dixa

BENCHMARK REPORT · 2026

The Dixa Customer Service Benchmark 2026

Platform data from 316 customer service teams, July to September 2026.

TEAMS	PERIOD	SOURCE
316	July to September 2026	Dixa platform data

INTRODUCTION

AI took over chat this year. Email is still waiting.

A third of customer service teams on Dixa now run an AI agent, nearly twice as many as a year ago. Where it runs, it resolves about six in ten of the conversations it gets, and customer satisfaction held. But the typical team still gives it less than a tenth of its conversations, and more than half of what it gets is chat. Email, which is six in ten of all conversations, is where next year's gains are.

This report doesn't ask anyone what they think. It measures what 316 teams did, and shows you where the typical team sits, what good looks like and what the best teams do.

—	The report on one page	03
—	Who is in this report	04
A	Part A. The AI year	05
01	AI agents run in a third of teams	06
02	A realistic AI resolution rate is about six in ten	08
03	Satisfaction didn't drop when teams switched on AI	10
B	Part B. The operation	12
04	Each seat handles a quarter more conversations	13
05	Email replies got almost a quarter faster	15
06	The typical team's busiest month is 15% above normal	17
—	How we measured it	20
—	Questions people ask about this data	21

HOW TO READ THE NUMBERS

Typical
The middle team.

Good
Better than three in four teams.

Best
Better than nine in ten.

THE REPORT ON ONE PAGE

01 **AI agents run in a third of teams, but the typical team gives them less than a tenth of its conversations.**

Section 1

02 **A realistic AI resolution rate is about six in ten, and the best quarter of teams pass three in four.**

Section 2

03 **Satisfaction didn't drop when teams switched on an AI agent.**

Section 3

04 **Each seat handles a quarter more conversations than two years ago.**

Section 4

05 **Email replies got almost a quarter faster this year, and satisfaction held at 81.**

Section 5

06 **The typical team's busiest month is 15% above normal. For a quarter of ecommerce teams it is 70% or more.**

Section 6

Who is in this report, and what it can't tell you

Platform data from 316 customer service teams on Dixa, most of them European and in ecommerce, July to September 2026, compared with the same months of 2025 and 2024. We report the typical team (median), good (top quarter) and best (top tenth), never an average across teams.

Every team in the report handled at least 100 conversations in every month of the quarter. Test, demo and internal accounts are left out. The full method is on the last pages.

WHAT THIS DATA CAN'T TELL YOU

These are Dixa customers: mid-market, mostly European, mostly ecommerce. They are not the whole industry.
Teams that left the platform are not in the earlier years.
AI conversations carry a channel only from late 2025, so the channel view is this year only.
The data can't split satisfaction by whether an AI or a person handled the conversation. Section 3 compares teams, not conversations.
Monthly data hides week-level peaks such as Black Friday week.

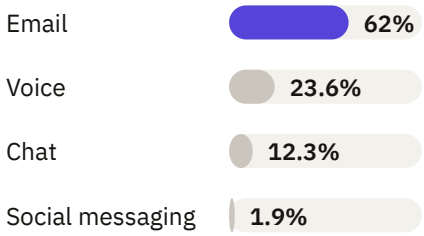
Every split in this report covers fewer than 200 teams. Read the splits as direction, not precision.

The AI year

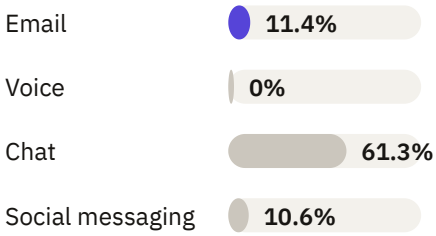
AI agents run in a third of teams, but the typical team gives them less than a tenth of its conversations.

In the third quarter of 2026, 97 of the 316 teams in this report ran an AI agent, up from 53 a year earlier. At the typical one of those teams, the AI agent handled 7.6% of all conversations. A year ago it was 3.8%.

Email is 62% of the work. SHARE OF ALL CONVERSATIONS



AI answers 11.4% of it. SHARE ANSWERED BY AN AI AGENT



7.6%

Typical team

Share of conversations an AI agent handles

19.7%

Good: top quarter

46.3%

Best: top tenth

01 AI AGENTS RUN IN A THIRD OF TEAMS

AI went from a pilot to a standard tool in a year. But it mostly answers chat. The typical team's email, which is most of its work, is still answered by people.

