Customer service metrics that matter

Metrics help us track how well we perform in our primary customer service mission— achieving and maintaining high customer satisfaction.

Customer service metrics help us track how well we perform in our primary customer service mission—achieving and maintaining high customer satisfaction.

To achieve this mission, we need data to understand our operations and the customer experience. Fortunately, abundant data is available these days.

All the metrics for measuring this data are valuable, but which ones do we rely on to effectively measure our success and guarantee that we're making smart business decisions? Which metrics make the difference when it comes to ensuring customer satisfaction? We have created this guide to help your business find those answers.

Try thinking about it like this: if we start a journey with a specific destination in mind, we'll want to plan our trip accordingly so that we can make the most out of it (and not get lost along the way). Achieving and maintaining high customer satisfaction means keeping your eyes open to determine which customer service metrics really matter and how to utilise them effectively.

It is like taking a road trip to somewhere new: you might have an idea of what to expect, but there will probably be a few surprises along the way. As long as you know your path and keep looking forward, the end goal will be within reach.

01   Customer relationship metrics

Because our primary goal is customer satisfaction (which in turn leads to healthy customer relationships), the metrics used to measure customer satisfaction are the best place to start. But what makes customers satisfied and happy? That's actually an easy question to answer. They don't want anything to go wrong ever. If something does go wrong, or if they need help using products or services, they don't want asking for help to be a time-consuming hassle.

As demonstrated by the authors of the Effortless Experience, customers want an easy, fast, low-effort support experience.

To gauge your achievement, measure both short and long-term customer satisfaction and happiness in the following areas:

- Look at satisfaction scores for support interactions.
- Measure the effort customers invest when contacting you.
- Survey customers about their overall experience with your company.

In this section, we'll explore these customer relationship metrics:

- Customer Satisfaction Ratings (CSAT)
- Customer Effort Score (CES)
- Net Promoter Score (NPS)
- Social media and churn metrics

Customer Satisfaction Survey (CSAT)

For each customer interaction, you can and should measure their satisfaction with the support they received. Do so using the Customer Satisfaction Survey (CSAT), a short survey sent to the customer after their problem is solved.

At Zendesk, we take a very simple and effective approach: ask only if the interaction was good or bad.

This is a quick, no-hassle way for customers to provide feedback. They also have the option to add a comment if they want to give you more feedback.

After you've gathered these satisfaction ratings, there's a lot you can do with this data. In fact, it is a metric you can look at from many angles. For example, you can track the following:

- 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

How would you rate the support you received?

Good, I'm satisfied

Bad, I'm unsatisfied

Tracking CSAT in these ways helps you spot trends that may be affecting customer satisfaction. Recently launched a new channel? Track CSAT for that channel to see how your customers like it.

Satisfaction ratings are also key performance indicators for your agents because each rating received is associated with the agent who solved the customer's problem. These are averaged to give each agent an overall CSAT score.

Because CSAT ratings are so important, it is a good idea to set CSAT ratings targets for your agents and teams and then track their performance against those targets.

We'll talk more about the use of CSAT ratings for agents and teams in Measuring agent satisfaction and performance.

Customer Satisfaction

Chat CSAT   Voice CSAT   Social media CSAT

Example of average customer satisfaction over time, by channels (sample data)

To what extent do you agree or disagree with the following statement:

The company made it easy for me to handle my issue:

☐ Strongly Disagree
☐ Disagree
☐ Somewhat Disagree
☐ Neither Agree nor Disagree
☐ Somewhat Agree
☐ Agree
☐ Strongly Agree

An example CES survey

Customer Effort Score (CES)

Some customer service industry experts, such as the CEB (the organisation that brought us the Effortless Experience), believe that asking customers about the effort they made to resolve their issue may be a better predictor of customer loyalty than the CSAT survey. That's why the customer effort survey, called the Customer Effort Score (CES), was developed.

Like CSAT, it is transaction-based—sent to the customer after their problem has been solved.

Companies that use the CES typically use it instead of CSAT ratings. Some move to CES after having used CSAT, feeling that they gathered all the useful feedback they could from that survey.

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, as a result, overall customer satisfaction.

Net Promoter Score (NPS)

To get to the bigger customer relationship story, beyond single support interactions, we use the Net Promoter Score (NPS). The NPS survey helps us understand if the customer is likely to return, stay loyal, and advocate for our brand to other customers.

Unlike CSAT and CES, it is not a transactional measurement of the support we provided; rather, it is a long-term measurement of their experience with our company.

The Net Promoter Score asks how likely they are to recommend our business to someone else.

Based on the 0–10 rating, customers are divided into three groups:

- Detractors (0–6)
- Passives (7–8)
- Promoters (9–10)

Detractors are dissatisfied 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 they probably are not sharing that with anyone else.

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

The NPS is calculated by subtracting the percentage of Detractors from the percentage of Promoters. (Passives are not counted.) If everyone who responded were Detractors, your score would be -100. If all of them were Promoters, it would be +100. But of course, neither of those outcomes are likely. The 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 customers to provide you with NPS survey feedback. NPS surveys are often managed by Customer Success and Marketing teams, not Customer Service.

In Zendesk Support, you can create NPS surveys and monitor the results in the Reporting dashboard.

How likely are you to recommend us to someone you know?

