



Top 5 Support Analytics Tools in 2026

Table of contents

Top 5 Support Analytics Tools in 2026	1
Table of contents	2
Introduction	3
Zendesk QA (Klaus)	7
Totango	12
SentiSum	16
Cresta	20
Isara	24
Conclusion	29

Introduction

Why Support Analytics Matters in 2026

Customer support has evolved from a cost center into a critical driver of customer experience, retention, and long-term revenue. In many organisations, support is now the most frequent and most honest point of contact with customers, capturing real-world feedback about product usability, service quality, and unmet expectations. Research consistently shows that experience matters: 86% of customers say they are willing to pay more for a superior experience, raising the stakes for support and customer success teams to deliver fast, consistent, and empathetic service at scale.

At the same time, modern support operations generate an unprecedented volume of data. Every ticket, chat, call, and email contains signals about customer sentiment, friction points, and emerging risks. Yet the core challenge is no longer data collection. Most teams already have access to dashboards filled with metrics. The real challenge lies in separating meaningful signals from background noise and turning raw interaction data into insights that teams can actually act on.

Traditional support metrics and reporting tools tend to focus on what happened after the fact: ticket volumes, response times, resolution rates, or CSAT trends. While these indicators are useful for tracking operational performance, they often remain backward-looking and siloed within support teams. They rarely explain why customers are contacting support, what patterns are forming across conversations, or how today's issues might affect churn, expansion, or compliance risk tomorrow. As a result, critical insights often fail to reach customer success, product, or leadership teams in time to influence outcomes.

In 2026, a new generation of support analytics tools is emerging to close this gap. These platforms use AI and machine learning to analyse support conversations at scale, extracting patterns that would be impossible to spot manually. Rather than relying solely on surface-level metrics, they help teams identify early warning signals for churn, detect anomalies and emerging product issues, surface compliance risks, and understand shifts in customer sentiment as they happen. The goal is not just better reporting, but earlier, more proactive decision-making.

This guide explores five leading support analytics platforms through that lens. For each tool, we examine where it excels, where its limitations lie, and how it fits into a modern support and customer success stack. The aim is to help leaders understand which capabilities are essential for their organisation today, and which gaps may require complementary solutions as expectations around customer experience continue to rise.

How to Use This Guide

This ebook compares five support analytics platforms to help support and customer success leaders choose the right approach for their needs.

Key Evaluation Themes

- Operational visibility and routing (are we measuring and acting fast?)
- Predictive signals for churn and revenue (are we seeing risk early?)
- Conversation-level intelligence (are we capturing root causes and anomalies?)
- Quality and compliance (are we consistently on-brand and within policy?)
- Cross-team impact (do insights flow from support to success and product?)

In the following chapters, we examine each tool using these themes and include practical context on where each excels and where teams might pair solutions.

Zendesk QA (Klaus)

AI-Powered Quality Assurance with
Deep Zendesk Integration

Zendesk QA (formerly Klaus) is a quality assurance platform that has evolved from a focused conversation review tool into a centralised QA layer embedded within the Zendesk ecosystem. Originally positioned as a standalone solution for manual ticket reviews, Klaus gained traction by addressing a persistent pain point for support leaders: the lack of consistency, fairness, and scalability in agent quality assessments. Its acquisition by Zendesk in 2024 marked a strategic shift, transforming the product into Zendesk QA and positioning it as the default quality framework for teams already operating on Zendesk.

This integration-first positioning is one of Zendesk QA's defining characteristics. Because it is native to Zendesk, the platform fits naturally into existing workflows, permissions, and reporting structures. Support teams can adopt structured QA practices without introducing a new vendor, data pipeline, or operational layer. For organisations prioritising simplicity, governance, and tight platform consolidation, this native alignment significantly lowers the barrier to implementing quality assurance at scale.

At its core, Zendesk QA is designed to formalise how support quality is defined, measured, and improved. The platform allows teams to build detailed scorecards that reflect their internal standards, brand voice, and policy requirements. These scorecards can include both objective criteria, such as whether required steps were followed, and more subjective dimensions, such as empathy or clarity. Conversations are then sampled and reviewed against these criteria, creating a repeatable evaluation process that replaces ad-hoc reviews or informal feedback.

