The ultimate customer service metrics guide

Our guide details crucial customer service metrics so you can evaluate your operations with confidence.

zendesk

Metrics help businesses track their performance to achieve and maintain high-quality customer service.

While all customer service metrics are valuable, some data helps organizations narrow in on staff and operational efficiency, as well as customer satisfaction and overall customer experience (CX). Use this guide as your roadmap to strengthen your CX strategy and boost the customer service experience.

Table of contents

1 Customer relationship metrics

In this section, we explore customer relationship metrics and explain how they impact your operations.

In order to achieve a happy and loyal customer base and an all-around great CX, measuring customer satisfaction is a good place to start.

Customers want an easy, fast, and low-effort support experience. To determine your customer relationship metrics, measure both short- and long-term customer satisfaction and happiness.

Topics covered in the section:

- Customer satisfaction (CSAT) ratings
- Customer Effort Score (CES)
- Net Promoter Score® (NPS)
- Social media and churn metrics

Customer satisfaction (CSAT) survey

After each support interaction, use a CSAT survey to measure the customer's satisfaction with the support they received.

At Zendesk, for example, we take a straightforward, effective approach: Ask if the interaction was good or bad. This is a quick, low-effort way for customers to provide feedback. We also give customers the option to add a comment if they want to provide more feedback.

While a simple CSAT survey may not seem like it offers a ton of insight, there's actually a lot you can do with this data.

For example, you can track:

- A customer's CSAT rating over time
- CSAT ratings by customer types
- CSAT ratings by channel
- CSAT ratings by product or service
- Average CSAT ratings for agents and teams

Tracking CSAT also helps you spot trends affecting customer satisfaction across various touchpoints. For instance, if you launched a new support channel, gather feedback to see the customer's reception of the channel.

Satisfaction ratings are also helpful key performance indicators (KPIs) for your agents, as each score links to the representative who resolved the customer's issue. Calcuate the average to give each agent an overall score.

To encourage your teams to hit CX goals, set CSAT targets for your agents and track their performance against those targets. We'll discuss using CSAT ratings to measure agent satisfaction and performance later.

Customer Effort Score (CES)

Like CSAT, CES is a transaction-based survey sent to customers after a support interaction. But while CSAT focuses on how satisfied customers are, you will use CES surveys to determine how much customer effort was required to contact or use support resources.

Companies that use CES typically use this metric instead of CSAT. Whether you use CES or CSAT, the message is clear: Focusing on customer effort reduction should be one of your primary goals to improve the customer experience and overall customer satisfaction.

Net Promoter Score® (NPS)

For even greater insight into the customer relationship, there's the Net Promoter Score (NPS). An NPS survey goes beyond a single support interaction and helps businesses understand if the customer is likely to return, stay loyal, and become a brand advocate.

Unlike CSAT and CES, NPS is not a transactional measurement of customer support but rather a long-term measurement of customers' experience with your company.

The Net Promoter Score asks how likely customers are to recommend your business to someone else.

Based on the 1 to 10 rating, customers get divided into three groups:

- Detractors (1–6)
- Passives (7–8)
- Promoters (9–10)

Detractors are unsatisfied customers who are unlikely to be loyal and often share their unhappiness publicly via social media and bad reviews.

Passives may be perfectly satisfied but likely won't share that with anyone else.

Promoters—your happiest customers—will most likely share their positive experiences with other customers and help you build your business. These are your brand advocates.

Calculate NPS by subtracting the number of detractors from the number of promoters. (Do not include passives.) If all respondents are detractors, your score is -100. If all of them are promoters, your score is +100. While neither of these outcomes is likely, an ideal score is on the positive side of the scale (for example, +50 would be an excellent score).

The NPS survey also lets customers add written feedback about why they chose a certain rating.

Because NPS is not a transaction-based survey, you must choose when and how often to request customer feedback

Social media and churn metrics

Other customer relationship metrics to watch include social media mentions and churn rate or why customers no longer want to do business with you.

Social media listening

Track positive and negative social media mentions to help you understand what people are publicly saying about your company. Social channels to monitor include X (formerly Twitter), Facebook, Instagram, and product review sites.

Using social media-monitoring tools, you can easily collect and analyze customer feedback.

Use this feedback to determine:

- How many comments appear to be written in moments of frustration, perhaps after a poor customer experience in person or online?
- How many are technical or account-specific questions?
- How many comments provide positive or negative feedback?
- How many questions can be answered using links to existing help content?
- How many brand mentions require or would benefit from a response?
- What time of day are customers most active on social media?

Customer churn survey

Another excellent way to collect customer feedback—especially if you're a subscription-based business—is to ask your customers why they're canceling their accounts. From this data, you can create a report of your churn activity.

Key takeaways

1. Determine if your customers are happy after customer support interactions by measuring CSAT and CES. Use NPS surveys to gauge long-term customer loyalty and possible brand advocacy.

2. Set CSAT targets for your agents and teams to keep them focused on the customer experience.

3. Monitor your social media channels for positive and negative feedback, and send out customer churn surveys when a customer cancels their subscription.

2 Essential performance and efficiency metrics

Measuring individual and team performance is essential for maintaining the efficiency of your customer service organization.

This section looks at productivity metrics, ranging from tickets solved to reply time, that can help businesses evaluate how well their teams are performing.

Topics covered in the section:

- Tickets solved
- Issues by resolution area
- Reply time metrics
- Resolution time metrics
- Ticket reopens
- Next issue avoidance

In this section, we look at the most important performance and efficiency metrics to aid your support strategy.

Tickets solved

Many customer service teams set a daily ticket-solved target for their agents. Based on this number, organizations can track how well their agents and teams are performing based on the percentage of the achieved target.

It's useful to look at this by team and by individual agent. Of course, some days, an agent will be under (or over) the ticket target. Tracking the daily average over time is most important for spotting trends in agent performance and resourcing needs.

Calculate this metric by dividing the total number of tickets solved in the last 30 days by an average of 18 working days a month.

Tickets solved vs. open tickets

Another way to use this metric is to compare tickets solved to the number of open tickets. This comparison can highlight how well the team keeps up with the ticket queue. We'll go more into the ticket queue and backlog in the "Ticket backlog" section.