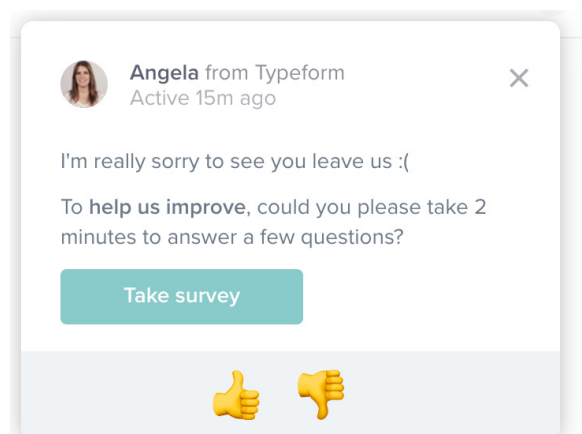| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |

Not at all likely   Extremely likely

Are your customers Detractors, Passives, or Promoters?

Because NPS is about overall, long-term satisfaction and not about the latest transaction, companies generally send NPS surveys at a regular interval—such as quarterly or every six months. The survey goes to a random sample of the overall customer base, not just to those who recently interacted with the company.

Your score

75

0%	25%	75%	66%
Detractors	Passives	Promoters	Response Rate
(0 users)	(1 user)	(3 users)	(4 / 6 users)

NPS Survey results (sample data)

When customers leave, ask them why

Customer Success teams want customers to be successful, to stay and renew their subscriptions. So, they really wish to understand the reasons that lead a customer to click on the dreaded "cancel subscription" button. By surveying customers as they are in the process of cancelling their accounts (or right after they've cancelled it), you can collect that vital feedback.

Social media and churn metrics

Other metrics to watch that reflect on your customers' experience include social media mentions and reasons why your customers no longer want to do business with you.

Social media listening

Track both positive and negative social media mentions to help you understand what's being said about you publicly. This includes Twitter, Facebook, your Instagram account and product review sites.

Using social media-monitoring tools, you can easily collect and analyse customer feedback. Use this feedback to determine the following:

- 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 feedback, positive or negative?

- 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 your 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 prompt your customers to tell you why they're cancelling their account. From this data, you can create a report of your churn activity.

This can be done within the user interface during the cancellation process or as a follow-up email request after they've cancelled. Your response rate will be much higher if you embed this survey into the user interface.

Key Takeaways

01   Discover if your customers are happy and will remain loyal by asking them if they're satisfied after you've solved their problem, by understanding the effort they put into it themselves, and by asking if they would recommend your company to someone else.

02   Setting customer satisfaction ratings targets for your agents and teams keeps the focus on the customer experience.

03   Monitor all your feedback channels. These include social media and the valuable feedback you get when your customers decide to stop doing business with you.

02 Essential agent and team performance and efficiency metrics

Measuring both individual agent and team performance is essential for keeping on top of the day-to-day running of your customer service organisation.

In this section, we look at the productivity metrics used to track that performance. Solving customer problems quickly and effectively is the goal, so these metrics are all about managing support issues to resolution.

You can monitor these metrics in the Zendesk Support reporting tools and in Zendesk Explore.
In this section, we look at these metrics:

- Tickets solved

- Issues, by resolution area

- First-reply time

- Average reply time

- First-contact resolution

- Time to resolution

- Resolution effort metrics

- Ticket reopenings

- Next issue avoidance

Tickets solved

Many customer service teams set a daily ticket solved target for their agents (15, for example).

Based on this number, they track how well their agents and teams are performing using the percentage of the target achieved. Here's an example of how two different teams are doing hitting their daily ticket solved target.

It is useful to look at this by team and by individual agent. Of course, some days, an agent will be under (or over) the target of 15 tickets solved per day. Tracking the daily average over time is most important (in this example, 30 days) for spotting trends in agent performance and resourcing needs.

This metric is calculated by dividing the total number of tickets solved in the last 30 days by an average of 18 working days a month.
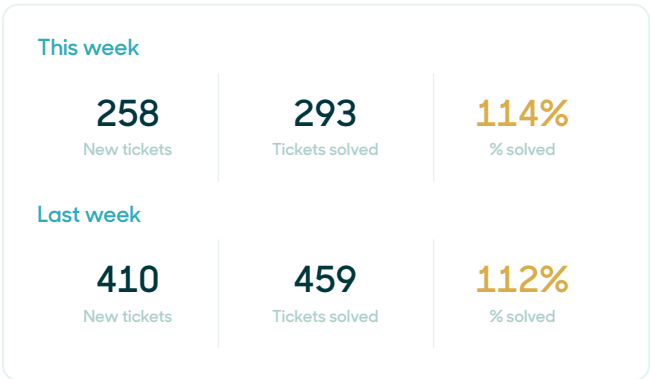
Tickets solved vs. open tickets

Another way to use the number of solved tickets is to compare tickets solved to how many open tickets you have. This is helpful in understanding how well the team is keeping up with the ticket queue.

We'll take a more in-depth look at the ticket queue and backlog in Ticket Backlog.

Hitting
112.2%
of 15 ticket/agent target

Hitting
110.1%
of 15 ticket/agent target

Are your agents and teams hitting their ticket solve targets? (sample data)

This week

258	293	114%
New tickets	Tickets solved	% solved

Last week

410	459	112%
New tickets	Tickets solved	% solved

Comparing new tickets to tickets solved (sample data)

Issues by resolution area

Customer feedback in the form of support requests may be the most valuable information you have available for improving your products and customer experience. Therefore, you should take advantage of this and categorise your support requests into the product areas generating support issues (e.g., Account Creation, Shopping Basket, and so on).

