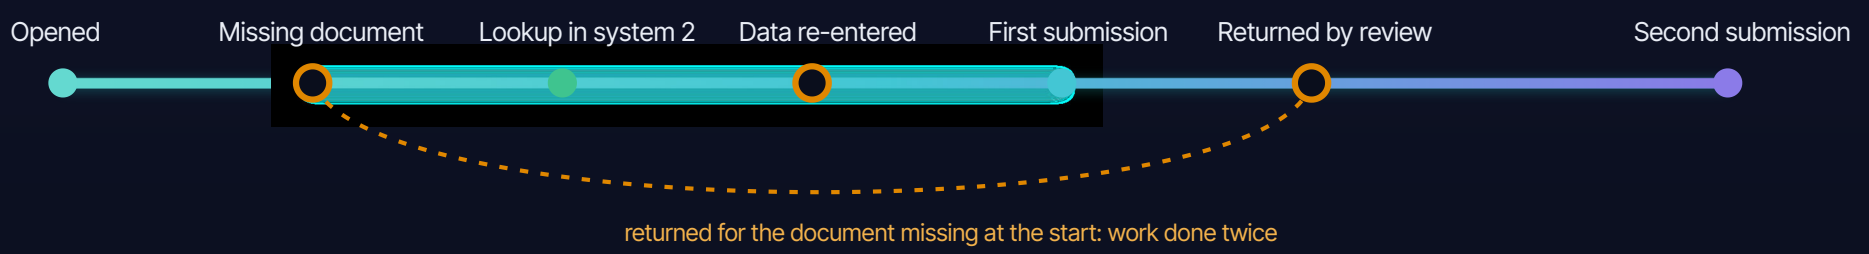
An adoption dashboard records what is easy to count. The case itself went through more than that. The cost sits in the second record.

ILLUSTRATIVE

WHAT THE DASHBOARD RECORDED



WHAT THE CASE WENT THROUGH



Neither record is wrong. The first answers "is the tool being used?" The second answers "did the work come out right the first time, and if not, where did it come back?" The amber points can be observed in aggregate in the browser without changing the application. Act 2 shows how to measure them.



ACT 2

Where the thread breaks

Follow one piece of work from start to finish and the breaks are visible: where context was lost, where someone went around the system, and how to count what that costs.

ACT 2 · WHERE THE THREAD BREAKS

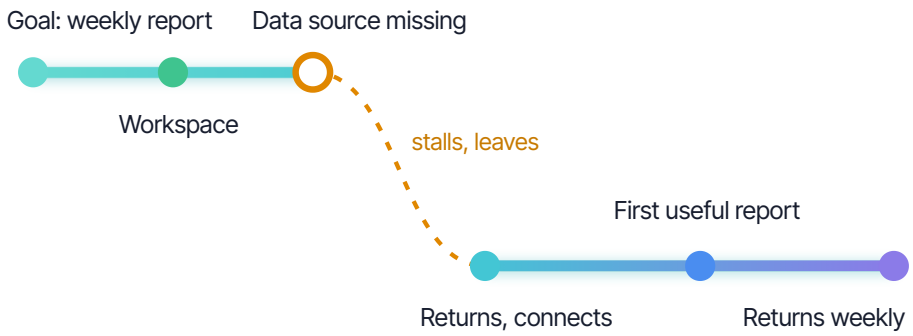
Two journeys, one pattern

A new user in an application you own, and a loan processor working across applications you do not. Different people, different software, the same kind of break.

In an application you own

ILLUSTRATIVE

An analyst at a customer signs up for a reporting product, wanting a weekly account report.

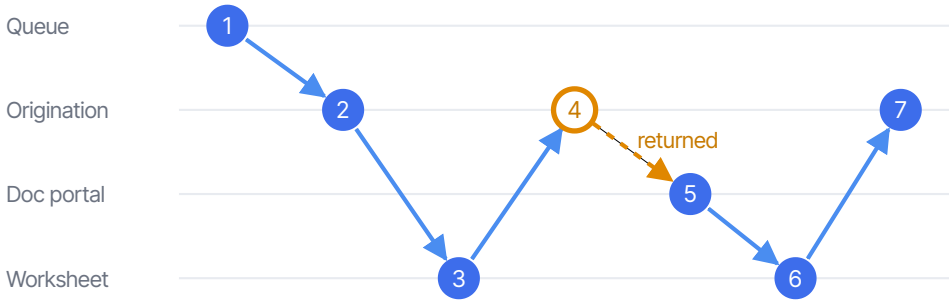


The prerequisite was known to the product, not to the user. Activation, the first useful verified result, waited days on one missing step.

Across applications you do not own

ILLUSTRATIVE

A processor prepares a loan file across a queue, an origination system, a document portal and an income worksheet.



1 Assigned
2 File opened
3 Income entered, one document
4 Submitted; returned by review
5 Second document attached
6 Worksheet corrected
7 Resubmitted

The worksheet accepted the file without the second borrower's document. Review caught it at step 4, and steps 5 to 7 are the work done again.

The common pattern. In both cases the information needed to get it right existed somewhere: in the product's own setup rules, in the lender's procedure. It was not available at the step where the person needed it.

ACT 2 · WHERE THE THREAD BREAKS

Signals worth measuring

Each one can be observed without asking anyone to fill in a form. None of them proves a problem on its own; treat each as evidence to classify, not a verdict.

FULLSESSION VIEW

In an application you own

Stalled before a prerequisite	New users who stop at a step that depends on something not yet done
Time to first useful result	From sign-up to the first result the user came for, verified, not just a page view
Repeated help on the same step	Questions or retries clustering on one screen or one setting
Return and sustained use	Whether users come back at the cadence the job needs, and do useful work when they do
Contact after self-service	A support contact soon after the user was shown an answer

Across applications you do not own

Returned to a prior step	A case goes back to an earlier stage or to the same person
Failed or retried submission	A form is rejected and submitted again
Repeat lookup	The same record is searched for again within one case
Re-entered data	The same value typed into a second system
Export out of the app	A download or copy mid-task, often to a spreadsheet

Classify before you count. A reopened case can be a legitimate update. A fast path can be wrong. A drop in support contacts can mean users gave up. Agree what outcome counts as correct with the process or product owner first, and count each signal against it.

PUBLIC FINDING

2 in 3

employees who use AI at work report having relied on its output without critically evaluating it.

31%

of employees intentionally use AI at work weekly or daily

2 in 5

report that their organization has a policy or guidance on generative AI use

KPMG and the University of Melbourne, Trust, attitudes and use of AI: A global study 2025. Survey of 48,340 people in 47 countries, November 2024 to January 2025; overall employed sample n=32,352; base for two in three: employees who use AI at work (item n not reported). Executive summary pp 5 to 6; fig 37 pp 67 to 68; fig 41 p 71; fig 44 pp 75 to 76. Self-reported; "relied without evaluating" counts any occurrence, not a habit. kpmg.com

ACT 2 · WHERE THE THREAD BREAKS

Checking needs to be part of the task

Two in three employees who use AI at work say they have relied on its output without checking it. Two in five employees say their organization has a policy or guidance on generative AI.

These are separate questions with different bases, so they are shown as separate figures rather than as one chart. The findings show gaps in checking and in employees' awareness of organizational guidance.