Issues by resolution area

Customer feedback in the form of support requests may be the most valuable information you have for improving your products and customer experience. Take advantage of this data point and categorize your support requests into product areas with support issues, like account creation, shopping cart abandonment, and so on.

You can then segment support requests and drill deeper into, for example, how many tickets are being generated, how long they take to solve on average, and the average CSAT ratings for each product area.

Reply time metrics

How quickly you reply to your customers significantly affects CX and customer satisfaction. For this reason, it's essential to focus on reply time metrics.

First reply time (FRT)

First reply time (FRT), or first response time, is the time from when a ticket is created to when an agent first replies to the customer. In other words, FRT measures how long it takes a real human, not an automated reply, to contact the customer.

Your performance target for this metric should align with customer expectations for first reply times, which are generally considered to be:

- 24 hours for support requests submitted via email and online forms
- 60 minutes for social media
- A few seconds for chat messaging

Of course, responding faster and exceeding these expectations is even better.

FRT is a good indicator of how efficient your support processes are and how well they hold up in handling fluctuating incoming ticket volumes. Monitor FRT by support channels to ensure you meet expectations.

Average reply time

Average reply time is the average amount of time for a support agent to respond to the customer when solving a support request.

Consider the following to determine what an acceptable average reply time is:

Use your FRT and average agent touch metrics to guide your average reply time goals. These metrics can highlight opportunities for improving employee performance, such as using better docs or tools or undergoing more training.

Look at your CSAT ratings and comments to identify any common customer concerns regarding your reply time.

Number of replies

The number of replies is the number of times an agent replies to the customer. This is an important metric because it's a good indicator of customer satisfaction. A lower number of replies could indicate good agent performance, while a higher number of replies can lower CSAT significantly.

Resolution time metrics

Customers like quick resolutions just as much as quick replies. Several different resolution time metrics help businesses identify how long it takes for agents to resolve issues.

First contact resolution (FCR)

First contact resolution (FCR), also known as first call resolution or first touch resolution, measures the percentage of support issues that agents resolve in a single interaction.

A single interaction is defined as one live-chat session, one telephone call with an agent, or an email inquiry an agent solved with a single reply. FCR greatly impacts customer satisfaction, helps reduce operating costs, and improves agent job satisfaction.

However, if you have a high percentage of FCR tickets, this may indicate many low-complexity support issues that you could resolve with self-service resources. Agents should work on more complex issues, not those a customer can easily solve by reading an FAQ page.

Resolution time

Also called "time to resolution" and "mean time to resolution/resolve (MTTR)," resolution time measures the time it takes for a support issue to be solved.

A support issue may be solved more than once before a ticket is closed (a ticket can be reopened before it's closed), and there can be multiple resolution times in the life cycle of a ticket.

The time it takes to solve an issue the first time is called "first resolution time." It does not count as FCR if multiple interactions existed before an agent solved the ticket. When an agent finally solves and closes a ticket without reopening it, it's known as full resolution time.

This is yet another metric that directly affects CX and customer satisfaction. But focusing solely on resolution time isn't necessarily a good idea. Speed doesn't always equal quality. We'll discuss this more later in this guide.

Resolution time metrics (part 2)

In addition to understanding how long it takes to resolve a support issue, determining how much effort went into that resolution is also important. Doing so reveals the actual time the agent spent resolving the issue. Total resolution time doesn't account for this since it includes stages in the ticket life cycle when the agent isn't working on it—like when a ticket is sent back to the customer for information and marked as pending.

Luckily, other metrics evaluate overall effort that can fill in the gaps.

Handle time

Handle time is the time that an agent spends working on a single support interaction. This metric is also referred to as average handle time (AHT), average call handle time (ACHT), and contact handle time.

Agent touches and ticket comments

An agent touch is any update an agent makes to a ticket. This includes ticket status changes and other updates that customers may not be aware of. A large amount of agent touches or comments may indicate that an agent needs more training if it takes them more touches than the team average, for example, to solve tickets.

Requester wait time

Requester wait time measures how long a ticket is set to ”new,” “open,” and “on hold.” Or, in other words, the time agents and support teams spend resolving a support issue.

When a ticket is pending, the agent needs more information from the customer. When a ticket is set to “solved,” it's up to the customer to close the loop or ask another question. An example could be a response to the question: Did our support solve your issue?

Unlike resolution time, this metric does not include all the stages in the ticket life cycle—only the stages when the customer is waiting to hear back from an agent (new, open, and on hold). Therefore, it's a measure of effort and an indicator of the customer experience during the support interaction.

Ticket reopens

A ticket reopen occurs when the ticket status is changed from “solved” to “open.” How often solved tickets are reopened may indicate that agents aren’t fully resolving support issues. This may be due to a focus on first contact resolutions and speed over quality. So, it’s best practice to routinely monitor your team’s ticket reopens.

You can compile your ticket reopen results into reports that should include:

- The total number of reopens
- The average number of reopens
- The percentage of tickets with reopens

Ticket reopens are more likely when agents deal with more complex support issues and may be higher for escalated tickets.

Next issue avoidance

Next issue avoidance examines how many customers have submitted multiple support requests for the same product or subject area. This is important because operations focusing solely on resolution speed and agent handle time can miss the bigger picture: solving the customer issue the first time.

Compare this metric to your first contact resolution results to see if there’s a correlation. If your FCR and next issue avoidance rates are suboptimal, examine how your agents interact with your customers to identify opportunities for improvement.

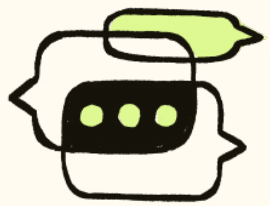

By listening to your community, you may better understand why you’re getting multiple support requests for the same product or subject area. If you’ve set up a help center and community forum, routinely analyze the top 10 community posts to see what your customers are asking and talking about.

Key takeaways

1. Set targets for tickets solved per day, but keep in mind that customer satisfaction is your goal, and hitting your numbers doesn't always ensure that.