With that information, you can segment support requests by these categories and drill in deeper to discover, for example, how many tickets are being generated, how long they take to solve on average, and the average CSAT ratings for each area.

At Zendesk, we collect this data using a custom ticket field we call the About field.

Ask agents to categorise tickets using the About field by selecting from the list of predefined product categories.

You can also add the About field to your support request form so your customers can select the relevant product areas when they submit a support request. They're happy to add this data for you in the hope that it will expedite their request.

You may also wish to compare data from the About field to overall ticket volume and resolution time to identify areas of the product that should be improved. Share that data with your Product development teams.

Reply-time metrics

How quickly you reply to your customers has a significant effect on customer satisfaction. For this reason, it is important to focus on reply-time metrics.

If you offer service level agreements (SLAs) to your customers, this is an essential metric to track.

First-reply time (FRT)

First-reply time (FRT) is the amount of time from when a ticket is created to when an agent makes the first reply to the customer.

In other words, FRT measures how long it takes a real human, not an automated reply, to contact the customer. FRT is also often referred to as *first-response time*.

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

- 24 hours for support requests submitted via email and online forms

- 60 minutes for social media

- A few seconds for Chat and Messaging

Of course, performing faster than these expectations is even better from the customer's perspective.

Monitor FRT by support channels to ensure that those expectations are being met. You should also measure FRT by product areas, if you use an About field, as a measure of team performance.

FRT is a good indicator of the efficiency of your support processes and how well they hold up in handling fluctuating incoming ticket volumes.

Average reply time

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

How do you determine what an acceptable average reply time is? Consider the following:

- If your FRT is 24 hours on email-submitted requests and your median number of agent touches is 8 per ticket, then your average reply time median should be 4 hours or less. Looking at tickets with an average reply time greater than 4 hours should highlight opportunities for improving agent performance (such as using better docs or tools or more training).

- What are your customers telling you about your reply time? Look at your CSAT ratings and comments.

In Zendesk Explore, you can calculate average reply time by dividing requester waiting time by the number of agent comments.

Number of replies

Number of replies is the number of times an agent replied to the customer.

This is an important metric because it is a good indicator of customer satisfaction; it drops significantly after three interactions with the same customer. It may also indicate how well the agent is performing. In Zendesk Explore, you capture this metric using *# Public Comments*.
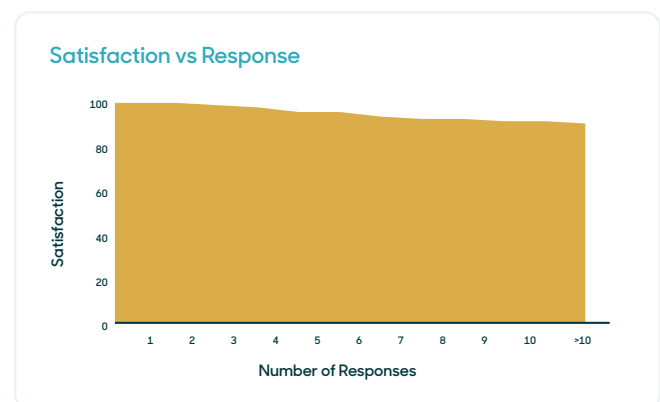
The data says it all: fewer replies equals higher CSAT (sample data)

Resolution time metrics

Resolution time is important to customers, tool; they like a quick resolution as well as a quick reply. We have several different metrics to help us look at how long it takes for agents to resolve issues.

First-contact resolution (FCR)
First-contact resolution (FCR) measures the percentage of support issues that were resolved in a single interaction.

This metric is also known as *first-call resolution* or *one-touch resolution*.

A single interaction is defined as one live-chat session, one telephone call with an agent, or an email enquiry that was solved with a single reply.

FCR has a huge impact on customer satisfaction, helps reduce operating costs, and improves agent job satisfaction.

If you have a high percentage of FCR tickets, however, this may indicate that you have lots of low-complexity support issues that could be better handled with self-service. Agents should work on the more complex issues, not those that a customer can easily solve on their own by reading an FAQ.

Analysing the data you get from using an About field can help you spot areas of the product that need self-service content.

You can read more about this important metric in First-contact resolution: beacon of good support?

Time to Resolution
Time to resolution measures the time it takes for a support issue to be solved.

This metric is also called *resolution time* and *mean time to resolution/resolve (MTTR)*.

A support issue may have been solved more than once before being closed—because a ticket can be reopened before it is closed. So, there can be multiple resolution times in the life cycle of a ticket.

The time it took to solve the issue the first time is called *first-resolution time*. This will not be a *first-contact resolution* if there were multiple interactions before it was solved.

When it is finally solved (the most recently solved) and closed (not reopened), that is the *full resolution time*.

This is another metric that has a direct effect on customer satisfaction: people don't want to wait a long time to have their problems resolved. However, focusing solely on resolution time is not necessarily a good idea. Speed does not always equal quality. We'll discuss that later in this guide.

To better understand the customer experience of resolution time (and agent effort), we also look at a metric called *requester waiting time,* which we'll discuss next.

An oft-cited study by the Service Quality Measurement (SQM) Group boasts that a 1% improvement in FCR can result in a 1% improvement in CSAT, a 1% decrease in contact centre operational costs, and a 1–5% increase in employee satisfaction.

Sorting out resolution time metrics

Resolution time metrics

In addition to understanding how long it took to resolve a support issue, going deeper into the details about how much effort went into that resolution is also important.