Over time, this structured approach helps organisations establish a shared language around quality. New agents are onboarded against clear expectations, experienced agents receive consistent feedback, and managers gain a defensible framework for performance discussions. For larger or distributed teams, where subjective assessments often lead to

mistrust or inconsistency, this standardisation can be a meaningful cultural improvement.

Zendesk QA's coaching workflows are tightly coupled to this evaluation model. Feedback is anchored directly to specific moments within a conversation, enabling agents to understand context rather than receiving abstract scores. Managers can track progress over time, identify recurring coaching themes, and measure whether targeted interventions lead to sustained improvement.

This longitudinal view of agent development makes the platform particularly useful for organisations investing in continuous enablement rather than one-off training sessions.

A major evolution in the platform has been the introduction of AutoQA. Historically, QA programmes relied on reviewing a small sample of conversations, often less than five percent of total volume. AutoQA addresses this limitation by applying AI models to evaluate every interaction against predefined quality criteria. This shift from sampling to full coverage dramatically increases visibility and reduces blind spots, especially in high-volume environments where issues can scale quickly.

AutoQA can surface conversations that deviate from expected quality patterns, detect sentiment deterioration, and flag missed steps or disclosures. In practice, this allows QA teams to focus their manual effort where it matters most, such as edge cases, high-risk interactions, or coaching opportunities, rather than spending time on routine reviews. It also enables earlier detection of systemic issues, such as when a new policy change or product update leads to widespread confusion in customer conversations.

Zendesk QA further strengthens its value by linking internal QA outcomes with customer-facing metrics like CSAT. This combined view allows leaders to explore how agent behaviour correlates with customer satisfaction, repeat contact, or escalation. While these relationships are inherently complex, the ability to analyse quality scores alongside customer outcomes provides a more nuanced understanding of what drives positive experiences beyond surface-level metrics.

Despite these strengths, Zendesk QA remains fundamentally focused on *how* support is delivered, rather than *why* customers are reaching out or *what* broader risks are emerging. Its analytics are centred on agent performance, coaching effectiveness, and quality trends, not on systemic customer issues. The platform does

not perform automated topic clustering across conversations, nor does it provide root-cause analysis that explains which product, policy, or process failures are driving contact volume.

Similarly, Zendesk QA is not designed to operate as a predictive analytics layer. While it can highlight quality risks and sentiment issues at the interaction level, it does not model churn, forecast revenue impact, or generate account-level health signals.

These limitations mean that insights tend to remain within the support organisation rather than flowing naturally into customer success, product, or executive workflows.

Compliance and risk detection represent another boundary. Although AutoQA can identify certain deviations from defined scripts or steps, organisations in heavily regulated industries may find that Zendesk QA lacks the depth required for continuous auditing of sensitive data exposure or regulatory obligations. For example, detecting PCI data leakage, GDPR data subject requests, or nuanced policy violations typically requires more specialised conversation intelligence tooling.

In practical terms, Zendesk QA is best understood as a quality enablement layer rather than a comprehensive support analytics platform. It excels at improving consistency, fairness, and scalability in agent evaluation and coaching, especially for teams deeply embedded in Zendesk. However, organisations seeking early warning signals for churn, real-time escalation monitoring, root-cause analysis, or cross-functional insight distribution will usually need to pair Zendesk QA with complementary analytics solutions.

In summary, Zendesk QA delivers strong value where quality assurance discipline and coaching maturity are the primary objectives. Its tight integration with Zendesk, structured scorecards, and full-coverage AutoQA make it a compelling choice for support teams focused on execution quality. Its limitations emerge when organisations expect support data to inform predictive risk, product strategy, or revenue outcomes beyond the support function itself.