> “Most providers just focused on content-related AI. Mim gave us the ability to do tasks which normally colleagues do.”
>
> JESSICA SCHULZ, CUSTOMER SERVICE SPECIALIST, SMARTPHONEHOESJES

DIXA’S POV

The next gain isn't a better AI. It's giving the AI the email.

If the data had gone the other way, with the typical team above 20%, we would be writing that the gap is closing and showing who closed it first.

WHAT OTHERS SAY

Nobody else publishes AI's share of all conversations. Vendors publish the resolution rate of what AI handles, which is a different number (see Section 2).

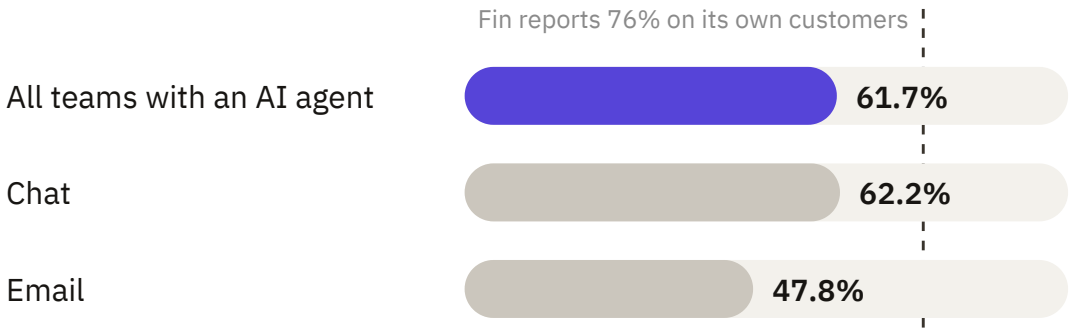
HOW IT'S MEASURED

Conversations an AI agent handled at any point, including ones it handed to a person, divided by all conversations. AI conversations carry a channel only from late 2025.

A realistic AI resolution rate is about six in ten, and the best quarter of teams pass three in four.

Across the 76 teams whose AI agent had at least 50 conversations with an outcome, the typical team's AI resolved 61.7% of what it was given. A year ago it was 55.6%, measured on 37 teams.

The typical AI agent resolves 61.7% of what it gets TYPICAL TEAM, JULY TO SEPTEMBER 2026



61.7%

Typical team

AI resolution rate

78.4%

Good: top quarter

87.4%

Best: top tenth

02 A REALISTIC AI RESOLUTION RATE IS ABOUT SIX IN TEN

Six in ten is the number to plan around. The best quarter of teams show what's possible: 78% and above. Email is harder than chat. The typical team's AI resolves under half of its email.
As resolution rose, handovers to a person fell: the typical team's AI handed over 35% of conversations, down from 42% a year ago. By time since go-live, resolution sits around 60% from the third month on.

> “They're all saying 'full AI.' And you think, okay, but what AI?”
>
> **JESSICA SCHULZ, CUSTOMER SERVICE SPECIALIST, SMARTPHONEHOESJES**

DIXA’S POV

Plan your business case on 60% and your roadmap on 78%.

If the typical team had been under 50%, we would be writing that realistic is well below headline rates, and showing what the first year looks like.

WHAT OTHERS SAY

Fin publishes a 76% average resolution rate on its own customers, without a published definition. Different customers and a different definition, so the two numbers can't be compared directly.

HOW IT'S MEASURED

Conversations the AI resolved divided by AI conversations with an outcome, per team. "Resolved" is the platform's label when the AI closes the conversation without a person. It is not a customer-confirmed resolution.

Satisfaction didn't drop when teams switched on an AI agent.

This page compares teams before and after they switched on an AI agent. It does not compare AI conversations with human ones: the data can't split satisfaction by who answered.