Doing so reveals the actual time that the agent spent resolving the issue. Total resolution time does not do that because it includes stages in the ticket life cycle when the agent is not working on it: for example, when it is sent back to the customer for information and marked as Pending.

Overall effort is described using several different names, a variation of *handling time* being the most common. Whichever way it is referred to, with this metric you can also determine the cost of each support interaction.

Handling time
Handling time is the time an agent spends working on a single support interaction.

This metric is also referred to as *average handling time (AHT), average call handle(ing) time (ACHT)* and *contact handling time.*

In Zendesk Support, agent handling time is captured using the Time Tracking app.

Time Tracking

00:05:04

Total time: 00:05:04

The Time Tracking app captures agent handling time

After you've captured this data using the Time Tracking app, you can report on it in Zendesk Explore using the custom metrics *Total time spent, Avg time spent per ticket,* and *Avg time spent per update.* Learn more about these custom metrics here.

Agent touches & ticket comments
An agent touch is an update that an agent makes to a ticket. This includes changes of ticket status and other updates that customers may not be aware of.

Agent touches is the terminology used in the Zendesk reporting dashboard. When using Zendesk Explore, you can report on "touches" using the *# Ticket Updates* metric.

To more exactly measure the customer's experience, we also look at the total number of public comments an agent adds to a ticket using the *# Public Comments* metric.

A high number of agent touches or comments may indicate that the agent needs more training because it takes them more than the team average (for example) to solve tickets.

Requester waiting time
Zendesk Explore provides a metric that captures how long the customer waits for their support issue to be resolved.

Requester waiting time measures how long a ticket is set to New, Open and On-hold.

In other words, the times when it is the agent's and support team's responsibility to work out a resolution for the support issue.

When a ticket is set to Pending or Solved, it is sent back to the customer to take the next step. Pending is used when an agent needs more information from the customer. When a ticket is set to Solved, it is sent back to the customer to close the loop. For example, a response to the question: Did the support we provided solve your issue?

Unlike *resolution time,* this metric does not include all the stages in the ticket life cycle—only those stages when the customer is waiting to hear back from an agent (New, Open and On-hold). Therefore, it is a measure of effort and an indicator of the customer's experience of the support interaction.
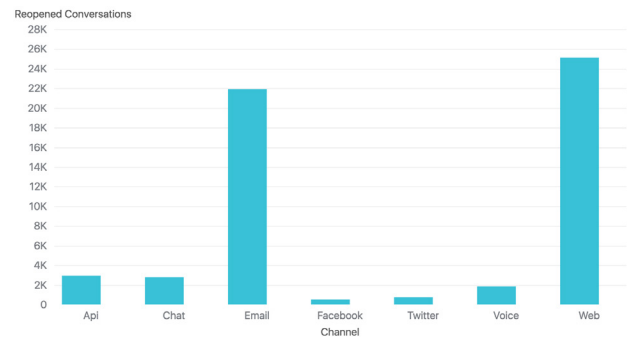
Ticket reopenings

How often solved tickets are reopened may indicate that agents are not fully solving the customer's support issues. This may be due to a focus on first-contact resolutions and speed over quality. Therefore, it is a best practice to routinely monitor your team's ticket reopenings.

A reopening occurs when ticket status is changed from Solved back to Open.

Ticket reopening reports should include:

- Total number of reopenings

- The average number of reopenings

- The percentage of tickets with reopenings

Ticket reopenings are more likely when agents deal with more complex support issues; therefore, reopenings may be higher for escalated tickets.

Reporting on ticket reopenings (sample data)

Multiple Support Requests

Fodrizzle Umbrellas	● Channels::Chat
Cabber	● Channels::Chat::Zopim API
HyperX	● Product::Embeddables::Web Widget
Sunshine Industries	● Channels::Chat::Other
Helios	● Channels::Voice
International Genetics	● Product::Help Center (HC)::HC Customisation

Number of Tickets

3 or more support requests, by product/subject area (sample data)

Next issue avoidance

Another argument against focusing on resolution speed and agent handling time is that it often results in more support requests from the same customer.

Your focus should instead be on engaging with customers when they do contact you to ensure that you're solving their problems the first time. This may result in fewer solves per day, but it will increase customer satisfaction.

The next-issue avoidance metric looks at how many customers have requested multiple support requests for the same product/subject area.

For example, you can create a report that tells you how many customers submitted 3 or more support requests for the same product/subject area.

Compare this metric to your first contact resolution results to see if there's a correlation. If both your FCR

and next-issue avoidance rates are high, work to improve that first contact with the customer.

You may also gain a better understanding of why you're getting multiple support requests for the same product/subject area by listening to your community. If you've set up a Help Centre and Community, routinely analyse the top 10 community posts to see what your customers are asking and talking about.

Key Takeaways

01

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

02

Quick first-reply times and first-contact resolutions lead to customer satisfaction, but make sure these metrics are not masking another problem such as ticket reopenings. Focus on quality over speed and efficiency.

03

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

03 Putting your metrics to work

Most of the metrics we have looked at are easy to access and analyse in the admin dashboards and reports provided by Zendesk Support.

If you wish to dig even deeper and create custom views of these metrics, you can use the advanced reporting and dashboards provided by Zendesk Explore.

But tools for slicing up and analysing data are just part of the job. Your approach to customer service determines what metrics are important to you—how you view your organisation's performance and your customer relationships. What are your goals and priorities?