FULLSESSION VIEW

Checking is easier when the assistant knows what a correct result looks like in this application. That is why the outcome check on page 22 is part of the design, not an afterthought: help that can say whether the result came out right gives the person something better than trust.

ACT 2 · WHERE THE THREAD BREAKS

People bring their own AI

Many people who use AI at work bring their own tools to it, and many prefer not to say how they use them.

Bringing your own tool is not the same as breaking a rule; Microsoft does not describe it that way, and neither do we. Those tools may sit outside the organization's connected context and controls. KPMG's study adds that almost one in two employees who use AI say they have used it in ways that contravene their organization's policies or guidelines.

FULLSESSION VIEW

The practical answer is not only a policy. It is an approved assistant that is better at the task than the one people bring, because it knows the application and works within its permissions.

PUBLIC FINDING

78%

of people who used AI at work in 2024 brought their own AI tools to work.

52% of people who used AI at work in 2024 were reluctant to admit using it for their most important tasks

Microsoft and LinkedIn, 2024 Work Trend Index Annual Report, AI at work is here. Now comes the hard part, 8 May 2024. Survey of 31,000 knowledge workers in 31 markets, 15 February to 28 March 2024, by Edelman Data & Intelligence. Base for both figures: AI users. Self-reported. microsoft.com. KPMG statement: Trust, attitudes and use of AI 2025, p 6.

ACT 2 · WHERE THE THREAD BREAKS

Connected applications are the exception

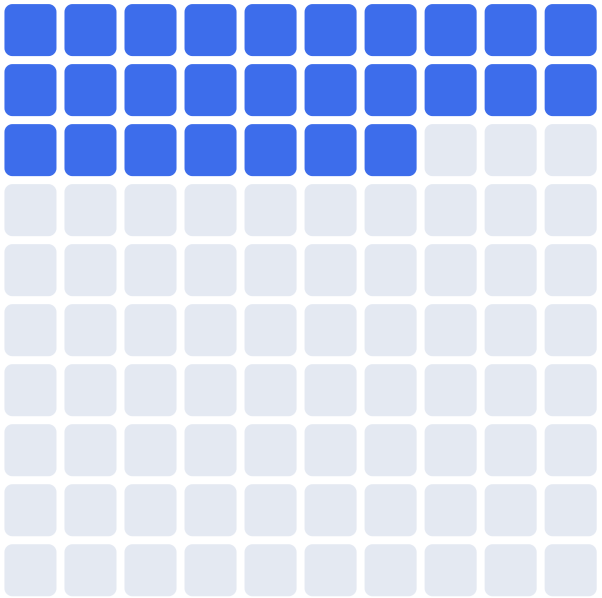
In MuleSoft's 2026 survey of enterprise IT leaders, the average organization runs 957 applications, and on average only 27% of them are connected.

82% of the same IT leaders cite data integration as one of the biggest challenges their organization faces when using AI. This is a vendor survey of perceptions in large enterprises, and it says nothing about any particular product. The finding shows why access to the relevant systems has to be established for each task an assistant supports.

FULLSESSION VIEW

That is why integration belongs in the adoption conversation. Every system connected to the layer where help is delivered widens what that help can understand. Page 30 shows what changes, and what does not.

PUBLIC FINDING



27%

of applications connected, on average

957

applications in the average organization

MuleSoft, 2026 Connectivity Benchmark Report, 5 February 2026, sections The expanding application landscape and The connectivity paradox. Vendor survey of 1,050 IT leaders at organizations with 1,000 or more employees in 9 countries, October to November 2025, by Vanson Bourne. Self-reported averages. mulesoft.com



ACT 2 · WHERE THE THREAD BREAKS · APPS YOU DO NOT OWN

Count every case, time only the work

Two counting rules prevent the most common misleading comparisons in workflow measurement.

FULLSESSION VIEW

Keep unfinished cases

Freeze the cohort of cases that entered in a period and follow all of them for the same horizon. A completed-case average drops the hardest cases and can make a worse month look cheaper.

Time active work, not elapsed time

A file open for three days may hold forty minutes of work. Count non-overlapping active minutes; keep waiting separate.

ILLUSTRATIVE

20.0	21.8	88%
active min per completed file (880 of 1,000)	active min per eligible file, 30-day horizon	completed within the horizon

> 17,600 active minutes across 880 completed files = 20.0. Adding 4,200 minutes on 120 unfinished files: 21,800 across all 1,000 = 21.8. Different questions, both useful, shown with completion. Time saved is capacity, not cash, until it absorbs volume, covers attrition or avoids an expense.

ACT 2 · WHERE THE THREAD BREAKS · APPS YOU OWN

Measure the first useful result

For an application you own, adoption has two moments worth defining: the first time a user gets the result they came for, and whether they keep getting it.

FULLSESSION VIEW

Define thresholds for your product and goal rather than importing a generic activation or time-to-value benchmark. What travels is the method: a fixed cohort, an outcome you can verify, a horizon, and every user kept in the denominator, including the ones whose outcome is unknown.

Support is where measurement most often flatters. Fewer contacts can mean users got their answer, or that they gave up. Count a resolution only when the user's outcome is confirmed, and look at escalations and repeat contacts beside it.

01 Eligible cohort

New users who started in a fixed window, by role and goal where known. Nobody leaves the denominator.

02 First verified useful result

Agreed per goal with the product owner: a report delivered, an invoice sent, an integration passing its check. Time to it, and share reaching it within the horizon.

03 Completion and quality

Whether the result was correct, not only reached. Errors, reversals and corrections count against it.

04 Sustained useful use

Return at the cadence the job needs, doing useful work, not just logging in.

