

# Inside the Intelligence

## Machine Learning and NLP, Purpose-built to Transform Support

### LEADING WITH TRANSPARENCY

Most artificial intelligence (AI) platforms hide their technology within a black box of proprietary claims. This paper does the opposite: revealing in detail how SupportLogic works so you can make a more informed decision. SupportLogic is built on a five-part AI design philosophy:



**AI should be like electricity—surrounding you with value but only noticed when it's absent.** To do this correctly, AI must be infused into workflows—otherwise it gets in the way and adoption fails. Every bit of SupportLogic machine learning (ML) and natural language processing (NLP) is built into workflows designed to enhance your day-to-day operations—the intelligence working behind the scenes to boost your resources.



**AI should be explainable—you should know exactly what causes every bit of insight, building confidence and trust in the predictions.** No black boxes or vague check engine lights. When SupportLogic displays an insight or recommendation, the signals are right in the interface—showing you the source and keeping you informed.



**AI should be domain-specific and purpose-built, never generic.** Domain expertise is crucial to getting value from AI. SupportLogic SX™ is trained on millions of support cases. This means the models can detect sentiments and urgency signals from your data right out of the box. Past AI products have tried metadata approaches with mixed results. Trial and error have proven that including your unstructured support case data in ML models improves the quality of the insights and predictions.



**AI should be fine-tuned to your products, customers, and team.** AI shouldn't catch everything, but what it does catch should always be worth your time—accuracy at the expense of recall. SupportLogic models are tuned to your team's bandwidth, products, and customers. If your organization can handle five escalations a day, the escalation prediction model is tuned to detect the five cases most likely to escalate, giving you the greatest opportunity to provide support where it's needed most.



**AI should be adaptable, and evolve with a feedback loop of improvement.** This is often where in-house AI solutions fail, as they aren't designed to change with your business. SupportLogic has a dedicated team of ML scientists specifically focused on improving signals found in support cases. You can't easily hire that skill nor could you match our years of experience doing it. You also benefit from the collective intelligence built from a global dataset of customers.

### CONTENTS

[Why AI? Why Now?](#)

[AI vs. Purpose-Built AI](#)

[AI Tuned to Your Data](#)

[Evolving NLP](#)

[Sentiment Labeling](#)

[Customer Sentiment Explained](#)

[Inside the Models](#)

[Actionable Customer Sentiment](#)

[Intelligent Escalation Prediction](#)

[Intelligent Case Assignment](#)

[Proactive Support Health Scoring](#)

[Coaching and Case Evaluation](#)

[Closing Points](#)

SupportLogic SX™ wasn't created by an analytics company looking for a purpose, nor was ML added to a support platform as an afterthought. From day one, the goal has been to transform customer support from reactive to proactive and predictive.

## WHY AI? WHY NOW?

Your customer support team is instrumental to defining customer experience, customer satisfaction, and identifying new business opportunities. The only way to achieve these objectives and continually stay ahead of the competition is through purpose-built AI. ML and NLP—two subcategories of AI—have the ability to unlock insights from the data sitting in every interaction you have with your customers. These insights can then provide actionable recommendations to the support, product, engineering, sales, and customer success teams charged with improving your customer experience.

The business outcomes are what is most important:

- Reducing your average resolution time
- Addressing your customers' needs before they ask
- Getting ahead of the smoke before the fire to leverage the right resources
- Unlocking the hidden value from your support team and case queue

ML and NLP have been used in applications ranging from self-driving cars and voice assistants to ad targeting and supply chain management. But these terms are more often used as marketing fodder than to denote genuine technological superiority. We believe ML and NLP are **foundational** to transforming B2B customer support.

## AI VS. PURPOSE-BUILT AI

AI helps you make better decisions. You rely on it to reroute your driving directions to avoid traffic jams. But that common example also demonstrates the need for purpose-built AI: you wouldn't want AI designed for navigation to make decisions for your customer support processes. Support teams need to consider customer value, related products, contractual obligations, available resources, and many other support-specific variables, all while grabbing relevant information from nuanced and highly variable human language.

Drilling deeper, support-specific, purpose-built AI must:

- Accurately capture the voice of the customer in every case
- Determine urgency and send the most urgent or customer-critical cases to the top
- Evaluate and track your team holistically, without bias
- Identify your best agents and provide guidance for getting more from your broader team

You rely on your team to review, interpret, and categorize every incoming customer communication, and further rely on them to identify best practices and discover hidden opportunities. But as case volumes and tribal knowledge increase, humans can't keep up efficiently or effectively. AI can, and your competitors have likely already made investments in AI to boost their own support capabilities.

Some opt to "build" rather than partner. Every pitfall in developing an in-house solution stems from missing the mark on the five tenets of effective AI discussed on page 1. Customer support infused with the right AI solution is the true competitive advantage—boosting a critical part of the business to anticipate customer needs, involve your entire org, and increase satisfaction.