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

- 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

Balancing quantity and quality

We track metrics to establish a clear understanding of our operational efficiency, team and agent performance, as well as how our customers feel about the job we're doing and how likely they are to continue doing business with us.

To do that well, we need to find a balance between the metrics that give us the quantitative data which helps us understand our organisation's throughput (tickets solved per day, resolution time, average handling time) and the metrics that reveal the more qualitative aspects of providing customer service (CSAT, customer effort and next-issue avoidance).

If you align your operational baseline with your customers' needs, their satisfaction, and the long-term relationship in mind, the appropriate metrics come more sharply into focus.

Blinded by fast and first

As we mentioned earlier in this guide: First-reply time improves customer satisfaction, and improving first-contact resolution improves customer and employee satisfaction and reduces operational costs.

The challenge is to make sure you're not blinded by what appear to be great performance stats. Fast and first can also reveal shortcomings and opportunities for delivering better customer service.

Getting it right the first time

A high FCR rate is great if ticket reopenings and next-issue avoidance rates are also low. The first-contact resolution needs to stick so that customers don't have to follow up with more support requests. Don't prioritise speed over quality.

When customers reach out to you for support, make sure that agents take the time needed to solve the issue. And while engaged with them, help them better understand how to use your products, avoid follow-up tickets, and improve your next-issue avoidance numbers.

Is self-service the better option?

A high percentage of issues solved as first-contact resolutions may indicate that you have lots of support issues that could be better handled by self-service (because they're low-complexity). Providing self-service helps reduce costs and saves your agents' bandwidth for the more complex issues.

Put your metrics to work to help you figure out where self-service will have the biggest impact. For example, see how many tickets are created for each product area you've defined using the About field, and compare that to the average handling time that goes into solving issues in those areas.

Beyond self-service, that information may indicate that certain areas need dedicated support expertise, more training to help agents through the issues or, more likely, changes to the product to make it easier to use.

Feedback effort without the ask

By breaking down the support you provide into product areas, you can also do a lot to reduce customer effort without having to even ask (via a CES survey, for example).

In addition to looking at the number of tickets by product area, you can also measure the average handling time, the number of agent updates and comments, full resolution time, the CSAT ratings, and direct feedback from customers in CSAT comments and churn surveys.

Setting operational baselines

To measure your success, establish baselines for your key performance metrics. The best way to start 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 one of the ways that you measure performance. With this, you can report on hitting target percentages by individual agents, for teams and as an organisation. This should be one of, rather than the only, performance factor to consider.

To determine a reasonable tickets-solved-per-day 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. Pick the middle number between each, and validate that across other agent solved rates per day. This baseline should challenge your teams without being unrealistic.

Customer satisfaction ratings

Set an overall CSAT goal for your organisation and CSAT ratings goals for your agents (e.g., 90-96% based on experience, job level and areas of expertise).

If your teams are doing well hitting their CSAT targets, you might also want to set 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, it is best to start by looking at your historical performance. What is your current FRT? What was it last year? Is there a trend up or down? From that, select an achievable performance goal that reflects your team's historical performance and also aligns with customer expectations. If you hope to drive your FRT down, set a stretch goal as well.

First-contact resolution

To set an FCR baseline, take the same approach as for FRT. Look at historical trends to establish an achievable goal. Also, keep in mind the FCR challenges that we discussed in Balancing quantity and quality.

Agent touches per ticket

As we saw earlier in this guide, CSAT declines with the number of ticket replies. Therefore, this is an area where you should also set performance measurements.

Start by looking at the median number of replies across your organisation, and establish that as the baseline. Then, monitor the tickets above the median to look for opportunities to improve performance, such as by providing more training and coaching for your agents and by creating more documentation or support tools.

If your median reply time is not syncing with customer expectations and is driving down CSAT ratings, set a more aggressive baseline and manage your staff to that goal.

Tickets solved in Tier 1

In a tiered customer service organisation, the rule of thumb is that the Tier 1 team solves 70-80% of support issues.

The other 20-30% of the support issues will be more complex and costly because they require more back and forth between agents and customers or agents with other agents, a higher level of technical skill, and generally more time troubleshooting.

Cost per contact

This is a measure of the cost of the effort that an agent puts into solving a customer's support issue.

A formula that may help you to work out your cost per contact is to divide your organisation's annual budget by the total number of requests for help across all products and channels that you support. You can then break it down by team as a way of setting more discrete costs.

As your organisation grows and provides more blended services, this will of course require more nuancing to maintain accuracy. Knowing your cost per contact, however, is an important part of demonstrating to your budget providers that they are getting a well-managed return on their investment.

Support channels

Many support channels provide metrics that are specific to those forms of communication. For example, when using a phone support channel, set baselines for metrics such as abandonment rate, average waiting time, average holding time, number of transfers, and so on. You'll find more information about these metrics in Getting Started with Zendesk Talk.

The uses of average and median

Before we get into how we look at these metrics as reports and in dashboards, let's talk about *average* (or mean) and *median*. In many of our reports, we look at both simultaneously because they're both revealing.

Averages are a useful middle point in a set of numbers, but they are easily skewed by numbers that are much higher or lower than the rest. That's why the median is a better midpoint measure for cases where a small number of outliers could drastically skew the average.

Here's an example to illustrate that point.