05 Support outcome, not support volume

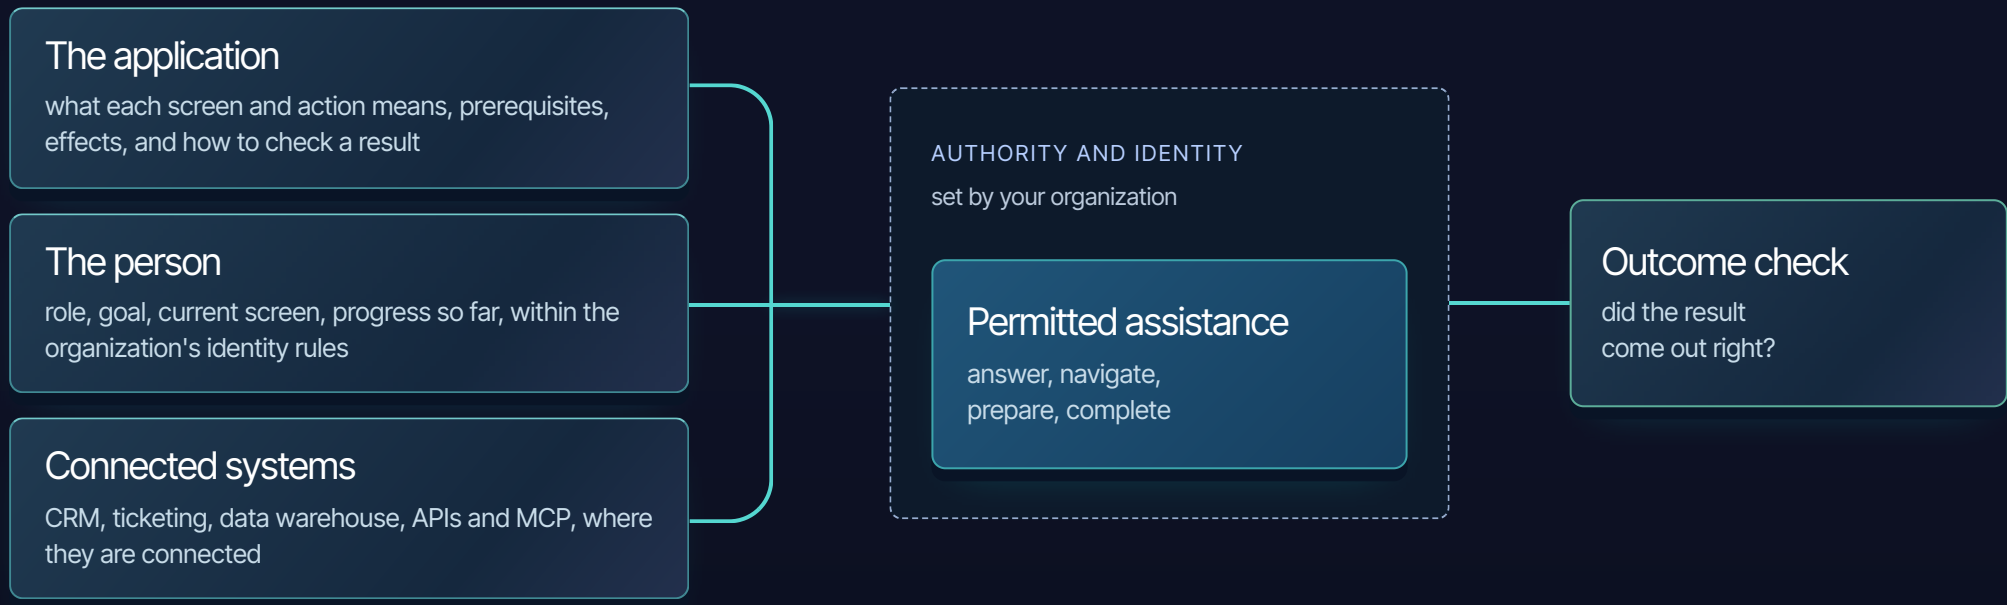
Confirmed resolutions, escalations and repeat contacts together; unknown outcomes reported, never counted as solved.

ACT 2 · WHERE THE THREAD BREAKS

What an assistant needs to finish work

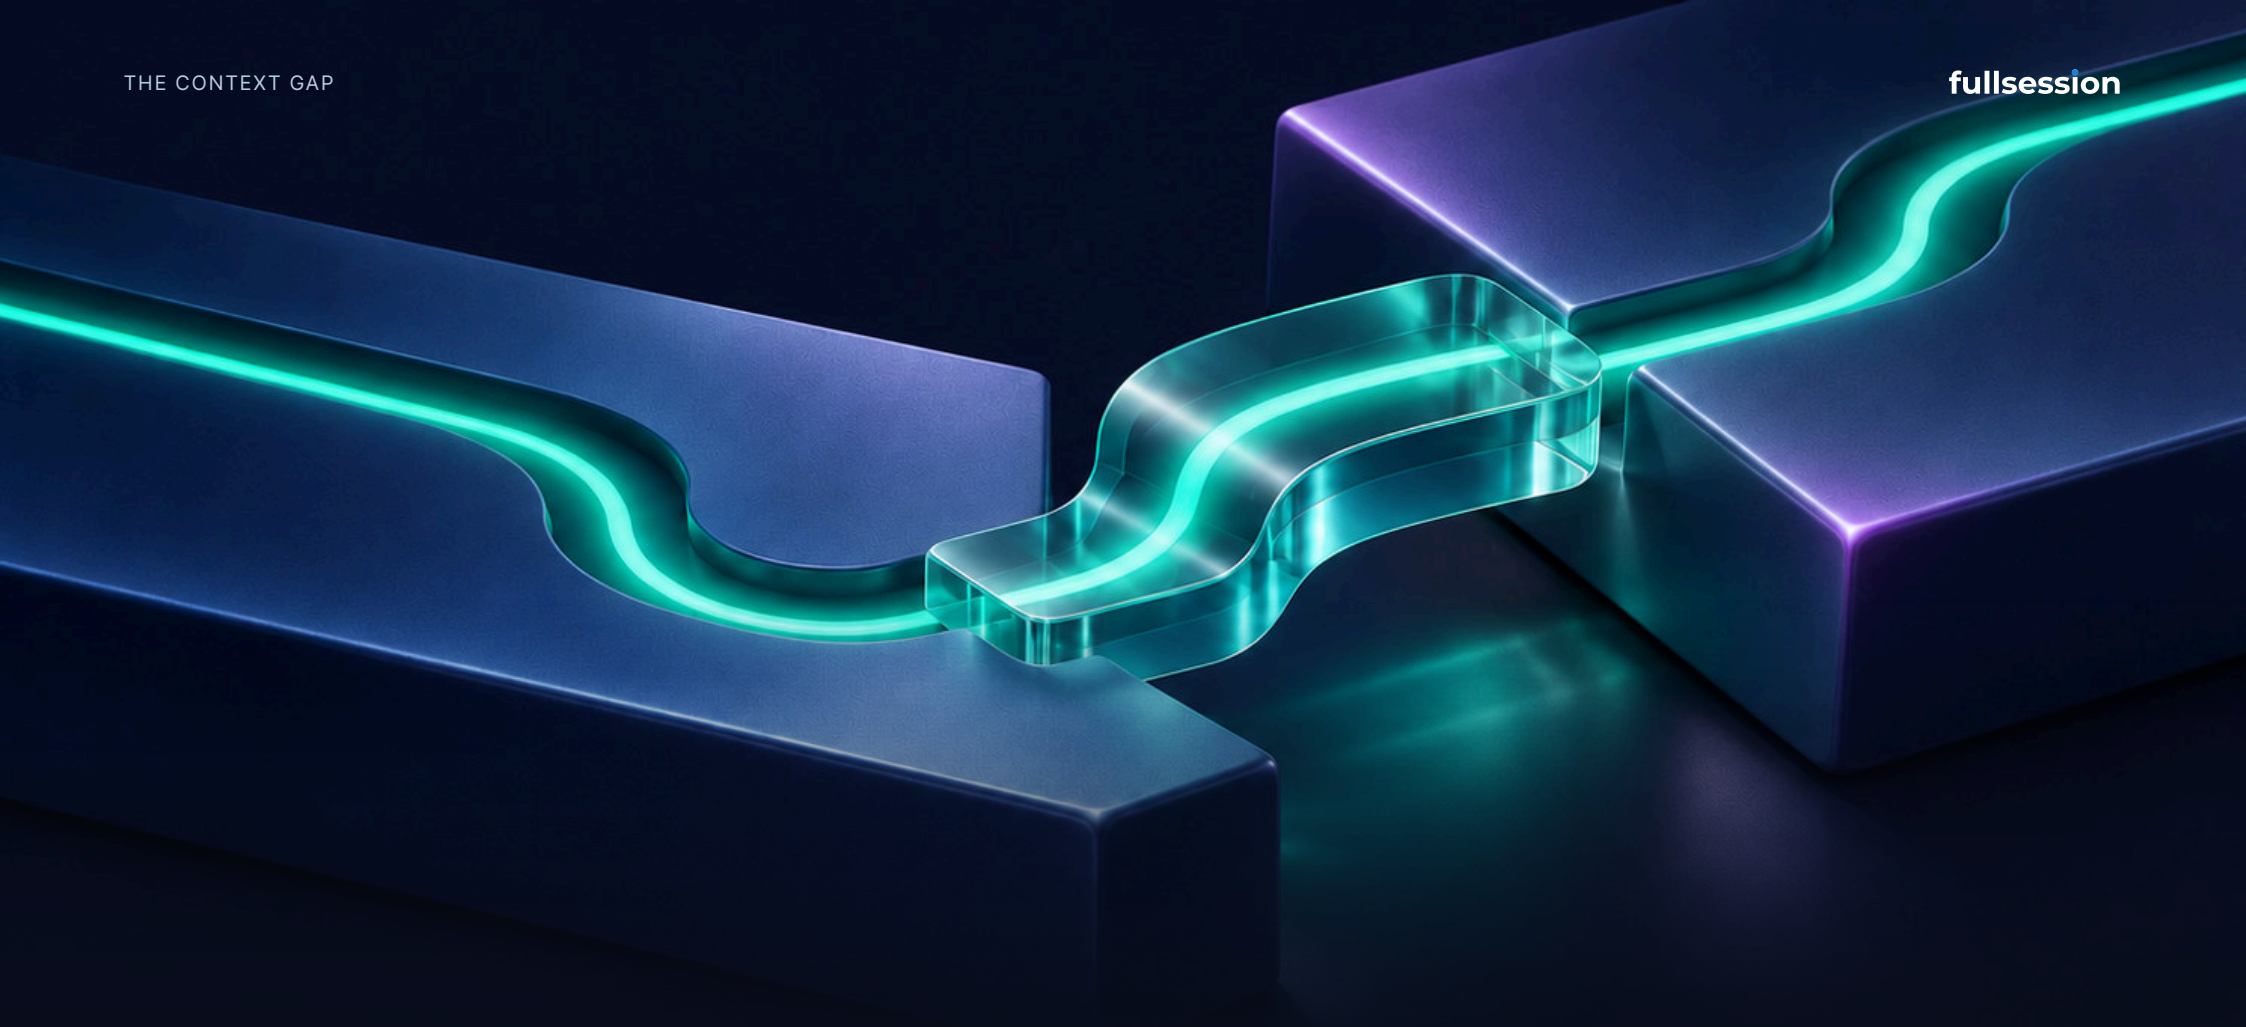
Grounded help on a single screen is already useful. Finishing a task takes more inputs, and a separate decision about what the assistant is allowed to do.