Where Zendesk QA Fits Best

- Teams already using Zendesk who want integrated QA without additional vendors
- Organisations building structured coaching programmes tied to real conversations
- Support leaders seeking full-coverage automated scoring rather than manual sampling

Gaps to Consider

- Limited operational analytics beyond QA scores and agent performance
- Root-cause analysis and topic clustering are not primary strengths
- No account-level churn prediction or customer success integration
- Compliance and sensitive data monitoring are not deeply developed



SUPPORT

Totango

AI-Driven Churn Intelligence for CS
Teams

Totango is a customer success platform that has increasingly positioned itself around AI-driven intelligence for retention and revenue growth. In recent years, Totango has expanded beyond traditional success workflows into analytics that incorporate signals from across the customer lifecycle, including support interactions. This shift reflects a broader industry trend: churn rarely originates in a single team, and support conversations often contain early indicators of risk that account health scores alone fail to capture.

At the core of this expansion is Totango's "Unison" growth intelligence engine. Unison uses machine learning to unify data from multiple sources, such as CRM records, product usage, customer communications, and support interactions, into a single customer model. Rather than treating support as a standalone operational function, Totango treats it as one input among many that contribute to overall account health. By analysing engagement data from calls, emails, meetings, and support tickets, Totango's models aim to detect patterns associated with declining engagement, dissatisfaction, or churn risk.

This approach enables Customer Success teams to receive earlier warnings about at-risk accounts, often well before renewal discussions begin. For example, Totango can adjust health scores based on negative sentiment detected in support conversations or an increase in support activity that deviates from historical norms. Accounts that might otherwise appear stable based on product usage alone can be flagged for proactive outreach, allowing CSMs to intervene before dissatisfaction escalates or becomes visible in traditional KPIs.

Totango's strength lies in its clear focus on outcomes that matter most to revenue-owning teams. Its analytics are designed to support prioritisation, forecasting, and executive decision-making. The platform can identify accounts that are trending toward churn months in advance, highlight expansion opportunities based on engagement signals, and surface which customers warrant immediate attention from success managers. For leadership teams, this provides a high-level view of portfolio health tied directly to retention and growth metrics.

However, as a broad Customer Success Platform, Totango's analytics are intentionally oriented around accounts rather than operational support execution. Its insights tend to aggregate support data into account-level risk signals, rather than examining the mechanics of support delivery itself. Day-to-day support workflows, agent performance, queue health, or escalation dynamics are not primary areas of focus. As a result, Totango may not surface issues such as process bottlenecks, recurring ticket handling problems, or agent-level quality gaps that affect customer experience in the moment.

Additionally, while Totango can ingest and analyse support ticket content, its use of that data is largely geared toward high-level risk scoring and forecasting. It does not typically provide deep, conversation-level analysis such as granular root-cause breakdowns, emerging topic detection, or automated compliance monitoring within support conversations. These kinds of insights are often essential for support leaders, product teams, or compliance stakeholders who need to understand exactly what customers are experiencing and why.

In practice, Totango works best as a strategic layer for customer success and revenue teams, using support interactions as one of several signals to inform retention and expansion strategy. Organisations that adopt Totango often pair it with dedicated support analytics or conversation-intelligence tools to gain deeper visibility into support operations and customer communications.

In summary, Totango delivers powerful, forward-looking churn intelligence at the account level, but it is not designed to replace tools that focus on real-time support performance, conversation-level insight, or compliance risk detection.

Where Totango Shines

- Early warning on churn at the account level
- Prioritization for renewals, expansions, and upsell plays
- Executive dashboards tied to revenue outcomes

What to Pair with Totango

- Conversation-level topic mining and anomaly alerts
- Operational support KPIs and agent performance analytics
- Compliance review of ticket text and sensitive data



SentiSum

Voice-of-Customer Insights from
Support Text

SentiSum is an AI-native Voice of Customer (VoC) analytics platform built to extract structured insight from unstructured customer support data. Instead of relying on surveys, manual tagging, or predefined taxonomies, it analyses the raw language customers use across support tickets, live chats, emails, app reviews, and other written feedback. This matters because surveys typically capture a small and often biased subset of customers, while support conversations reflect real, in-the-moment friction. By working directly from that source, SentiSum provides a more continuous and representative view of what customers are experiencing.

At the core of the platform is real-time text analysis powered by natural language processing and machine learning. SentiSum automatically processes large volumes of conversations, identifies sentiment, clusters related issues, and surfaces emerging topics without requiring manual tagging or rule configuration. Customers may describe the same problem in many different ways, but SentiSum groups these variations into coherent themes. For example, a login issue described as "can't access my account," "stuck on sign-in," or "password loop" can be recognised as a single underlying problem. This allows teams to quantify issues that would otherwise remain fragmented across thousands of tickets.