How a few outliers can skew your averages (sample data)
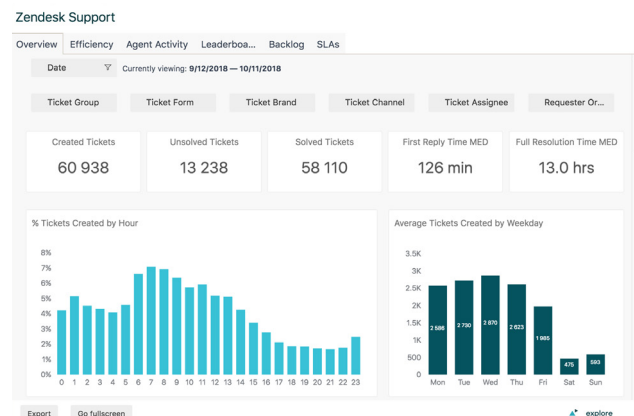
Here is our average vs. median rule of thumb:

- If the data you're comparing is mostly uniform, use *average*

- If your number set has some outliers, use *median* (or filter out values that otherwise skew the results).

With your available reporting tools, you can look at both.

Average is useful when you look at metrics that have few outliers, such as phone and chat first-reply times. It is also a good way to measure general ticket volume.

Tracking metrics in Zendesk

The Zendesk suite of customer service tools provides lots of options for tracking metrics and building custom reports and dashboards, which enable you to look at your data from many different angles.

The Knowledge Base, Community and Search dashboards provide you with a view into the performance of the self-service you provide in your Help Centre. We look at these in more detail in Measuring your self-service channel.

The Reporting overview in Zendesk Support (sample data)

In Zendesk Support, essential customer service metrics are automatically calculated and presented to you in the admin Reporting dashboard.
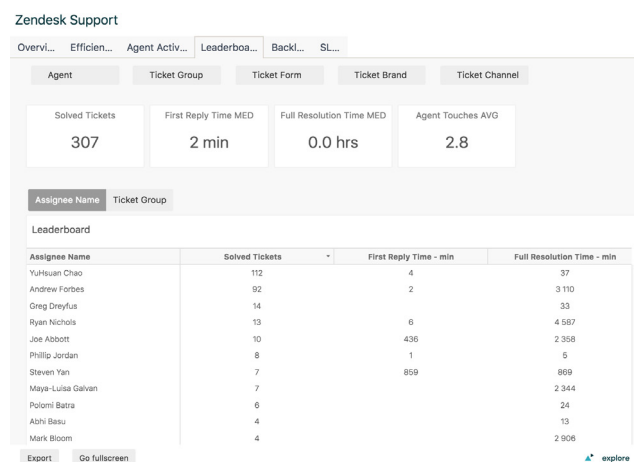
The Leaderboard in Zendesk Support

Reporting in Zendesk also includes dashboards for Customer Satisfaction and Net Promoter Score (two essential measures of the health of our customer relationships).

Out of the box reporting to understand the impact of AI

In Zendesk Explore, you can use an expanded set of metrics to deep-dive into agent and team performance, ticket queue status, and the customer experience.

In Zendesk Explore, you can easily create highly visual dashboards by connecting to your Zendesk Support data.

Zendesk Explore is the new analytics platform from Zendesk. It provides analytics for all Zendesk products—Support, Talk, Chat and Guide—and enables you to analyse and report on Zendesk data holistically, as well as integrate customer data from sources like your CRM, website or data warehouse.

With Zendesk Explore, you can measure and understand the entire customer experience. Find out more about Zendesk Explore.

A few examples: key metrics by support channels

Here are a few examples of how you can easily monitor your key customer service performance metrics using Zendesk Explore.
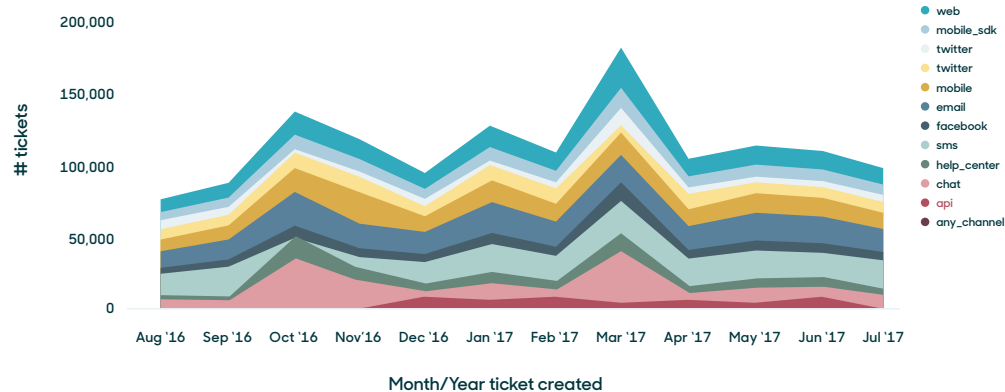
↩ **Median first reply time**

1.2 hrs

Measuring median first reply time (sample data)

⭐ **Satisfaction score**

91.1%

Customer satisfaction ratings, by channels (sample data)

Tickets created by channel

Keeping track of ticket volume, by channels (sample data)

Key Takeaways

01

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

02

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

03

Setting operational baselines helps you establish realistic performance goals for your teams and agents.

04

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

04   Ticket backlog

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, given your resources.

As we have pointed out in this guide, 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.

The longer it takes to solve a customer's issue, however, the more likely customer satisfaction will suffer. In this section, we look at ways to monitor your ticket backlog.

Reporting on your backlog

Backlog tracking is an essential metric in Zendesk and is included in the Reporting overview and as a collection of metrics in the advanced reporting provided by Zendesk Explore.