FULLSESSION VIEW



Context and authority are different things. More context lets the assistant understand more of the case; it never gives it permission. What it may do, for whom and in which application stays a setting the organization controls, checked against the person's own identity and permissions every time.

Readiness is also not one switch: a source can be connected without being understood, understood without being tested, and tested without being authorized or live.



ACT 3

Two ways to close it

Put help where the work happens, give it the context it needs, keep authority with the organization, and measure what changes. Whether you own the application decides the route.

ACT 3 · TWO WAYS TO CLOSE IT

Adoption is completed, correct work

Logins and course completions say a tool is in use. They do not say whether it helped. We define adoption by the work.

FULLSESSION VIEW

Activation	The first time a person gets the useful result they came for, and it is verified: the report delivered, the integration passing its check, the file accepted.
Sustained use	Returning at the cadence the job needs and doing useful work each time. There is no universal threshold; the product or process owner sets it.
First-time-right work	For operations in applications you do not own: the case completed correctly without coming back, with active time counted honestly.
Readiness is not activation	An assistant that is connected, configured and live has not yet helped anyone. Count the result, not the deployment.



PUBLIC FINDING · ONE FIRM, OBSERVED WORK

15%

more customer issues resolved per hour, on average, after support agents got an AI assistant inside their workflow.

~30% for less experienced and less skilled agents

Brynjolfsson, Li and Raymond, Generative AI at Work, Quarterly Journal of Economics 140(2), 889 to 942, 2025, doi:10.1093/qje/qjae044. Staggered rollout across 5,172 customer-support agents at one firm; observed productivity. Effects varied; the most experienced agents saw small speed gains and small quality declines. academic.oup.com

ACT 3 · TWO WAYS TO CLOSE IT

In one firm, help inside the task lifted output by 15%

One substantial field study of AI assistance in observed work, not a survey: support agents with an assistant inside the conversation resolved more issues per hour.

It is one setting: human agents with assistance, in one company's contact center. It is not evidence about software onboarding, autonomous support, or any vendor's product, including ours, and we do not use it as a forecast.

What it does show is the shape of a good test. The assistant was introduced to groups at different times, so the researchers could compare like with like. Results differed by experience: the largest gains went to the least experienced agents, with small quality declines for the most experienced. The authors find suggestive evidence that the assistant spread more skilled agents' practices. So examine results by experience level, not only on average. Page 32 applies the same design to your own rollout.

ACT 3 · TWO WAYS TO CLOSE IT

Two routes, decided by ownership

If you own the application, FullSession Digital Adoption works inside it, whether the people using it are your customers or your employees. If you do not own it, FullSession Workforce works in the browser your team already uses.

YOU OWN THE APPLICATION

FullSession Digital Adoption

Where	Inside your product, portal or internal application, through a snippet you install
For whom	The people using it: your customers or your employees
Who helps	**Max**, the in-app agent, with each user; **Daleel**, the shared layer, with your team
Measured by	Activation, task completion and sustained useful use

YOU DO NOT OWN IT

FullSession Workforce

Where	In the approved web applications your teams use, through a browser extension your administrator deploys; no change to the applications
For whom	Your teams working across third-party systems
Who helps	**Detect** finds recurring rework; **Coach** guides the approved path in the task
Measured by	First-time-right work and re-measurement on equivalent work

Both routes share the same principles in this report: context at the step where the work happens, authority set by the organization, aggregated reporting by default, and measurement against a comparison group.