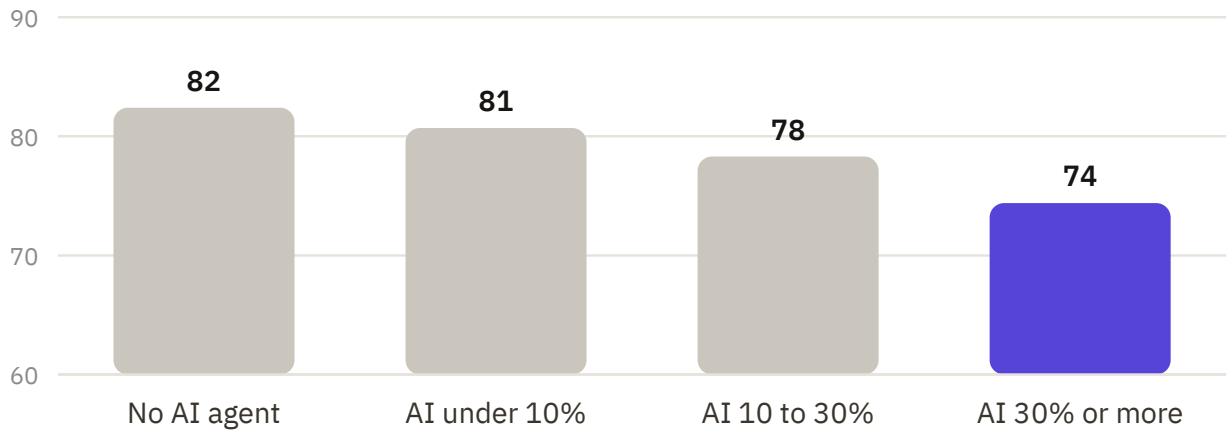
−1.1
Switched on an AI agent
CSAT points · 37 teams

−1.7
Never used an AI agent
CSAT points · 69 teams

+1.1
Already using AI
CSAT points · 30 teams

We followed 37 teams that switched on an AI agent between April 2025 and March 2026. Their customer satisfaction three to six months later was 1.1 points lower than in the three months before. Teams that never used AI moved 1.7 points over a comparable period. Teams already using AI moved up 1.1. None of these moves is large enough to call an effect.

But teams that give AI more work score lower today TYPICAL CSAT BY AI SHARE, Q3 2026



The second chart is the finding we didn't expect. Teams that hand AI a larger share of their conversations score lower today, eight points lower at the top than at teams with no AI agent. Since switching on AI didn't move scores, that gap was most likely there before the AI was.

> "You have to be in the game, really strict on what's going good, what's not going so good."
>
> JESSICA SCHULZ, CUSTOMER SERVICE SPECIALIST, SMARTPHONEHOESJES

DIXA'S POV

The fear that AI hurts satisfaction isn't in the data.

The lower scores at heavy-AI teams are a reason to watch quality as AI takes on more, not a reason to wait. If adopters had dropped two points or more against other teams, we would be writing that satisfaction dips after go-live, and showing how long recovery takes.

WHAT OTHERS SAY

Gladly's 2026 consumer survey of 1,000 US adults found 88% say an AI, or an AI with a person, resolved their issue, and 22% say it increased their preference for the company. That's a survey of feelings, not platform data.

HOW IT'S MEASURED

Share of positive survey answers, weighted by answers, months with 30+ answers. Each team's AI start month comes from its full AI history, back to June 2024.

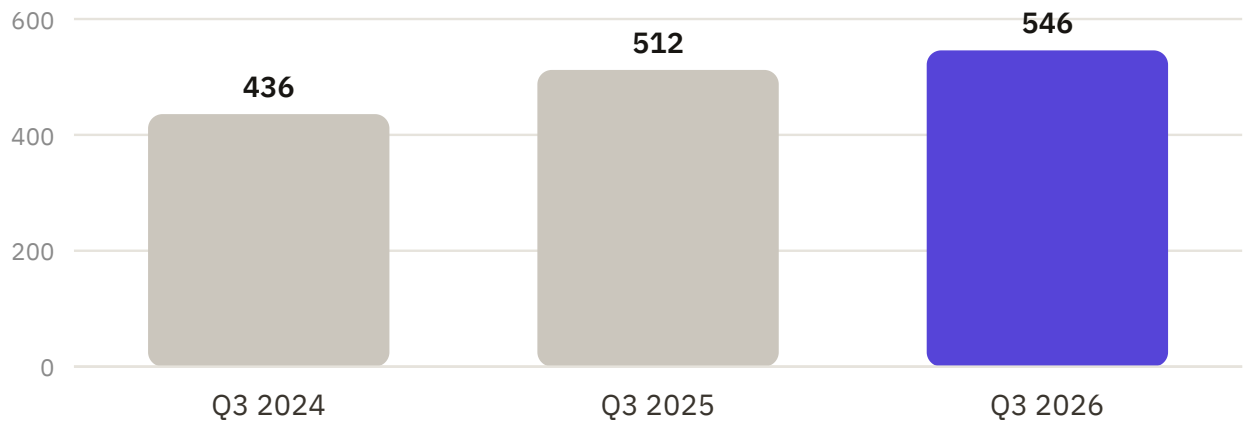
The operation

Teams got faster without getting worse. Email replies got almost a quarter faster this year, and satisfaction didn't move. This part looks at how much each seat handles, what good looks like, and how big the peak really is.