Here, we look at some of the ways you might wish to keep track of your backlog to understand its impact on your team and customers, as well as why it happens in the first place.

Your current backlog

A backlog consists of both unassigned tickets and assigned tickets that have not yet been solved. The two combined equal your total ticket backlog.

It is also a good idea to cross-reference backlog volume with ticket age and first-reply time. A large backlog is not necessarily a bad thing if your team is capable of a high throughput.
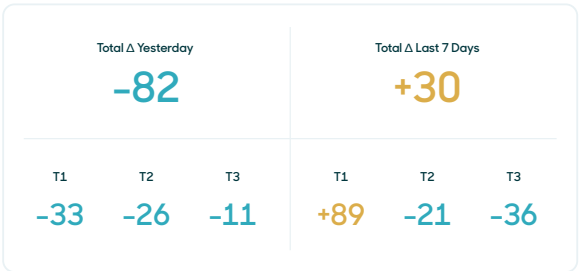
A large backlog can be bad, but one with a large number of high-priority tickets is much worse.

Current unassigned	T1	T2	T3
140	100	30	10

Unassigned tickets, by teams (sample data)

Teams:
- T1
- T2
- T3

Total backlog by teams (sample data)

Total Δ Yesterday			Total Δ Last 7 Days		
−82			**+30**		
T1	T2	T3	T1	T2	T3
−33	−26	−11	+89	−21	−36

Backlog trends by team (sample data)

Ticket status:
- New (not visible)
- Open
- Pending
- Hold

Total backlog by ticket status (sample data)
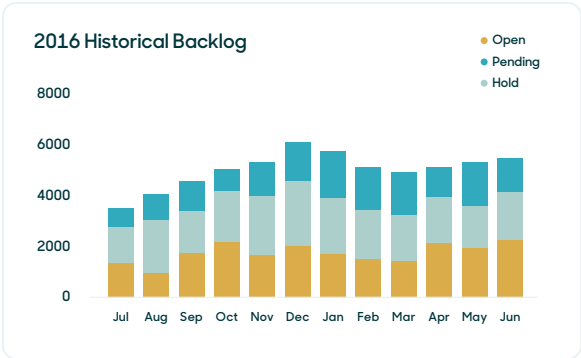
Historical backlog by ticket status - one year (sample data)

Historical backlog by ticket status - last 30 days (sample data)

Historical backlog

Looking at the historical backlog, as well as the current backlog, helps you spot trends in ticket volume and team performance.

Managing your backlog

Support requests stuck in the backlog means that customers are waiting longer to get help. To manage that, it is best to keep the customer informed about the status and progress of their request.

Your backlog numbers may be driven up by tickets that are in limbo—for example, if the customer has not replied to a request for more information or to confirm that a proposed resolution worked. In the chart above, these are the Pending tickets.

To handle this, you can create an automation to remind the customer of the pending status of their issue. If they don't respond after a set number of days you can automatically resolve the ticket to get it out of the backlog and queue.

Key Takeaways

01

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

02

Your total ticket backlog is both unassigned and assigned unsolved tickets combined.

03

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.

05 Measuring agent satisfaction and performance

Agents who are happy and satisfied in their roles deliver better customer service. They're more engaged, motivated and productive. They also stay on the team longer. As you know, managing turnover in a customer service organisation is always a challenge. By regularly measuring job satisfaction, you get the feedback you need to create a better job experience for everyone on the team.

The insight you gain from understanding how satisfied your agents are, combined with their day-to-day performance stats, gives you the big picture you need to understand how well they're doing today and where they may be headed in the future.

In this section, we look at how to measure agent satisfaction, which we call ASAT, and the metrics we typically use to measure 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 looks like this:

☐ **How much do you like or dislike your current job?**

- Like a great deal

- Like a moderate amount

- Like a little

- Neither like nor dislike

- Dislike a little

- Dislike a moderate amount

- Dislike a great deal

☐ **What do you like most?**

☐ **What do you dislike most?**

The first question is multiple choice, which can be evaluated as numerical data (the same numeric scale as NPS), and the two follow-up questions allow employees to provide more information about what they like and dislike the most about their jobs. The feedback in the follow-up questions helps you act on the survey results.

You might also include these two questions:

- How much do you like or dislike working as a customer service professional?

- How likely are you to recommend a job on your team to someone you know?

Here's the process for conducting ASAT surveys:

- Survey agents, leads and managers

- Survey everyone, every quarter

- Have someone outside the support team conduct the survey and evaluate and summarise the data

- Act on the survey feedback

Evaluating agent performance

Many of the key metrics we use to evaluate the customer service organisation's performance apply to individual as well. Here are the metrics we've found useful for evaluating their performance.

• Number of tickets solved per day

• CSAT ratings

• Agent touches/updates

• Ticket volumes by channel

• Number of escalations

• Areas of focus (About field)

You will probably have set a baseline for most of these, so they can be tracked easily over time.

First-reply time is not on this list because it's often outside 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 out of an agent's control (e.g., a customer's frustration that a product does not include a feature they want, resulting in a bad CSAT rating).

For a deeper level of detail about why customers choose the ratings they do, Zendesk provides a feature called Satisfaction Reasons that helps break it down. When a customer leaves a bad CSAT rating, you can prompt them to provide a reason why.

Using the comments provided in bad CSAT ratings and the reasons provided, you can more fairly assess an agent's CSAT ratings performance.