fullsession

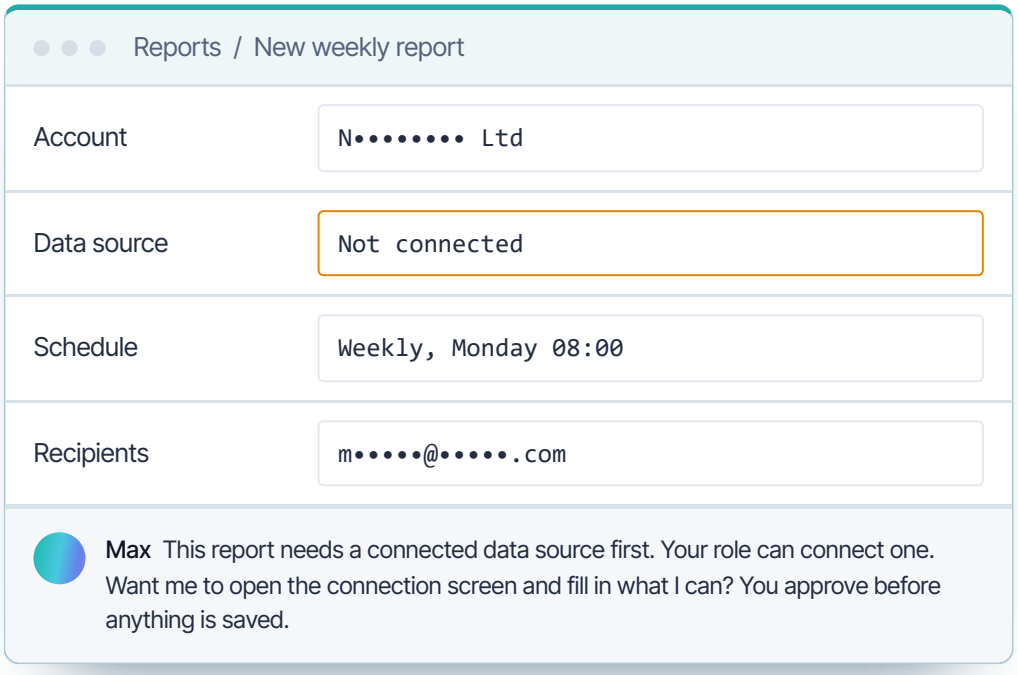
DIGITAL ADOPTION · APPLICATIONS YOU OWN

Max, inside your application

Max works with each user in your application, starting from what they are trying to achieve and resuming where they left off.

Onboarding	Starts from the person's goal and plans next steps by role.
Help on this screen	Answers, navigates or co-drives, so the user sees each step done and can take over at any time; hands off to your team when needed.
Feature discovery	Introduces relevant features and guides first use.
Task completion	Prepares or completes permitted actions and returns a receipt.
Authority you set	Answer, Navigate, Prepare or Complete, per surface and per action. Memory of progress for identified users, within your identity and permission rules.

ILLUSTRATIVE · MASKED VALUES



DIGITAL ADOPTION · APPLICATIONS YOU OWN

Daleel, the shared layer

Max serves the people using your application. Daleel works with your team: it is the shared AI layer behind FullSession's solutions and your copilot for setting up and improving the experience.

Keeps the map current	Maintains the application map with your team: what screens and actions mean, what each step needs, how to check a result.
Finds what users miss	Spots where people stall or skip value, and proposes how to improve the path.
Hands value at risk to Max	When a user has stalled short of a result, Max can offer the fix and check whether the user reached the result.
Makes readiness visible	Shows what is connected, what is understood, what has been tested, what is authorized and what is live, so a gap in one is not mistaken for readiness in all.

> Measuring activation, task completion and sustained use does not require session replay or heatmaps; they are optional.



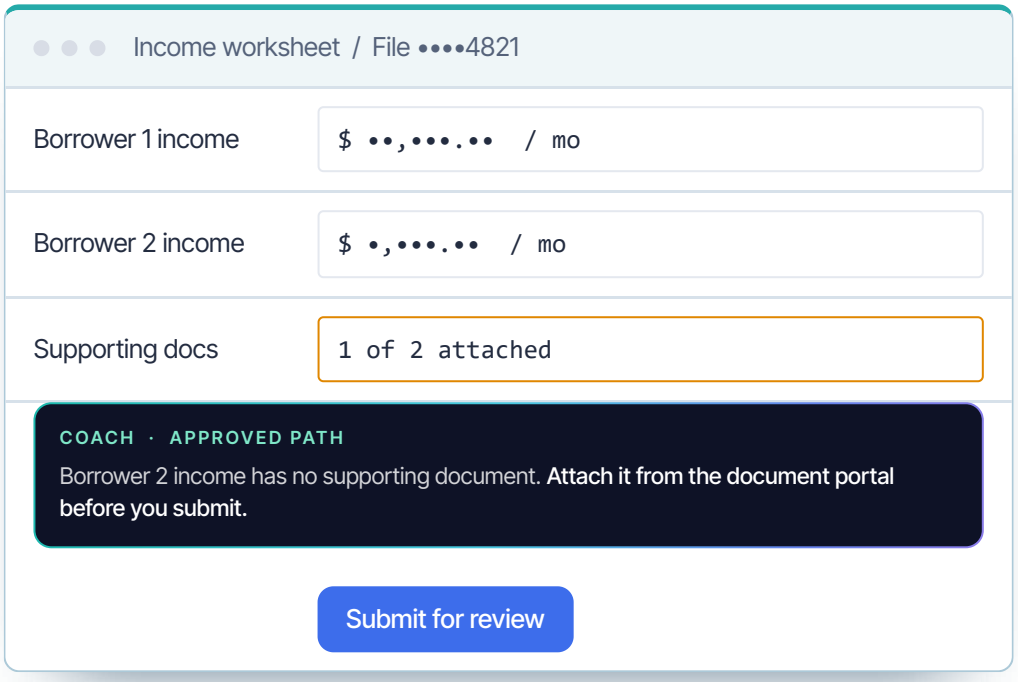
WORKFORCE · APPLICATIONS YOU DO NOT OWN

Detect and Coach, in the apps your team uses

When the application belongs to someone else, help goes where your team already works: the browser.

Detect	Finds recurring rework in a named process across approved web applications and estimates what it costs, with the assumptions shown.
Coach	Guides the approved path inside the task, at the step where cases go wrong. Paired with Detect; supported actions depend on the application and its permissions.
Re-measure	Compares equivalent work before and after, so an observed change can be told apart from an estimated effect.
Deployment	A browser extension your administrator deploys through managed-browser policy, limited to approved applications. No change to the applications themselves.

ILLUSTRATIVE · MASKED VALUES



ACT 3 · TWO WAYS TO CLOSE IT

More connected context. The same authority.

The value of an assistant inside the application grows with what it can understand. Each connected system widens that. What it is allowed to do stays exactly where you set it.

FULLSESSION VIEW · CAPABILITY AND DIRECTION

The application

what screens, actions and steps mean

- Explain this screen and field
- Guide or co-drive the next step
- Catch a missing prerequisite
- Check the result against the app

+ The person

role, goal, progress, within identity rules

- Plan from their goal and role
- Resume where they left off
- Recover a stalled user
- Hand off with the history attached

+ Connected systems

CRM, ticketing, data warehouse, APIs, MCP

- Help that knows the account and the open ticket
- Answer with the record, not a guess
- Show what a stall or a return is costing
- Support a decision: what to fix first, and why

Authority and identity, set by your organization: Answer, Navigate, Prepare, Complete, per surface and per action

Checked against the person's own permissions every time. Context never grants permission.

Not every connection works the same way. Some information is looked up when a person needs it; some has to flow regularly so it can be analyzed. Which systems are connected, and how, is agreed during setup for each deployment, not assumed.

ACT 3 · TWO WAYS TO CLOSE IT

From help to decision support

The same layer that helps one person at one step can answer the question their leader asks at the end of the month: what is this costing, and what should we fix first?

ILLUSTRATIVE · A PRODUCT TEAM ASKS DALEEL

Daleel

Which onboarding step held back the most new accounts this month?