2. Quick first reply times and first contact resolutions lead to customer satisfaction, but make sure these metrics aren't masking another problem, such as ticket reopens. Focus on quality over speed and efficiency.

3. Take every opportunity to engage with customers and to anticipate and solve the next issue they may face.

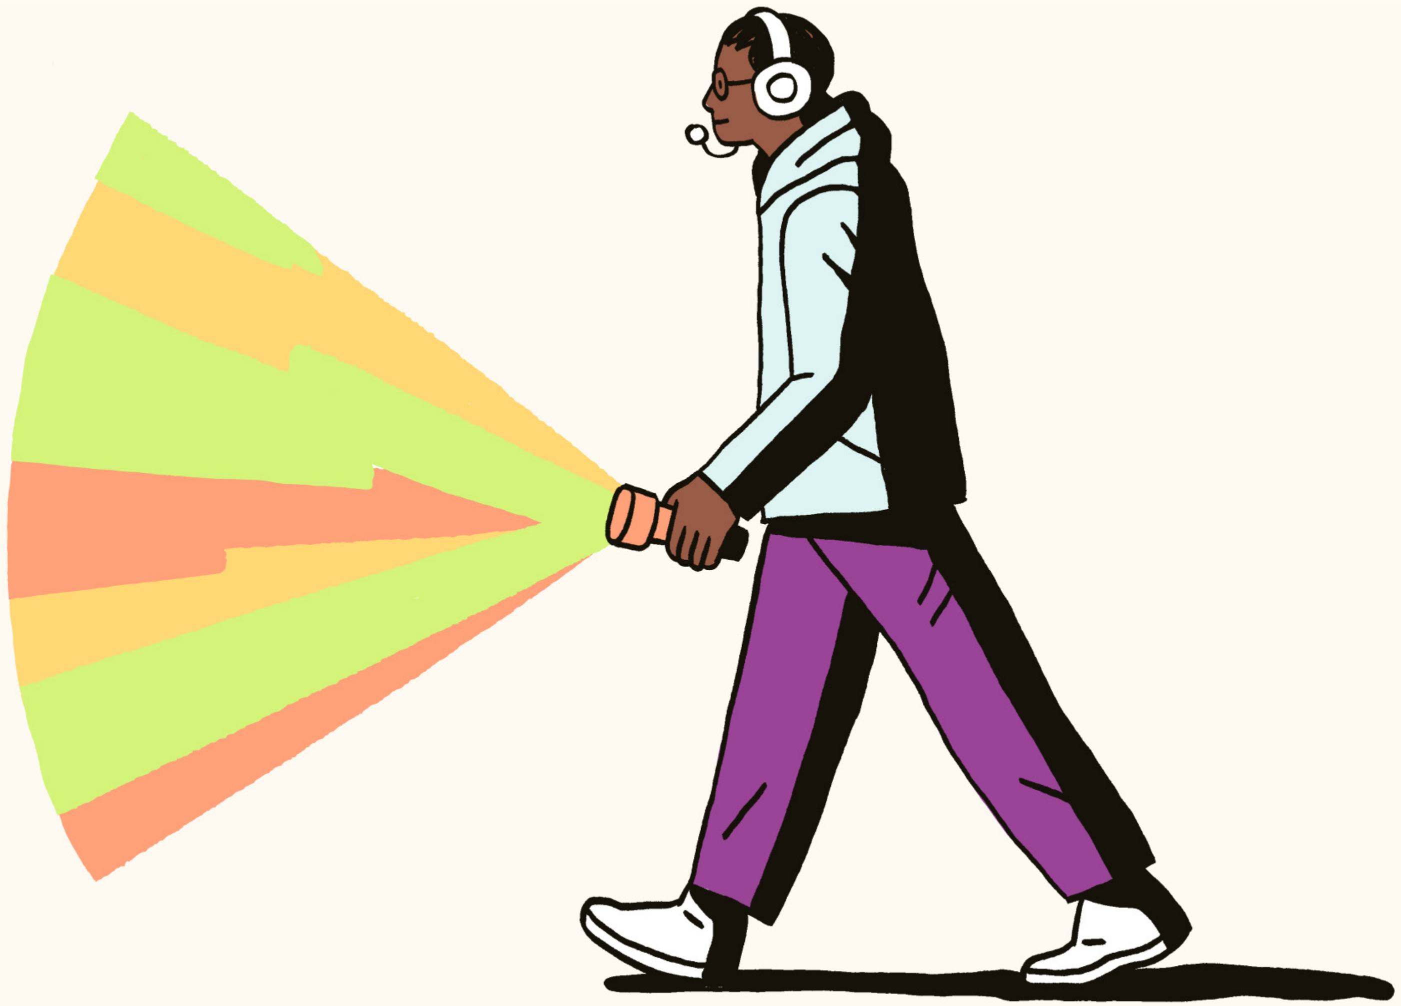

3 Putting your metrics to work

Most of the metrics we just looked at are easy to access and analyze. However, your approach to customer service determines what metrics are important to your business, how you view your organization's performance and customer relationships, and what your goals and priorities are.

Topics covered in the section:

- Balancing quantity with quality
- Setting operational baselines
- The uses of average and median
- Tracking metrics in Zendesk
- A few examples: Key metrics by support channels

In this section, we consider factors that can help shape your support approach and some of the practical ways to set goals for and actively manage your key customer service metrics.

Balancing quantity and quality

Businesses track metrics to clearly understand operational efficiency, team and agent performance, how customers feel about the company, and how likely they are to continue doing business with it.

It's important to find a balance between the metrics that provide the quantitative data that helps us understand our organization's throughput (tickets solved per day, resolution time, average handle time) and the metrics that reveal the more qualitative aspects of delivering customer service (CSAT and customer effort).

If you align your operational baseline with customer needs, satisfaction, and loyalty in mind, the appropriate metrics come more sharply into focus.

Don't get blinded by fast and first

As we mentioned earlier in this guide, first reply time can improve customer satisfaction, and first contact resolution can improve customer and employee satisfaction and reduce operational costs.

The challenge is to make sure you're not blinded by what appear to be great performance stats. Focusing only on being fast (a low FRT) and first (a good FCR) can open the door to shortcomings and lapses in customer service.