One of SentiSum's most valuable strengths is its ability to surface root causes behind contact volume. Rather than only showing that ticket volume is increasing, it explains what is driving that increase. Support and CX leaders can see which topics are linked to the most negative sentiment, which issues generate repeated contacts, and which problems disproportionately consume agent time. This makes it particularly useful for identifying product defects, UX friction, policy confusion, and documentation gaps that create avoidable support demand.

SentiSum also includes anomaly detection and alerting that continuously monitors incoming

support data. It can flag unusual spikes or shifts in topics or sentiment, helping teams spot emerging issues earlier than CSAT trends or backlog metrics might reveal. If a new release triggers a sudden increase in complaints, SentiSum can surface that signal quickly, acting as an early warning system for customer pain and enabling more proactive intervention.

Because its outputs are designed to be interpretable and shareable, SentiSum often becomes a bridge between support, product, and leadership. Product teams can prioritise based on real customer friction rather than anecdotal feedback. Operations teams can identify what drives support load.

Executives can track sentiment and theme trends without reading raw ticket data. In practice, it helps align stakeholders around a common, evidence-based view of what customers are saying.

However, SentiSum is intentionally focused on insight generation rather than operational execution. It does not manage tickets, route conversations, or provide detailed agent-level operational analytics such as handle time, SLA compliance, or workforce performance. It answers “what is going wrong and why,” rather than “how efficiently is support operating,” so it typically complements rather than replaces a helpdesk platform.

Similarly, while SentiSum can reveal signals associated with churn risk such as sustained negative sentiment or repeated unresolved issues, it does not explicitly model churn or generate account-level forecasts. It does not maintain health scores, renewal timelines, or customer success workflows. Teams that want to translate VoC signals directly into churn prediction or expansion modelling usually need to connect SentiSum outputs into a broader customer success or analytics stack.

It also focuses primarily on textual insight rather than compliance or policy enforcement. While it can highlight problematic themes and emotionally charged conversations, it is not designed to automatically flag regulatory risks, sensitive data exposure, or policy violations within support interactions. Those use cases generally require more specialised conversation intelligence or compliance-focused tools.


In summary, SentiSum is a strong, AI-driven VoC layer that makes unstructured support feedback understandable, measurable, and actionable, with notable strengths in theme clustering, root-cause analysis, and anomaly detection. Its main limitations are that it stops short of operational support analytics, churn forecasting, and compliance monitoring, so it is most effective when paired with tools that cover execution and prediction.

Best-Use Scenarios for SentiSum

- Rapidly understanding emerging issues without manual tagging
- Quantifying drivers of contact volume and negative sentiment
- Providing VoC trends to product, marketing, and ops leaders

Gaps to Consider

- No native churn prediction or account-level health scoring
- Limited operational support KPIs and agent performance analytics
- No built-in compliance or sensitive-data risk detection



Cresta

**Real-Time Agent Performance & Quality
Management**

Cresta takes a distinct approach to support analytics by focusing on real-time conversation intelligence designed to improve agent performance and interaction quality as conversations unfold. Rather than treating analytics as a retrospective reporting layer, Cresta embeds AI directly into live customer interactions, influencing outcomes in the moment rather than after the fact. This positioning makes it particularly well suited to high-volume, real-time support environments such as contact centers and chat-heavy support teams.

Cresta is also particularly strong in the area of quality assurance and compliance. The platform applies automated quality checks across every conversation, continuously evaluating whether agents and AI bots adhere to approved scripts, policies, and regulatory requirements. If an interaction includes an unapproved claim, misses a required disclosure, or violates a defined compliance rule, Cresta flags the issue automatically. This approach replaces traditional manual QA processes, which typically review only a small fraction of interactions, with full-coverage monitoring that significantly reduces regulatory and brand risk.

An additional advantage of analysing all conversations is the ability to compare performance across human agents and AI-driven automations. Cresta provides a unified view of outcomes across both, helping organisations understand where automation performs well and where human intervention remains essential. This is especially valuable for teams operating hybrid support models that combine bots, assisted agents, and fully human interactions.