Connect a data source. 41% of accounts that started in October stopped there for more than two days. Accounts with an open support ticket on this step took longer to reach their first report.

From the CRM: most are on plans where an admin, not the person signing up, can connect sources.

Suggested change: ask for the admin at sign-up, and let Max offer to send them the connection step. Roll it out to half of new accounts first and compare time to first report with the other half; a before-and-after alone is an observed change.

ILLUSTRATIVE · AN OPERATIONS LEADER ASKS ABOUT A PROCESS

Workforce / Loan files

Which return reason cost the most active time last month?

Missing second income document. About 90 files returned, roughly 25 active minutes each on the second pass, from an eligible cohort of 1,000 (a separate example from page 20).

The return clusters at the worksheet step, before the document portal is opened.

Suggested change: an approved Coach prompt at the worksheet, rolled out to one queue first, with the other as comparison.

All values on this page are invented to show the kind of answer, not a result. Answers like these depend on which systems are connected and on the permissions of the person asking; aggregated by default.

ACT 3 · TWO WAYS TO CLOSE IT

Measure, help, measure again

The loop is the same on both routes. What makes it credible is the comparison.

FULLSESSION VIEW

01 Baseline

Fix the cohort, the outcome and the horizon (pages 20 and 21). Record what else is changing.

02 Help at one step

Approved guidance at the step where the work goes wrong, owned by a named person.

03 Stagger the rollout

One group first, another later, chosen at random where feasible. The later group is the comparison.

04 Compare changes

Each group's change from its own baseline, then the difference between those changes, with sample sizes and a range.

Observed change is what you see before and after. **Estimated effect** is the coached group's change compared with the uncoached group's, assuming comparable case mix and trends, no one-sided concurrent change and limited spillover. State those assumptions; re-check after the first window, because a gain that fades is not a sustained improvement.

ACT 3 · TWO WAYS TO CLOSE IT

What counts as evidence

Agree these points with the product or process owner and with finance before the results arrive, not after.

FULLSESSION VIEW

01 Eligible population

Which users or cases count, fixed in advance; unfinished and unknown kept in.

02 Coverage

Which applications and steps were observed, and what was not.

03 Outcome and quality

What a correct result is, including a quality check, not only speed.

04 Comparable work

Like-for-like cohorts, task types and complexity across periods and groups.

05 Time, measured honestly

Non-overlapping active time, with waiting kept separate.

06 Reasons, not only counts

Why each return, stall or contact happened, classified against the outcome.

07 Baseline and comparison

A baseline for every group and an uncoached group where feasible.

08 Concurrent changes

Releases, policy, staffing and seasonality logged during the window.

09 Follow-up

A date to re-check after the first window. Report sample sizes and a range, and keep capacity separate from cash until it is used.

ACT 3 · TWO WAYS TO CLOSE IT

Privacy and scope, set before anything runs

An assistant with context is only acceptable if the boundaries are clear to the people it serves and the people it observes.

Approved scope	Digital Adoption runs in the application you install it in. Workforce runs only in the web applications your administrator approves.
Aggregated by default	Reporting is aggregated by default and built around the process and the product, not around ranking individuals.
Controlled access	In Workforce, reviewing an individual session needs elevated permission, a stated reason and an audit entry.
Authority you set	What Max may do is set per surface and action. Learning about the application never grants permission.
Masked by design	Sensitive fields can be masked from capture; the masked values in this report's illustrations show what that looks like.



DEFINED BOUNDARIES

fullsession

WORKSHEET

Are you ready to measure?

For one process or one onboarding goal. Answer with the people who own it.

Every "no" or "partly" is a gap to close before the first measurement, not a reason to wait for a tool. Questions marked DA apply to applications you own, WF to applications you do not; unmarked questions apply to both.

	QUESTION	YES	PARTLY	NO
01	We can name one goal or process, and the person who can approve a change to it.	☐	☐	☐
02	We have a written definition of a correct, complete result.	☐	☐	☐
03	We know which applications and steps the work touches.	☐	☐	☐
04	We can fix a cohort for a period, including unfinished cases and unknown outcomes.	☐	☐	☐
05	DA: we can verify a user's first useful result, not just a page view.	☐	☐	☐
06	WF: we can separate active work time from waiting time.	☐	☐	☐
07	We can say why a case came back, or why a user stalled.	☐	☐	☐
08	We know which connected systems hold the rest of the case.	☐	☐	☐
09	We have agreed what the assistant may do, per screen and action.	☐	☐	☐
10	We can stagger rollout so one group is a comparison.	☐	☐	☐
11	Finance has agreed how capacity becomes a benefit: volume, attrition or avoided expense.	☐	☐	☐
12	We have a date to re-check the result after the first window.	☐	☐	☐

fullsession

CLOSING

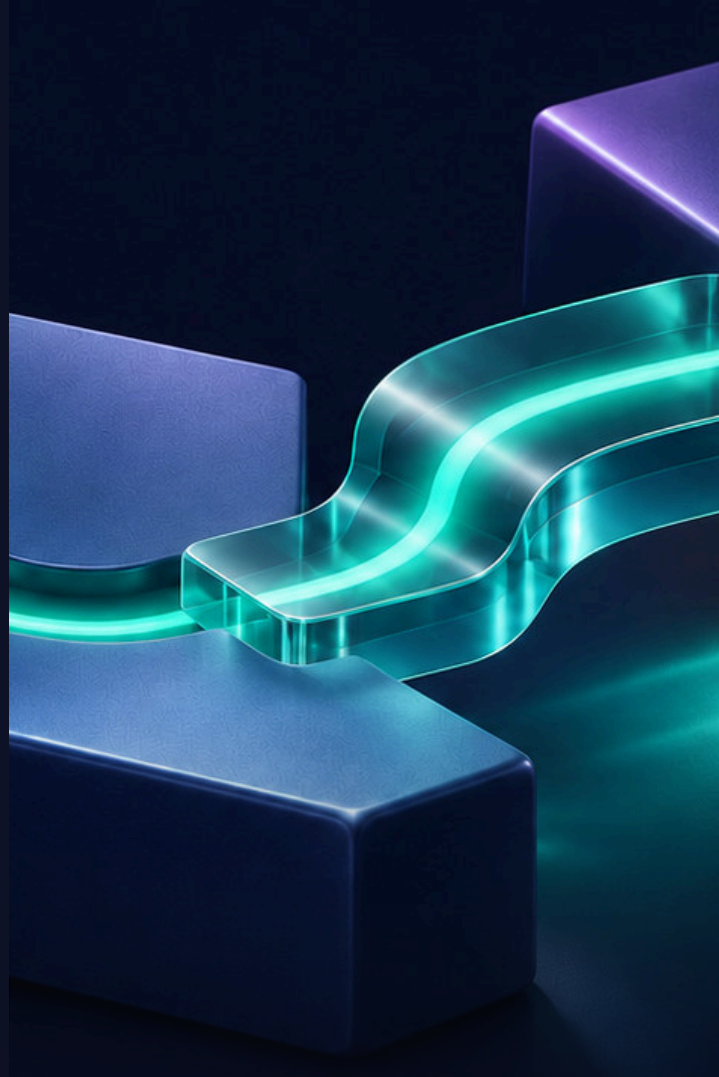
The thread, carried

For organizations adopting AI, the decisions in front of them are where it sits, and how they will know it helped.