Remember to prioritize quality over quantity. A high FCR rate is great if your ticket reopens and next issue avoidance rates are low. When customers reach out to you for support, make sure agents take the time needed to solve the issue and help customers better understand how to use your products. This approach can avoid follow-up tickets and improve your next issue avoidance numbers.

Is self-service the better option?

A high first contact resolution percentage may also indicate you have a lot of low-complexity support requests that could be better handled by self-service options. Self-service can help reduce costs and reserve your agents' time and expertise for more complex issues.

Put your metrics to work to figure out where self-service will have the biggest impact. For example, see how many tickets are created for each product area and compare those numbers to the average handle time for solving those issues.

The resulting information should not only indicate the product areas that would benefit from a self-service option, but also which areas need dedicated support expertise, more training to help agents through the issues, or, more likely, changes to the product for greater ease of use.

Gather feedback without the ask

Use your metrics as a source of indirect feedback to monitor the status of your service operation.

In addition to looking at the number of tickets by product area, you can measure the average handle time, number of agent updates and comments, and full resolution time. Compare these figures with past CSAT ratings and direct customer feedback to gain insight into recent performance without asking your customers.

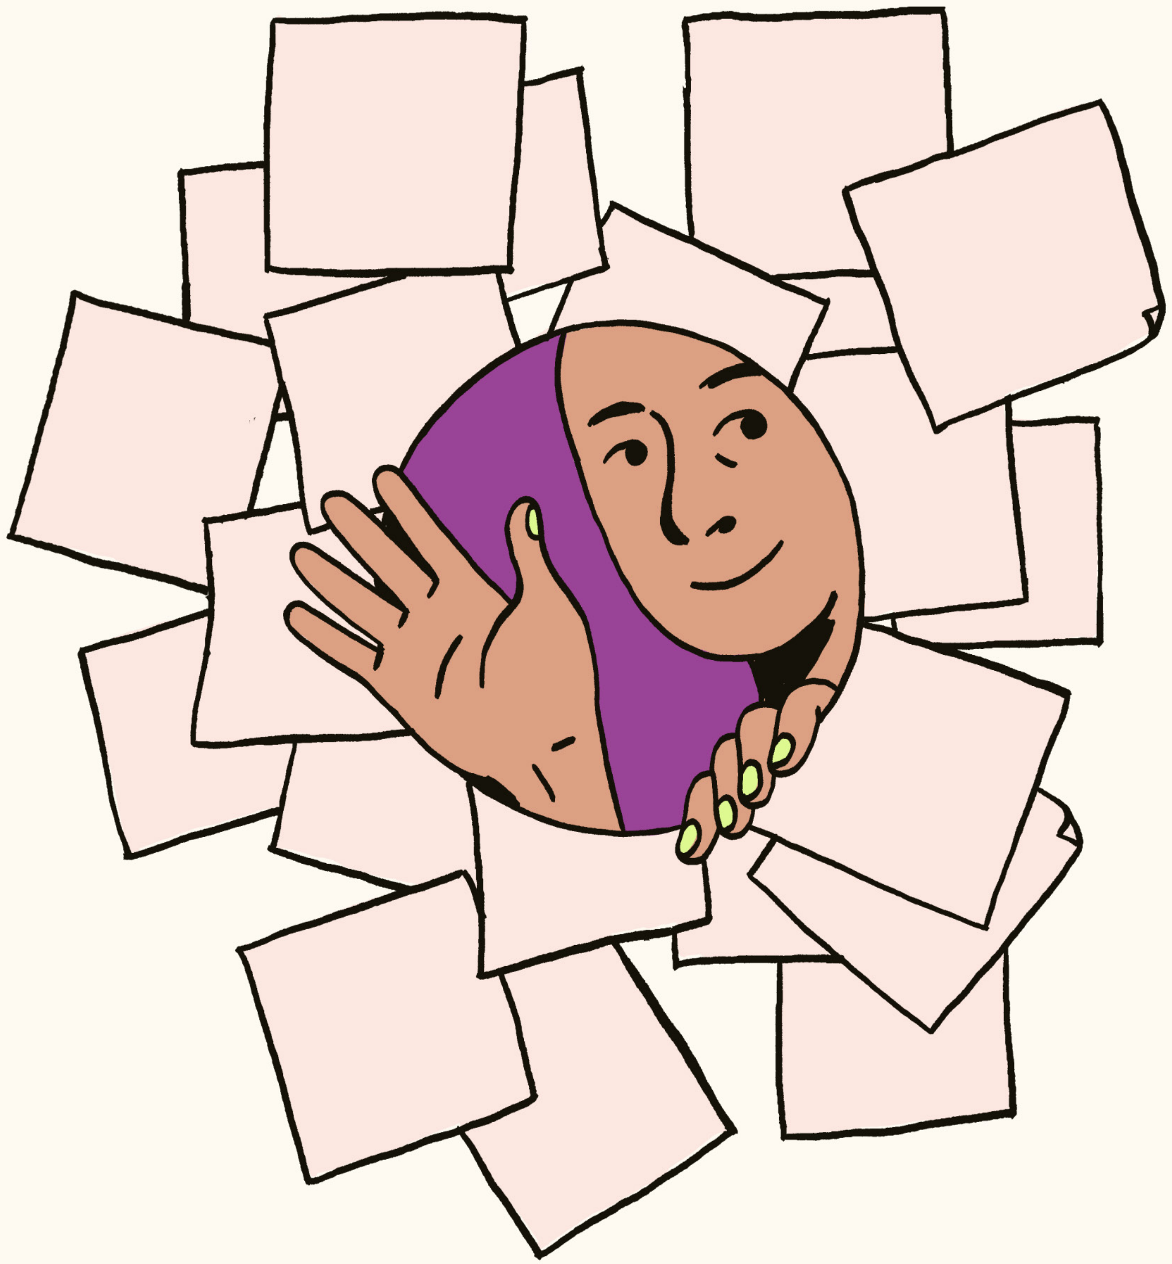

Setting operational baselines

To measure your success, establish baselines for your key performance metrics. The best way to get started is to look at your historical data. Here are key metrics that you should consider setting baselines for.