While Cresta can also aggregate customer feedback and surface recurring themes similar to a Voice of Customer tool, this is a secondary outcome of its design rather than its primary purpose. The platform's core focus remains operational excellence: improving agent effectiveness, ensuring consistent quality, and optimising conversations in real time. As a result, Cresta is most often adopted by support

operations and contact center leaders rather than customer success or revenue teams.

In terms of limitations, Cresta is not designed to provide account-level health scoring, long-term churn forecasting, or renewal-focused analytics. Its insights are anchored at the interaction and agent level, not the customer account or lifecycle level.

For organisations looking to predict churn months in advance or tie support signals directly to renewal and expansion strategy, Cresta typically needs to be paired with customer success or support analytics platforms that operate at a higher level of abstraction.

In practice, Cresta excels in environments where live interaction quality, compliance, and agent performance are mission-critical. It complements other analytics tools by improving the quality and consistency of support conversations themselves, ensuring that downstream data used for VoC analysis, churn modelling, or product insights is cleaner, more reliable, and more actionable. For organisations with large-scale live support operations, Cresta serves as a powerful real-time control layer for customer interactions.

Cresta first gained visibility through its AI-powered Agent Assist capabilities, which provide live guidance to agents during customer conversations. These include suggested responses, coaching cues, and prompts based on what the customer is saying and how the conversation is progressing. Over time, the platform has expanded beyond assistance into a broader conversation intelligence layer that analyses both human- and AI-assisted interactions across voice and digital channels. Its Conversation Intelligence module processes 100 percent of conversations, enabling comprehensive coverage rather than sampling-based analysis.

A key strength of Cresta lies in its ability to measure quality and outcomes at the conversation level. The platform uses predictive models to monitor signals such as customer sentiment, resolution confidence, escalation risk, and conversion likelihood in real time. Rather than waiting for post-interaction surveys or delayed CSAT results, supervisors can see early indicators that a conversation is going off track. For example, if sentiment drops sharply mid-call or an agent deviates from proven resolution patterns, Cresta can surface that signal immediately, allowing for intervention while the interaction is still active.

This real-time visibility is paired with rich supervisory dashboards that highlight trends across agents, teams, and channels. Leaders can identify which behaviors correlate with successful resolutions, shorter handle times, or higher customer satisfaction, and which patterns lead to friction or escalation. Over time, this enables data-driven coaching, where training and feedback are grounded in observed conversation outcomes rather than anecdotal reviews or limited QA samples.

Cresta's Sweet Spots

- Real-time coaching and on-call guidance for agents
- Automated quality assurance and compliance monitoring
- Unified insights across human agents and AI automations

Gaps to Consider

- No churn prediction or account-level health scoring
- Limited VoC and root-cause analysis beyond individual interactions
- Less suited to asynchronous ticket data and cross-team success workflows



Isara

Conversation Intelligence Bridging
Support and Success

Isara is a next-generation support analytics platform designed specifically to bridge the long-standing gap between customer support and customer success. Where most tools optimise for either operational efficiency or account-level forecasting, Isara treats support conversations themselves as the most reliable early signal of customer risk, satisfaction, and growth potential. Its core premise is simple but powerful: the earliest indicators of churn, expansion, and trust erosion almost always appear in support interactions before they show up in usage data, surveys, or renewal metrics.

One of Isara's most distinctive strengths is its **built-in compliance and risk monitoring**. The platform continuously audits support conversations for regulatory and policy risks, such as the presence of sensitive personal data, payment card information, or data subject requests that may otherwise go unnoticed. Unlike traditional compliance processes that rely on manual sampling, Isara analyses every interaction, giving organisations a far clearer view of their true exposure. This is particularly important for companies operating in regulated environments, where compliance failures often originate in everyday support exchanges rather than formal workflows.

Isara also brings a more outcome-oriented approach to **agent performance and coaching**. Instead of focusing solely on surface metrics like handle time or CSAT, it evaluates interactions based on whether issues were genuinely resolved and how customer sentiment evolved throughout the conversation. This allows support leaders to coach agents on behaviours that actually drive positive outcomes, not just operational efficiency. Over time, this leads to higher-quality interactions, more consistent customer experiences, and cleaner downstream data for analytics and forecasting.