Beside the work, it answers questions. Inside the application and the task, with the meaning of the screen, the person's goal and the relevant systems available to it, it can help finish the work, and the organization can see whether it did.

Where you own the application, that means help inside it for every person who uses it. Where you do not, it means help in the browser, at the step where the case goes wrong. In both cases it means authority you set, reporting that respects people, and results measured against a comparison, not assumed.

That is how the context gap closes: one step, one process, one measured result at a time.



THE THREAD, CARRIED



About FullSession

FullSession helps organizations see how work and journeys actually happen in web applications, and act on it. This report covers two of its solutions.

FullSession Digital Adoption

For applications you own. Max works with each user inside the application; Daleel, the shared AI layer, works with your team. fullsession.io/solutions/digital-adoption

FullSession Workforce

For applications your teams use but do not own. Detect finds recurring rework; Coach guides the approved path in the task. fullsession.io/solutions/workforce

Discuss an evaluation

Bring one onboarding goal or one process. We will help you scope it against the measures in this report, including what a comparison would need.

fullsession.io

SOURCES

Research cited in this report

Each study answers its own question about its own population. We do not combine them. Locators point to where each figure appears. All were read at the publisher between 1 and 2 November 2026.

McKinsey & Company. The state of AI in 2026: On the road to ROI. 25 August 2026.
Online survey, 1,719 participants, 97 nations, 4 May to 8 June 2026. Self-reported. Figures: 80% individual productivity; 37% at least some EBIT impact; nearly three quarters of high performers (respondents attributing 5% or more of EBIT to AI and reporting significant value) redesigned workflows. Pages 5, 7. mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Microsoft WorkLab. Breaking down the infinite workday. Work Trend Index special report, 17 June 2025.
Survey of 31,000 full-time knowledge workers, 31 markets, 6 February to 24 March 2025, Edelman Data x Intelligence. Self-reported. Figure: 48% of employees and 52% of leaders say work feels chaotic and fragmented. Pages 5, 8. microsoft.com/en-us/worklab/work-trend-index/breaking-down-infinite-workday

BCG. AI at Work 2025: Momentum Builds, but Gaps Remain. June 2025.
Survey of 10,635 leaders, managers and frontline employees, 11 countries and regions; frontline subgroup n=3,537. Self-reported. Figures: regular use 85% / 78% / 51% (deck p6; regular = several times a week or daily); 36% training enough (deck p13); frontline with vs without clear leadership support, regular use 82% vs 41%, very positive about impact on job enjoyment 55% vs 15% (deck p15, association). Pages 5, 10, 11, 12. bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain

KPMG and the University of Melbourne. Trust, attitudes and use of AI: A global study 2025. doi:10.26188/28822919.
Survey of 48,340 people, 47 countries, November 2024 to January 2025; overall employed sample n=32,352. Self-reported. Figures: two thirds of employees who use AI at work report having relied on its output without critically evaluating it (p6, fig 44 pp75 to 76; item n not reported); 31% intentionally use AI at work weekly or daily (fig 37 p67); two in five report an organizational policy or guidance on generative AI (p71, fig 41); almost one in two employees who use AI report use that contravenes policies or guidelines (p6). Pages 17, 18. assets.kpmg.com (global report PDF)

Microsoft and LinkedIn. 2024 Work Trend Index Annual Report: AI at work is here. Now comes the hard part. 8 May 2024.
Survey of 31,000 knowledge workers, 31 markets, 15 February to 28 March 2024, Edelman Data & Intelligence. Self-reported. Figures, base AI users: 78% bring their own AI tools to work; 52% reluctant to admit using AI for their most important tasks. Page 18. microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part

MuleSoft. 2026 Connectivity Benchmark Report. 5 February 2026.
Vendor survey of 1,050 IT leaders at organizations with 1,000 or more employees, 9 countries, October to November 2025, Vanson Bourne. Self-reported averages. Figures: 957 applications on average; on average 27% connected; 82% cite data integration as one of the biggest challenges when using AI. Sections The expanding application landscape and The connectivity paradox. Page 19. blogs.mulesoft.com/agentic-perspectives/connectivity-benchmark-report

Brynjolfsson, E., Li, D. and Raymond, L. Generative AI at Work. Quarterly Journal of Economics 140(2), 889 to 942, 2025. doi:10.1093/qje/qjae044.
Staggered rollout of an AI assistant to 5,172 customer-support agents at one firm; observed productivity. Figures: 15% more issues resolved per hour on average; about 30% for less experienced and less skilled agents; small speed gains and small quality declines for the most experienced. Final paper figures. Page 25. academic.oup.com/qje/article/140/2/889/7990658

FullSession product statements

Descriptions of Max, Daleel, Detect and Coach on pages 26 to 29 and 34 follow FullSession's published solution pages, fullsession.io/solutions/digital-adoption and fullsession.io/solutions/workforce, as of November 2026. They describe the product, not measured results. Method pages are FullSession recommendations.

DEFINITIONS

Terms as we use them

Activation. The first verified useful result a person came for.

First-time-right. A case completed correctly without returning to an earlier step.

Eligible cohort. All users or cases that entered in a fixed window, followed for the same horizon, unfinished and unknown included.

Observed change. The measured difference between a baseline and a later period.

Context. What an assistant can draw on: the application's meaning, the person's goal and progress, and connected systems.

Readiness states. Connected, understood, tested, authorized and live are separate; one does not imply the next.

Sustained use. Returning at the cadence the job needs and doing useful work.

Repeat touch. Additional active work on a case caused by its earlier step being incomplete or incorrect.

Active time. Non-overlapping minutes of observed work; waiting excluded.

Estimated effect. A coached group's change from its own baseline compared with an uncoached group's, under stated assumptions.

Authority. What an assistant may do, set by the organization per surface and action, checked against the person's permissions.

Capacity. Time freed by fewer repeat touches. It becomes a financial benefit only when used.

PHOTOGRAPHY

p3 Roman Mohren, FullSession
p12, p20, p24, p37 Vitaly Gariev on Unsplash

Used under the Unsplash License. People shown are models in stock photographs; they do not endorse FullSession and are not customers.