Each seat handles a quarter more conversations than two years ago.

The typical team handled 546 conversations per licensed seat per month in the third quarter of 2026. A year earlier it was 512, and two years earlier 436.

Conversations per licensed seat per month TYPICAL TEAM, JULY TO SEPTEMBER



546

Typical team

Conversations per licensed seat per month

856

Good: top quarter

1,355

Best: top tenth

Productivity per seat rose over two years, and the spread between teams is wide: the good tier handles 856 a month, and the best tenth more than 1,300. In the last year, the typical team added about one seat while its volume held.

The gains don't line up with how much a team uses AI. Teams that give AI 10 to 30% of their conversations gained the most, and teams that give it 30% or more lost ground. For most teams, headcount is still the main lever.

04 EACH SEAT HANDLES A QUARTER MORE CONVERSATIONS

AI SHARE OF CONVERSATIONS	TEAMS	CHANGE PER SEAT, 2025 TO 2026
No AI agent	198	0.7%
AI under 10%	47	-3.2%
AI 10 to 30%	18	21.6%
AI 30% or more	18	-20.8%

> “Working in this way has meant that our team has nearly doubled their contacts per hour, our customer CSAT stays consistently high and our team feels empowered to independently work their way through the day.”
>
> **BECKY HICKEY, HEAD OF CUSTOMER SERVICE, OLIVER BONAS**

DIXA’S POV

AI hasn't yet turned into fewer seats for most teams, because most teams haven't given it enough of the work (see Section 1).

If the number had been flat, we would be writing that the productivity gain hasn't arrived and seats grew with volume.

WHAT OTHERS SAY

Nothing we can cite. The only platform figure in circulation is a 2013 retail study.

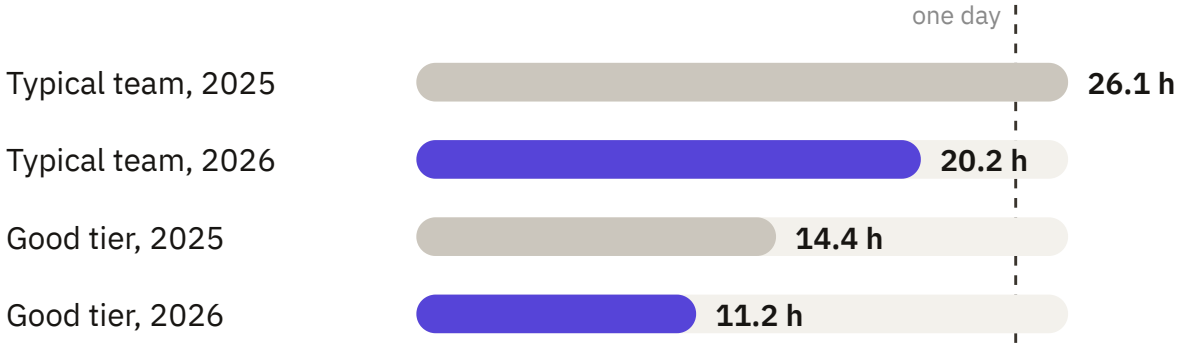
HOW IT'S MEASURED

Average monthly conversations divided by average licensed seats, teams with 3 or more seats. A seat is a licence, not a working hour.

Email replies got almost a quarter faster this year, and satisfaction held at 81.

The typical team's first reply to an email took 20.2 hours in the third quarter of 2026, down from 26.1 hours a year earlier. The good tier, the fastest quarter of teams, replied in 11.2 hours.