Key Takeaways

01

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

02

As you do with your customers, routinely survey your agents to get their feedback. Find out if they like their jobs and careers, as well as if they'd recommend to a friend to join the team.

03

Set realistic goals for your agents based on your key performance metrics.

04

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

06 Measuring your self–service channel

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

A self-service channel is what you provide in your Help Centre, as well as all other forms of learning and reference content you provide to your customers.

In this section, we look at the metrics that help determine if customers can find and use your self-service content, if they find the content useful, and if your Help Centre serves as an effective self-service portal and helps deflect tickets.

We categorise self-service metrics into the following areas:

- Views and engagement metrics

- Knowledge base and community metrics

- Search metrics

- Self-service score

Views and engagement metrics

These are the typical metrics you track for Help Centre performance (or for any website) using both Google Analytics and the Knowledge Base reporting available in Zendesk Support.

Views

This is the number of page views in your Help Centre. You can track views in both Google Analytics and the Knowledge Base reporting tab in Zendesk Support.

Users

This is the number of unique users who visit your Help Centre. Each visit to your Help Centre counts as a session, and each session (usually) results in multiple page views. Tracking the number of users visiting your Help Centre gives you some perspective about its use, compared to the total number of views in a specified period. A total of 10,000 monthly views compared to 1,000 unique users within that same period tells you that those users are viewing on average 10 pages per session. This helps you understand how many of your customers use your self-service content.

% New Sessions

Understanding how many new versus returning users visit your Help Centre helps you focus on the content that addresses those users' needs. For example, rolling out a new product may result in a spike in new users, which you can address by providing the information needed to use the new product.

Average session duration

The average duration of a user session in your Help Centre tells you how much time they spend in your Help Centre, as well as how much time they spend reading specific articles and FAQs, if you drill down deeper. Ideally, they spend enough time to read through the information you provided them. If they don't, that tells you something as well—that perhaps your content is not engaging or is not the information they need.

Pages per session

This is the average number of pages viewed during a session on your Help Centre. Once again, this 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 Centre. A bounce means that the customer left your Help Centre after viewing the first page that they landed on. A user may have visited the Help Centre unintentionally, or they may not have liked what they saw when they got there.

Knowledge base and community engagement metrics

Within the Knowledge Base and Community reporting tabs in Zendesk Support, you can drill into more metrics that capture how your Help Centre is being used. Here, you'll find the following engagement metrics.

Votes

This is the total number of Up and Down votes applied to your Help Centre articles. You can view either Up or Down votes or the total number combined. This is an excellent way to track how your customers feel about your content. You'll want to review the articles that received Down votes to see how to improve them.

Subscriptions

This is the number of customers who have chosen to follow articles in your Help Centre so that they can receive email updates when the content is updated or comments are added. It is a good measure of engagement, and many customers who have expertise in specific areas use this to engage your community of users and provide them with help.

Comments

Comments are added to articles when customers need more information or a clarification about the information you've provided. While this might indicate to some that the content is insufficient (because some customers need more information), it is actually a good way to enrich your knowledge base with more detail and unexpected use cases. It is also a good way to help customers before they submit a support request—deflecting support ticket creation.

Community posts

On the Community side of your Help Centre, monitoring the number of posts customers create is a great way to determine your self-service portal's effectiveness. Your Help Centre should be the go-to place for customers to find the information you provide them, as well as to engage with other customers, share their expertise, and learn from each other.

Search metrics

If your customers cannot find the information they're looking for in your Help Centre, you should improve the self-service you provide. Fortunately, we have metrics that help us track what customers search for and what actions they take after searching for answers.

Searches with no results

This is the number of searches that supplied no results. In other words, there was nothing in your knowledge base or community that contained the search keywords. While these might indicate that you need to create content that covers those searches, it also could mean that you need to spend some time making sure you and your customers are using the same terms.

Searches with no clicks

This metric generally indicates your articles are not needed—or more likely, that their titles need to be revised so that customers have a better understanding of their content.

Tickets created after search

This is probably the most important of the three metrics because it is a strong indication that the content provided was not enough to help the customer solve their problem on their own. This could also indicate a trouble spot with the product—that customers are having difficulty using it even with sufficient self-service content.

Self-service score

Whereas the metrics above provide insight into the performance and quality of your Help Centre content, the self-service score (also known as the self-service ratio) is an attempt to measure the impact that your Help Centre has as a support channel, how it is helping customers solve their problems and preventing them from opening support requests that then need to be handled by agents.

You determine your self-service score using this formula:

Self-service score =
Total users of your help centre(s) / Total users in tickets

This gives you a ratio such as 4:1, meaning that for every four customers who attempt to solve their own issues using self-service, one customer submits a support request.

You can read more about self-service and its role in ticket deflection in Ticket deflection: the currency of self-service.

Key Takeaways

01

Many customers prefer self-service. Providing it improves customer satisfaction and helps you scale your organisation.

02

In addition to the usual website and content engagement tools provided by Google Analytics, Zendesk provides powerful purpose-built metrics and tools for evaluating your self-service portal.

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

Conclusion

Sometimes, we can be overwhelmed with too much information and lose our way. This is especially true with all the data and metrics we now have available to measure how well we deliver customer service and build our customer relationships.

In this guide, we have defined a path through all that data and provided some expert advice about how to best use the metrics that we have found to be essential for managing our own support operations. We hope you have found it useful.

You can learn more about the tools we use to measure all of our essential metrics at www.zendesk.co.uk/explore.

zendesk