Tickets solved per day

Set a target for tickets solved per day as a way to measure performance. By doing so, you can report on hitting target percentages by agent, for teams, and as an organization.

To determine a reasonable target, start with your historical data. Look at the numbers for both your agent with the highest solved rate per day and the agent with the lowest. Determine the median and validate it across other agents' daily solved rates. This baseline should challenge your teams to hit a realistic goal.

Customer satisfaction ratings

Set an overall CSAT goal for your organization and agents. For example, this can look like a 90 to 96 percent CSAT rating based on agent experience, job level, and areas of expertise. If your teams are hitting their CSAT targets, consider setting a goal for your CSAT response rate.

First reply time

First reply time performance goals vary by support channel and other factors such as issue complexity. To set FRT goals, start by reviewing your historical performance. Review your current and previous FTR, and see if numbers are trending up or down.

Once you determine this information, select an achievable performance goal that reflects your team's historical performance and also aligns with customer expectations. And if you're aiming to decrease your FRT, set a stretch goal.

First contact resolution

To set an FCR baseline, take the same approach as FRT. Look at historical trends to establish an achievable goal.

Agent touches per ticket

Since CSAT declines with the number of ticket replies, agent touches are yet another area that benefits from setting a performance benchmark.

Start by looking at the median number of replies across your teams and establish that as the baseline. Then, monitor the tickets above the median to identify opportunities to improve performance, such as providing more agent training, internal documentation, or support tools.

If your median reply time isn't syncing up with your customer expectations and is driving down your CSAT ratings, set a more aggressive baseline.

Tickets solved in Tier 1

In organizations with a tiered customer service model, a Tier 1 team—or non-specialized support agents—typically solves 70 to 80 percent of support issues.

The other 20 to 30 percent of support issues are more complex and costly, requiring more time for agent troubleshooting and a higher technical skill. Gauge how often your general support agents solve support issues, and use that information to set an operational baseline. If issues are moving past Tier 1 support more often than normal, update your processes or establish training programs.

Cost per contact

Cost per contact measures the cost of the effort an agent puts into solving a customer support issue. A formula that can help you figure out your cost per contact is dividing your organization's annual budget by the total number of help requests across all products and channels. You can then break that number down by team to set more discrete costs.

As your organization grows and provides more blended services, this will require more nuance to maintain accuracy. But knowing—and setting a goal number for—your cost per contact is an important part of demonstrating to your stakeholders that they are getting a well-managed return on their investment.

Support channels

Many of your support channels provide metrics specific to those communication forms. For example, when using a phone support channel, set baselines for metrics such as abandonment rate, average wait time, average hold time, number of transfers, and so on. Monitoring these metrics can help you set goals for improvement—and when you improve the metrics, the customer experience will naturally improve as well.

The uses of average and median

Before we get into how to look at these metrics as reports and dashboards, let's talk about average (or mean) and median. As a refresher, the average is obtained by dividing the sum of a set of quantities by the total number of quantities. In contrast, the median is the middle number in a series arranged in order of value or size.

Calculating the average is useful to find the midpoint in a set of numbers, but if some numbers are significantly higher or lower than others, the results can easily be skewed. In instances like these, the median is the better measurement point.

A good rule of thumb is:

- If the data you're comparing is mostly uniform, use the average.
- If your number set has some outliers, use the median or filter out values that otherwise skew the results.

Finding the average is useful when looking at metrics with few outliers, such as phone and chat first reply times. It is also a good way to measure general ticket volume.

Key takeaways

1. Use your metrics to measure the quantitative and qualitative aspects of customer service, but stay focused on improving your customer relationships.

2. Effort reduction leads to higher customer satisfaction. That's why metrics like first reply time and first contact resolution are so important.

3. Set operational baselines to help you establish realistic performance goals for your teams and agents.

4. Use a robust set of reporting tools to continuously monitor your organization's performance against the baselines you've set for managing your ticket queue and the health of your customer relationships.

4 Ticket backlog

Your ticket backlog is your total number of unsolved tickets. This is important to follow because it provides insight into your incoming ticket volume and how well you keep up.

As we've discussed, speed is important, but not at the expense of quality. Sometimes, support issues take longer to solve than the customer expects or the performance targets you've set for your team. That said, the longer it takes to solve a customer's issue, the more likely your CSAT will suffer.

Topic covered in the section:

- Reporting on backlog

In this section, we touch on how you can monitor your ticket backlog to strike the perfect balance between speed and quality.

Reporting on your backlog

Backlog tracking can help you understand how your backlog impacts your team and customers and why a backlog happens in the first place. Here are the two types of backlogs and how you can manage them.

Your current backlog

A backlog consists of unassigned and assigned tickets your team hasn't solved yet.

When tracking your current backlog, cross-reference its volume with ticket age and first reply time. A large backlog is not necessarily bad if your team is capable of high throughput, but one with many high-priority tickets may indicate there's a problem.

Historical backlog

Looking at the historical and current backlogs helps you spot trends in ticket volume and team performance.

Managing your backlog

Support requests stuck in the backlog means customers wait longer to get help. If this happens, make sure to keep the customer informed about the status and progress of their request. Doing so keeps customers in the loop, which can improve your CX and CSAT.

Tickets in limbo may drive up your backlog numbers if, for example, the customer hasn't replied to a request for more information or confirmed that a proposed resolution worked.

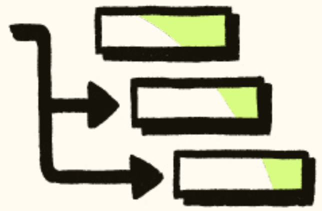

To handle these kinds of tickets, create an automation to remind the customer of the pending status of their issue. If they don't respond after several days, you can automatically resolve the ticket to get it out of the backlog and queue. Additionally, review your entire backlog periodically to see how your team performs compared to your expectations.

Key takeaways

1. Ticket backlogs happen. When they do, manage customer satisfaction by informing customers of the status of their requests.

2. Your total ticket backlog is a combination of unassigned and assigned unsolved tickets.

3. Monitor your backlog to spot ticket volume trends so you can better manage your resources and measure overall team performance against your solved ticket baseline expectations.