Crucially, Isara is **designed to operate across teams rather than reinforcing silos**. Insights generated from support conversations are structured so they can be consumed by

customer success, product, operations, and leadership stakeholders alike. Product teams can use Isara to identify recurring friction points and prioritise roadmap decisions based on real customer pain. Success teams gain actionable churn and expansion signals. Executives receive a clear, continuously updated view of customer health grounded in lived customer experience rather than abstract metrics.

What ultimately differentiates Isara is the way it combines capabilities that are typically spread across multiple tools. It brings together **conversation-level intelligence**, **churn prediction** from support signals, **real-time escalation alerts**, automated tagging and anomaly detection, and continuous compliance auditing within a single platform. Rather than optimising one slice of the customer lifecycle, Isara provides a unified layer of intelligence that turns support conversations into a strategic asset for retention, growth, and risk management.

For organisations looking ahead to 2026 and beyond, Isara represents a shift in how support data is valued. Instead of being treated as operational exhaust or retrospective evidence, support conversations become a live input into decision-making across the business. This **holistic, AI-driven approach** enables companies to act earlier, align teams more effectively, and ensure that critical customer signals never fall through the cracks.

At its foundation, Isara connects directly to support channels such as helpdesks, email, and chat, ingesting both historical and real-time conversations. Using AI and machine learning, it analyses the **full context** of these interactions, not just keywords or sentiment scores. This allows Isara to detect subtle but meaningful patterns in customer language, behaviour, and escalation dynamics that are typically invisible to traditional reporting. Rather than asking teams to define rigid rules or tagging frameworks upfront, **Isara continuously learns from incoming data**, adapting as customer language and products evolve.

For Customer Success teams, Isara functions as a **true early warning system for churn and revenue risk**. The platform automatically flags accounts that exhibit rising frustration, repeated unresolved issues, abnormal contact patterns, or sudden changes in tone. These alerts are generated before health scores dip or renewal conversations begin, giving CSMs valuable time to intervene proactively. In practice, this means

success teams can prioritise outreach based on real customer signals rather than lagging indicators. Isara also identifies expansion and upsell opportunities by detecting moments where customers express interest in advanced features, higher usage limits, or new use cases during support interactions, turning what are often overlooked conversations into actionable revenue signals.

On the customer support side, Isara helps teams move decisively from reactive firefighting to **proactive service management**. Real-time sentiment analysis and escalation detection surface at-risk conversations as they unfold, allowing managers to step in before issues escalate into complaints or churn events. Isara's automated topic clustering groups related conversations without manual tagging, enabling support leaders to immediately see when a specific bug, workflow change, or policy update is driving contact volume. This near real-time visibility is especially valuable during product launches, incidents, or periods of rapid change, when traditional reports arrive too late to be useful.

While Isara offers a broad and forward-looking set of capabilities, there are areas where teams should be aware of trade-offs and maturity considerations.

Isara's focus on predictive and cross-functional intelligence means it is not designed to replace core helpdesk systems or real-time agent assist tools. Teams looking primarily for live agent scripting, response suggestions, or real-time call coaching may still require complementary solutions alongside Isara. The platform is optimised for insight generation and decision support rather than in-the-moment agent guidance.

Finally, as with any AI-driven system, trust and adoption evolve over time. While Isara surfaces early warning signals and risk indicators ahead of traditional metrics, organisations often need to build internal confidence by validating these signals against historical outcomes. This learning curve is typically short, but it requires teams to shift from reactive reporting to proactive decision-making based on probabilistic insights rather than absolute certainties.

In summary, Isara represents a new class of support analytics platform that treats customer conversations as a strategic source of intelligence rather than operational by-products. By combining real-time sentiment analysis, automated topic clustering, compliance monitoring, and churn prediction, it enables organisations to detect risk earlier, act faster, and align support, success, and product teams around the same customer signals. While it requires thoughtful onboarding and is best paired with execution-focused tools, Isara's strength lies in turning everyday support interactions into actionable insights that drive retention, trust, and long-term growth.