Email replies got almost a quarter faster in a year FIRST REPLY BY A PERSON, HOURS



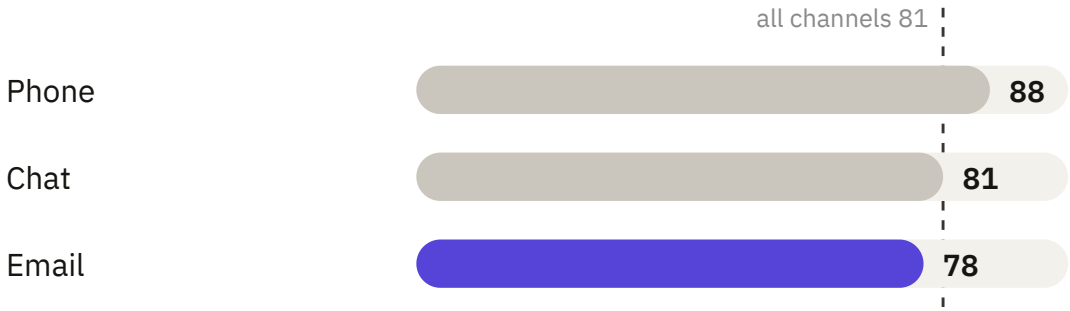
20.2 h	11.2 h	81	88
Typical team	**Good: top quarter**	**Typical team**	**Good: top quarter**
Email first reply	Email first reply	CSAT, share of positive answers	CSAT · best tenth 91

The typical team now answers email within a day, and the good tier within half a day. Chat and phone run on a different clock: the typical team answers chat in 3.5 minutes and picks up the phone in under a minute.

Satisfaction didn't move: the typical team scored 81 this year and 80 last year. Customers rate phone highest (88) and email lowest (78). Small teams score higher than large ones: 86 at teams of 3 to 9 seats, 75 at teams of 50 or more.

05 EMAIL REPLIES GOT ALMOST A QUARTER FASTER

Satisfaction held at 81, and email scores lowest TYPICAL CSAT BY CHANNEL, Q3 2026



DIXA'S POV

Speed improved without hurting satisfaction.

Email is both the slowest channel and the lowest-scoring one, which makes it the obvious place for AI to take on more (see Section 1). If reply times had been flat, we would be writing that email is slower than buyers assume.

WHAT OTHERS SAY

The "12 hours 10 minutes" email reply figure that many blog posts repeat comes from an older survey in which companies reported their own times. Measured on the platform, the typical team is at 20 hours and the good tier at 11.

HOW IT'S MEASURED

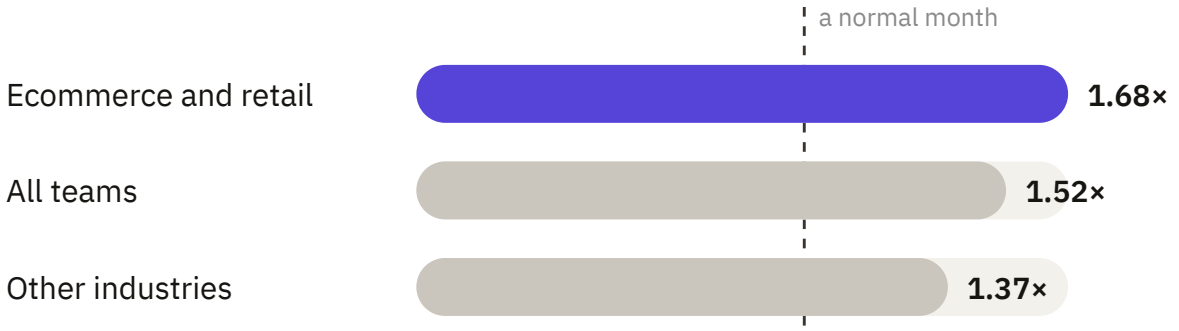
First reply by a person; an AI agent's first answer isn't counted. Each team's figure is an average of its monthly averages, weighted by replies, which runs higher than a median. CSAT is the share of positive survey answers at teams with 30 or more answers; response rates are typically 10 to 20%.

06 THE TYPICAL TEAM'S BUSIEST MONTH IS 15% ABOVE NORMAL

The typical team's busiest month is 15% above normal. For a quarter of ecommerce teams it is 70% or more.

In 2025, the typical team's busier month of November and December was 1.15 times its typical month from January to October. In 2024 it was 1.14 times. The peak isn't getting steeper.