5 Measuring agent satisfaction and performance

Agents who are satisfied in their roles deliver better customer service. They're more engaged, motivated, and productive and are less likely to churn.

Managing turnover in a customer service organization is always a challenge. Regularly measuring job satisfaction gives you the feedback you need to create a better job experience for everyone on the team.

Understanding your agents' satisfaction and their performance stats can determine how well they're doing today and where they may be heading in their careers.

Topics covered in the section:

- Measuring agent satisfaction (ASAT)
- Evaluating agent performance

In this section, we look at how to measure agent satisfaction (ASAT) and the metrics typically used to evaluate agent performance.

Measuring agent satisfaction (ASAT)

Measuring agent satisfaction is done using a survey that asks how much the support team employees like or dislike their jobs. It can look something like this:

There are three survey questions. The first question is multiple choice, which can be evaluated as numerical data on the same scale as NPS. The second and third questions are open-ended, allowing employees to provide more information about what they like and dislike the most about their jobs. The feedback collected through these follow-up questions can help you take action based on the survey results.

Here's the general process for conducting ASAT surveys:

- Survey agents, leads, and managers every quarter.
- Have someone outside the support team conduct the survey and evaluate and summarize the data.
- Act on the survey feedback.

Evaluating agent performance

Many of the key metrics used to evaluate an organization's performance also apply to individuals.

Here are useful metrics for evaluating agent performance:

- Number of tickets solved per day
- CSAT ratings
- Agent touches
- Ticket volumes by channel
- Number of escalations
- Areas of focus

Set a baseline for these metrics to track them easily over time.

First reply time isn't on this list because it's often outside of the agent's control and may be a better indicator of overall process efficiency. CSAT ratings can sometimes be unfair to agents because of other factors that have nothing to do with the quality of service (e.g., a customer is frustrated with a product that doesn't include a feature they want).

You can more accurately assess an agent's performance by using the negative comments left in CSAT surveys.

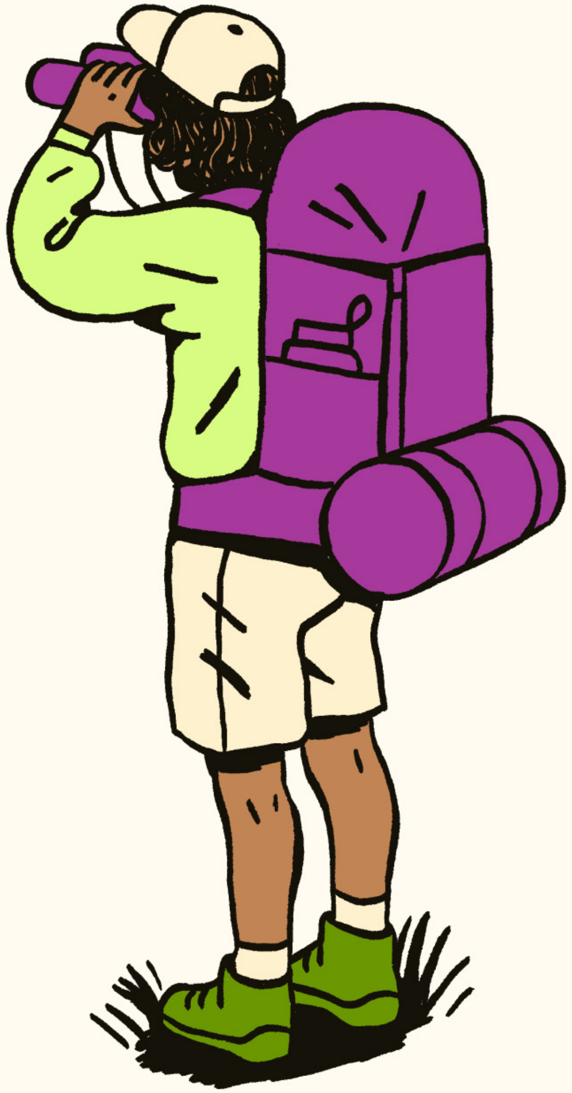

Key takeaways

1. Agents who are happy and satisfied in their roles deliver better customer service. They also stay on the team longer, reducing turnover.

2. As you do with your customers, routinely survey your agents for feedback. Find out if they like their jobs and careers.

3. Consider factors out of an agent's control when evaluating their performance.

6 Measuring your self-service channel

Self-service is a great way to improve overall customer satisfaction (customers like to help themselves) and scale your customer service organization.

A help center or knowledge base is a self-service channel customers can use to resolve issues independently. But if your resource hub is inadequate or customers can't access it, this may negatively affect the overall customer experience.

Topics covered in the section:

- Views and engagement metrics
- Knowledge base and community metrics
- Search metrics
- Self-service score

In this section, we examine the metrics that help determine if customers can find your self-service content, if they find it useful, and if your help center is an effective self-service portal that deflects tickets.

Views and engagement metrics

These are the typical metrics you track for help center performance (or for any website) using both Google Analytics and the knowledge base reporting available in a customer service tool like Zendesk.

Views

This is the number of page views in your help center.

Users

This is the number of unique users who visit your help center. Each visit to your help center counts as a session, and each session usually results in multiple page views. Tracking the number of users visiting your help center can provide insight into its use compared to the total number of views in a specified period. For example, a total of 10,000 monthly views compared to 1,000 unique users within that same period tells you that those users view 10 pages per session on average.

Percent of new sessions

Understanding how many new versus returning users visit your help center can help you determine if you have or need to create content to address user needs. For example, rolling out a new product may result in a spike in new users. Make sure you have self-service content available that adequately covers the product and enhances the customer experience.

Average session duration

The average duration of a user session tells you how much time users spend clicking through your help center and reading specific articles and FAQs. Ideally, they spend enough time reading through the provided information. If they don't, this tells you that your content isn't engaging or your help center lacks the necessary information.

Pages per session

This is the average number of pages viewed during a session on your help center and tells you how much of your self-service content is being used.

Bounce rate

This is the percentage of single-page sessions in your help center. A bounce means that the customer left your help center after viewing the first page they landed on. A user may have visited the help center unintentionally or may have yet to find the information useful.