When to Choose Isara

- You need churn warnings from support conversations before health scores dip
- You want support and success working from one stream of conversation insights
- You must surface compliance risks like PCI/PII exposure inside tickets



Conclusion

Choosing the Right Mix for 2026


As customer expectations continue to rise, support and success leaders in 2026 must rely on analytics tools to stay proactive rather than reactive. The five platforms covered in this guide each address a different part of the support analytics landscape.

Established solutions such as Zendesk provide the operational dashboards and workflow automation that underpin day-to-day support execution. Customer success platforms like Totango add predictive intelligence to help teams focus on the right accounts and anticipate churn earlier in the lifecycle. Specialized AI tools such as SentiSum turn large volumes of unstructured feedback into clear Voice of Customer insights, making it easier to understand what is driving customer frustration. Cresta focuses on real-time interaction quality, helping teams improve agent performance, consistency, and compliance as conversations happen.

Isara brings these threads together by treating support conversations as a shared source of intelligence across teams. By translating conversation data into early churn signals, escalation risks, coaching opportunities, and growth indicators, Isara bridges the gap between support and customer success. This unified approach ensures that critical customer signals are surfaced early and can be acted on by the right stakeholders.

When evaluating support analytics tools, leaders should start with their most pressing goals. Teams looking to improve operational efficiency may prioritize workflow automation and SLA visibility. Those focused on retention should look for early warning signals tied to real customer behavior. Organizations struggling to understand why customers are contacting support may benefit from AI-driven text analytics. Companies operating in regulated environments may need real-time quality and compliance monitoring. And teams aiming to break down silos between support, success, and product should look for platforms that unify conversation insights across the business.

By choosing the right mix of tools, support organizations can move beyond reactive troubleshooting and become proactive contributors to customer success and growth. The outcome is not only better customer experiences, but also measurable improvements in retention, trust, and long-term revenue.



Thank you for reading.

**We hope this guide was useful in
shaping how you think about support
analytics.**

Sources and Further Reading

- Customer experience benchmark: <https://www.salesforce.com/ap/blog/customer-experience/>
- Customer experience strategy: <https://www.superoffice.co.uk/resources/articles/customer-experience-strategy/>
- Future of customer experience: <https://www.pwc.com/us/en/services/consulting/library/consumer-intelligence-series/future-of-customer-experience.html>
- Zendesk intelligent triage: <https://support.zendesk.com/hc/en-us/articles/4471123173402-Intelligent-triage-resources>
- Zendesk Answer Bot metrics: <https://support.zendesk.com/hc/en-us/articles/4408824748698-Metrics-and-attributes-for-Zendesk-Answer-Bot>
- Totango Unison: <https://www.totango.com/products/unison>
- Totango AI churn intelligence: <https://www.totango.com/press/totango-announces-ai-driven-churn-intelligence-solution-following-acquisition-of-technology-and-team-from-parative-ai>
- Totango press coverage: <https://www.prnewswire.com/news-releases/totango-announces-ai-driven-churn-intelligence-solution-following-acquisition-of-technology-and-team-from-parative-ai-302282336.html>

- Unison support: <https://unison-support.totango.com/hc/en-us/articles/32426812291348-About-Totango-Unison>
- SentiSum ticket analytics: <https://www.sentisum.com/support-ticket-analytics>
- SentiSum anomaly detection: <https://help.sentisum.com/what-is-anomaly-detection>
- Cresta quality management: <https://www.cresta.com/cresta-quality-management>
- Cresta conversational analytics: <https://cresta.com/guides/conversational-analytics>
- Cresta AI agent optimization: <https://www.cresta.com/ai-agent-optimize>
- Cresta conversation intelligence: <https://www.cresta.com/conversation-intelligence>
- Isara compliance risks in support: <https://www.isara.ai/blog/the-hidden-compliance-risk-inside-everyday-support-conversations>
- Isara PCI DSS in support: <https://www.isara.ai/blog/preventing-card-data-exposure-in-customer-support-pci-dss-compliance-with-isara>
- Isara CCPA in support tickets: <https://www.isara.ai/blog/how-to-ensure-ccpa-compliance-in-customer-support-tickets-with-isara>