A quarter of ecommerce teams see a peak month 70% above normal TOP QUARTER OF TEAMS, NOV–DEC 2025



1.15×

Typical team

Busiest month vs a normal month, 2025

1.25×

Typical ecommerce team

Busiest month vs a normal month, 2025

1.68×

Top-quarter ecommerce team

Busiest month vs a normal month, 2025

06 THE TYPICAL TEAM'S BUSIEST MONTH IS 15% ABOVE NORMAL

For most teams the peak is a busy month, not a crisis. For a quarter of ecommerce teams it's a different job. Here is how teams absorb it. The typical team adds 7% more licensed seats in November and December, and the top quarter adds 21% or more. Email reply time holds at the typical team but rises by 28% or more in the top quarter. And teams with an AI agent push three times their usual share to it: 6.9% of conversations in the peak, against 2.1% the rest of the year.

DIXA'S POV

The peak is where AI already earns its place: teams that have an AI agent lean on it hardest in November and December.

If peaks had grown steeper, we would be writing about who absorbed them without hiring.

WHAT OTHERS SAY

We found no platform-measured figure we can cite.

HOW IT'S MEASURED

The busier of November and December divided by the team's median month from January to October, for teams with 100+ conversations in all ten months (81 teams with partial years are left out). Monthly data hides week-level peaks such as Black Friday week: a 1.5× month can hide a 3× week.

COMPARE YOUR TEAM

See where your team sits on every number in this report.

Enter five numbers you already have: conversations a month, licensed seats, email first reply time, CSAT, and the share of conversations your AI agent handles. You'll see where you sit against typical, good and best, for teams of your size and industry.

dixa.com/customer-service-benchmark →

How we measured it

WHO IS IN IT

Teams with 100 or more conversations in every month of the quarter. Internal, test, demo and sandbox accounts and suspended accounts are excluded. 316 teams in Q3 2026, 305 in Q3 2025, 257 in Q3 2024. Industry and country come from our CRM and are known for 91% of teams.

AI AGENT

A team counts as running an AI agent in a quarter if Mim or an automation bot handled at least one conversation that quarter. AI share is the share of all a team's conversations that an AI agent handled at any point, including ones it handed to a person. A conversation counts as resolved when the platform records the AI closing it without a person; this is not a customer-confirmed resolution.

SEATS

The average number of licensed agent seats over the quarter. A seat is a licence, not a working hour.

REPLY TIMES

First reply by a person; an AI agent's first answer is not counted. Figures are averages of each team's monthly averages, weighted by the number of replies, which run higher than medians.

SATISFACTION

Share of positive survey answers, weighted by answers, for teams with 30 or more answered surveys in the quarter. Response rates are typically 10 to 20%.

PEAK

A team's busier month of November and December divided by its median month from January to October, for teams active all ten months.

TYPICAL, GOOD, BEST

The median team, the 75th percentile and the 90th percentile, calculated by linear interpolation. Where lower is better, as for reply time, good is the fastest quarter.

SMALL GROUPS

Groups with fewer than 10 teams are not shown. Every group smaller than 200 teams, which is every split in this report, should be read as directional.

REPRODUCIBILITY

We wrote down every definition before computing any number. One script produced every figure, and an independent second run that saw only the written method reproduced them.

Questions people ask about this data

WHAT IS A GOOD AI RESOLUTION RATE FOR CUSTOMER SERVICE?

About 62% is typical. The best quarter of teams pass 78%. Measured on 76 teams on Dixa, July to September 2026.

HOW MUCH OF CUSTOMER SERVICE DO AI AGENTS HANDLE?

7.6% of all conversations at the typical team that runs one. Across those teams, AI handles 61% of chat but 11% of email.

DOES AI LOWER CUSTOMER SATISFACTION?

Not measurably. Teams that switched on an AI agent moved about one point, the same as teams that didn't.

WHAT IS A GOOD EMAIL FIRST REPLY TIME?

About 20 hours is typical. The good tier replies within 11 hours.

WHAT IS A GOOD CSAT SCORE?

81% positive is typical, 88% is good and 91% is best. Phone scores highest (88) and email lowest (78).

HOW MANY CONVERSATIONS DOES A SUPPORT AGENT HANDLE A MONTH?

About 546 per licensed seat at the typical team. The good tier handles 856.

HOW MUCH DOES CUSTOMER SERVICE VOLUME RISE IN Q4?

The typical team's busiest month is 15% above normal. For a quarter of ecommerce teams it is 70% or more.

WHERE DOES THIS DATA COME FROM?

Platform data from 316 customer service teams on Dixa, mostly European and in ecommerce. It is not a survey.