Knowledge base and community engagement metrics

Most customers want to be able to solve their problems independently. Offering a knowledge base and other self-service capabilities helps improve CX and eases the pressure on your support team.

When constructing and analyzing your knowledge base, look for the following engagement metrics:

Votes

This is the total number of upvotes and downvotes applied to your help center articles. You can view vote totals separately or combined. This is an excellent way to track customer sentiment about your content. Make sure to review downvoted articles and improve them.

Subscriptions

This is the number of customers who subscribe to receive updates when somebody adds new comments or content to your help center. Subscriptions are a good measure of customer engagement.

Comments

Customers typically leave comments when they need more information or clarification. While this might indicate that the provided content is insufficient, comments highlight if you need to add more content and use cases to your knowledge base. Replying directly to comments is also a good way to help customers before they submit a support request.

Community posts

Monitoring the number of community posts can help you determine if your self-service portal is effective. Your help center should be the go-to place for customers to find information, engage with other customers, share their expertise, and learn from each other.

Search metrics

If your customers can't find the information they're looking for in your help center, you may need to improve your self-service offerings. Fortunately, some metrics track what customers search for and the actions they take after searching for answers.

Searches with no results

This is the number of searches that supplied no results. In other words, nothing in your knowledge base or community forum contained the search keywords.

While this might indicate you need to create content covering those searches, it also can mean you need to expand your keyword terms.

Searches with no clicks

This metric typically indicates your articles aren't needed or, more likely, that their titles need to be revised for greater context and clarity.

Tickets created after search

This is probably the most important of these three search metrics because it strongly suggests that the provided content wasn't enough to help the customer solve their problem independently. This could also hint at a problem with the product if customers are finding it difficult to use even with sufficient self-service content.

Self-service score

While the metrics above provide insight into the performance and quality of your help center content, a self-service score (also known as a self-service ratio) measures the effectiveness of your help center as a self-service channel: whether it's helping customers solve their problems and preventing them from submitting support requests.

Determine your self-service score using this formula:

$$\text{Self-service score} = \frac{\text{Total users of your help center(s)}}{\text{Total tickets}}$$

Key takeaways

1. Many customers prefer self-service. Providing a self-service option improves customer satisfaction and helps you scale your organization.

2. Tracking your metrics can tell you how effective your self-service channels are at deflecting tickets and indicate areas for improvement.