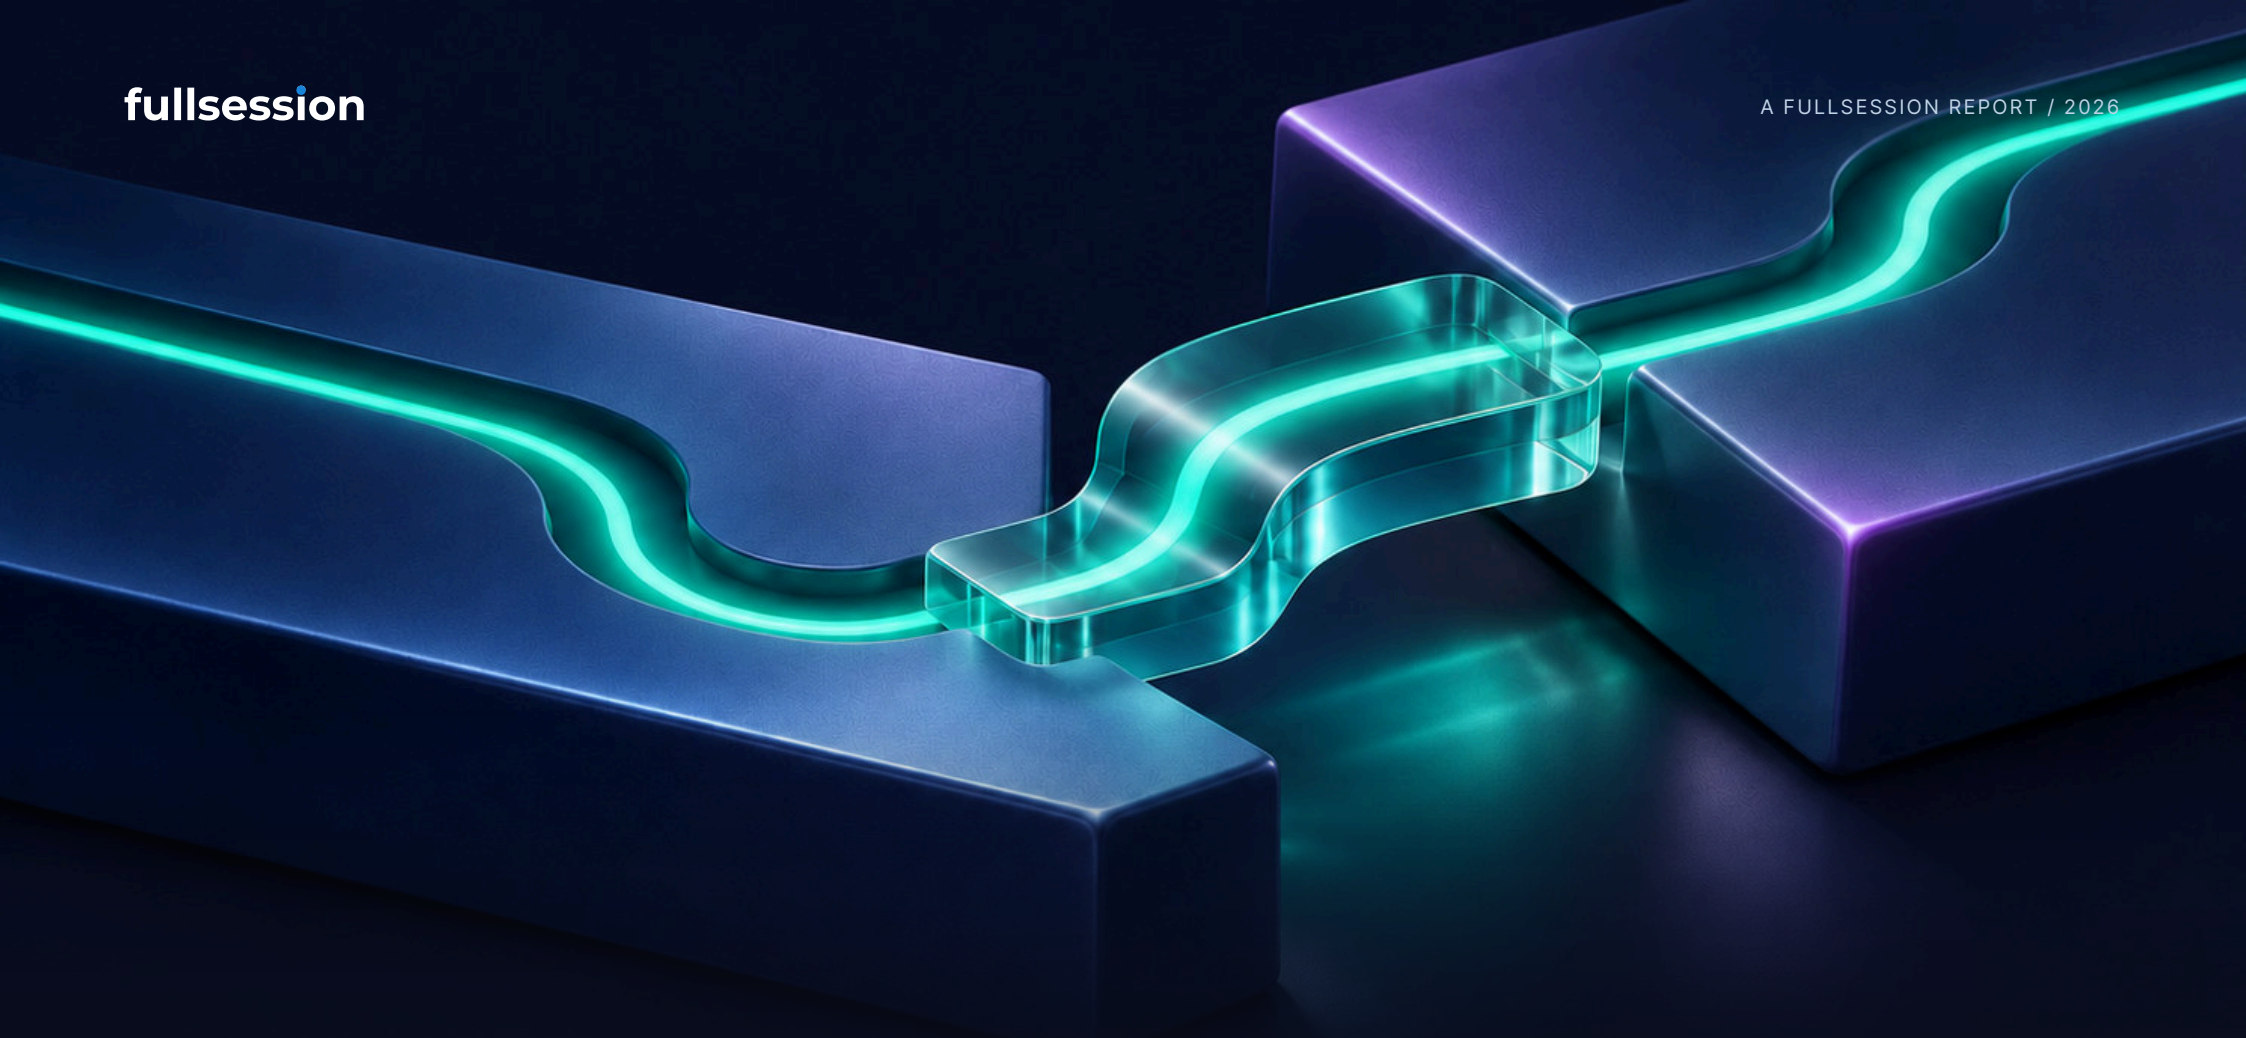
Conceptual artwork
Connected-path illustrations created for FullSession with AI-assisted image generation and art direction. These are visual metaphors, not product screenshots or measured results.

ABOUT THIS REPORT

Published by FullSession, 2026. The research figures are the cited publishers' and are reproduced with attribution; charts are drawn by FullSession from the published values. Illustrative examples use invented values and masked data. No FullSession customer data was used.

fullsession

A FULLSESSION REPORT / 2026



The Context Gap

Why AI inside the application is the next digital adoption decision.

FullSession Digital Adoption · FullSession Workforce

fullsession.io