3. Measuring your self-service score helps you evaluate your self-service portal's impact on ticket volume.

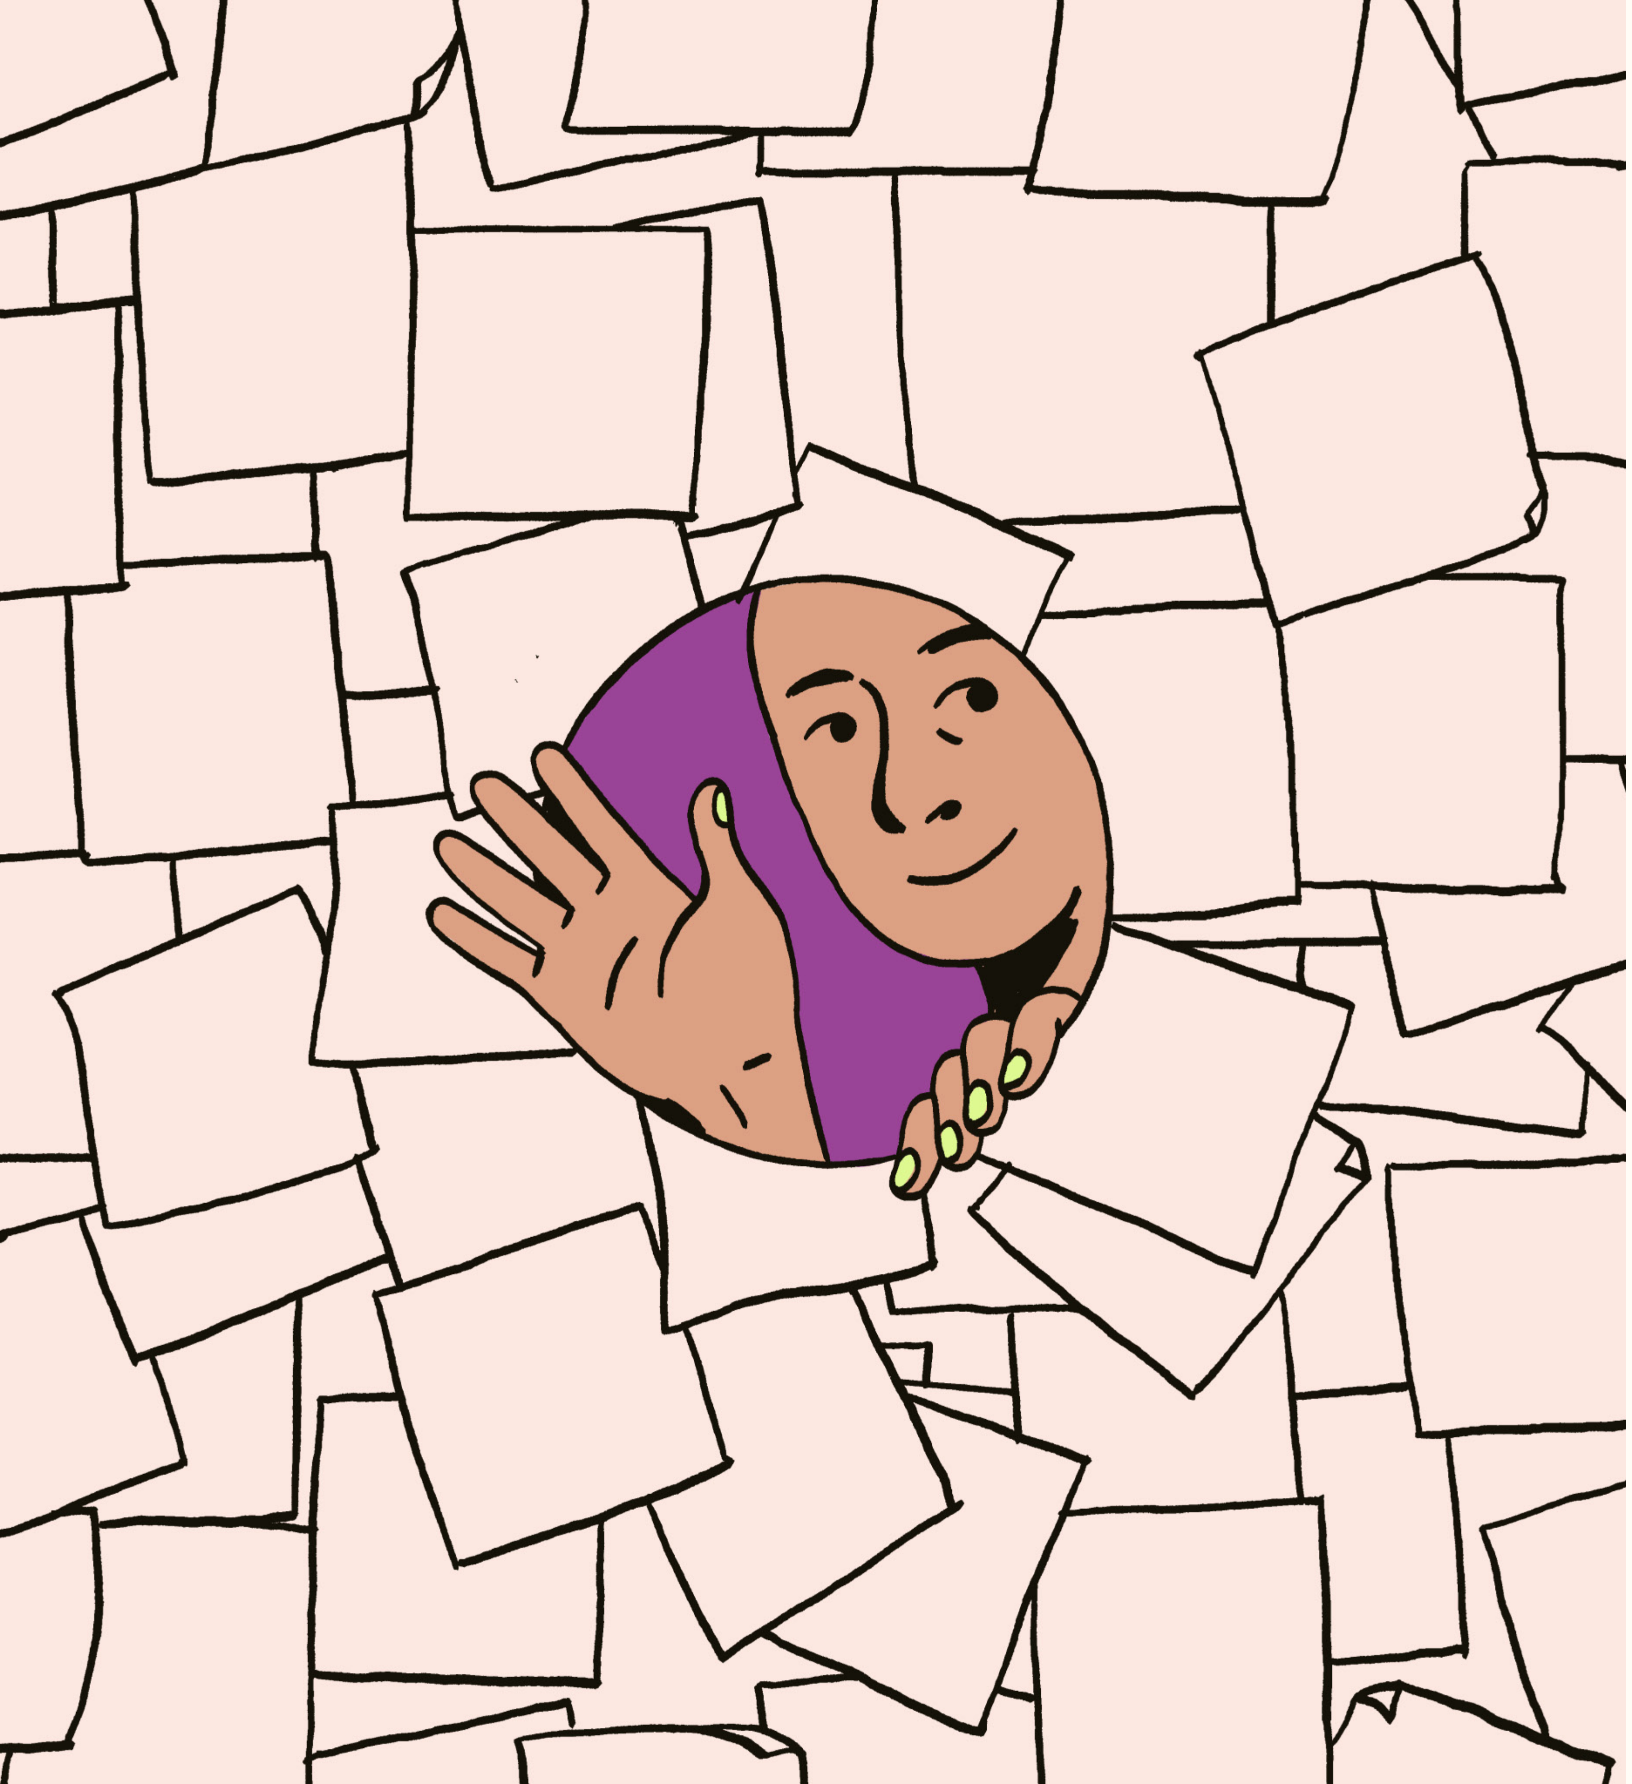

Track, measure, and improve your customer service metrics with Zendesk

While keeping track of customer service metrics can be overwhelming, doing so helps businesses deliver quality customer service and identify how to operate more efficiently.

At Zendesk, we offer several features that can help you track, measure, and improve your customer service metrics—so you can enhance your customer experience.

View a demo of Zendesk below to learn more about our reporting capabilities.

View demo

zendesk