## AI TUNED TO YOUR DATA

Purpose-built AI will fall short unless it is trained from the data in your ticketing system. Support-specific NLP reads customer conversations and turns them into useful information. This is part of what separates generic AI from purpose-built AI. Generic language processing may recognize generally satisfied versus unsatisfied sentiment by tagging terms like “happy” and “thrilled” versus “unhappy” and “frustrated”. Support-specific NLP, however, understands words and phrases as they relate to support like “delay”, “urgent”, “contract”, “SLA”, “escalate”, and so on.

**Effective AI should be tailored to consider your resource constraints and costs as well as your customer hierarchy and targeted goals.**

**To be truly useful to you and your business, NLP must recognize your product names, numbers, versions, industry terms, account names, and other language unique to your business.**

These two layers of SupportLogic NLP—support-specific and business-specific—are highly-tuned algorithms that use your customer support data to pick out customer sentiment, tag relevant words, products, and terms, and highlight the relevance to support. From there, support-specific ML models predict future customer actions and recommend proactive support actions to arrive at the most desirable outcomes, all while simultaneously considering your resource costs, process times, operational rules, and more. This separates purpose-built AI from the basic keyword detection or rules-based technology in your CRM that require manual maintenance and adjustment.

## EVOLVING NLP

Your business continues to add product names and numbers, terms, customers, partners, and more, all while your customers (and humans in general) evolve in how they use, describe, and interact with support. That evolution requires NLP and ML models that can keep up. SupportLogic data scientists constantly update the platform models and your custom-tuned models based on support-specific terms and phrases. They then comb through customer data for examples and use customer feedback to tune the accuracy of the models based on how they tag, categorize, and label your support data. To check for precision, data scientists pull random sets of labeled data for evaluation—counting true and false positives. The SupportLogic platform also allows you the user to manually label phrases, which are then added to the model after an audit for relevancy.

In the NLP model, support-specific NLP determines *what your customers mean*, and then purpose-built ML predicts *your customers' likely action* and recommends your best action. Deep learning is a method used for customer sentiment because of its advanced ability to read language – uncovering the meaning behind individually ambiguous words by using the surrounding words to define context rather than relying solely on sequential words. BERT and GPT both enable NLP to read forwards and backwards to better discern context and flow, helping predict subsequent sentences.

Additionally, training on millions of words helps capture the relationships between words. Meaning, that if the negative sentiment labeler is trained to capture “this is terrible”, it will likely capture “that is bad” without having to specifically list synonyms. This is crucial because it eliminates having to “boil the ocean” on language understanding, or make lists upon lists of rules. It’s important to note that most NLP is trained on generic language. To further tune to support, over 60 million B2B customer support interactions have been used to train SupportLogic models for your applications.

## SENTIMENT LABELING

Labeling is used to continually improve the accuracy of the NLP models as changes occur in support, your business, and the way your customers interact. Developers, linguists, and data scientists identify, label, and add value to relevant words and phrases as they occur in your support interactions, helping the ML models better find and use different key terms, expressions, and sentence structures that can better classify the NLP data.

**If the NLP and ML models remain static while your data, customers, and business change, the results will lose accuracy over time.**

**Any AI platform, no matter how good, must be continuously updated to keep up with change.**

The SupportLogic platform recognizes subtle changes in customer language that may have been previously unimportant but are now critical to avoiding case escalations. It also recognizes and adjusts to changes in your data schema, internal support processes, and other factors. This refinement can be broken down into two methods:

- **Labeling:** Using your feedback to discover specific phrasing that is relevant to your business. While SupportLogic data scientists are highly experienced in customer support, you know the language most relevant to your business. Those labeled phrases are then added along with phrases discovered through combing the data to make patterns that capture similar phrases with different syntax and vocabulary.
- **Scoring:** Labelers manually label the data, creating a sufficiently large dataset to apply a model that uses word embeddings (meaning the relationships between words) on unlabeled data to find new phrasings that have not yet been thought of or seen in your data.

## CUSTOMER SENTIMENT EXPLAINED

Customer sentiment is the discovery and measurement of how customers feel about your product, brand, or service. In the context of support, it's generally more difficult to determine sentiment because the reason a customer calls support is because something is wrong—a negative event is the topic of the call.

It can be both art and science, as "I'm frustrated" obviously conveys a negative sentiment while something like "it's not what we expected" could be either positive or negative. Three levels of analysis, combined with support domain expertise and millions of actual customer support examples, continually improve the NLP and evolve the ML models.

- **Fine-grained analysis** measures positive or negative sentiment by determining the polarity score of words in a text.
- **Emotion detection** builds on that by identifying customers' emotions and moods, such as satisfied or frustrated.
- **Aspect-based sentiment analysis** breaks text into phrases, sentences, and tokens to measure the accompanying sentiment.

When combined, it can accurately determine your customer's emotions, understand their actions, and help you make informed decisions that improve the customer experience.

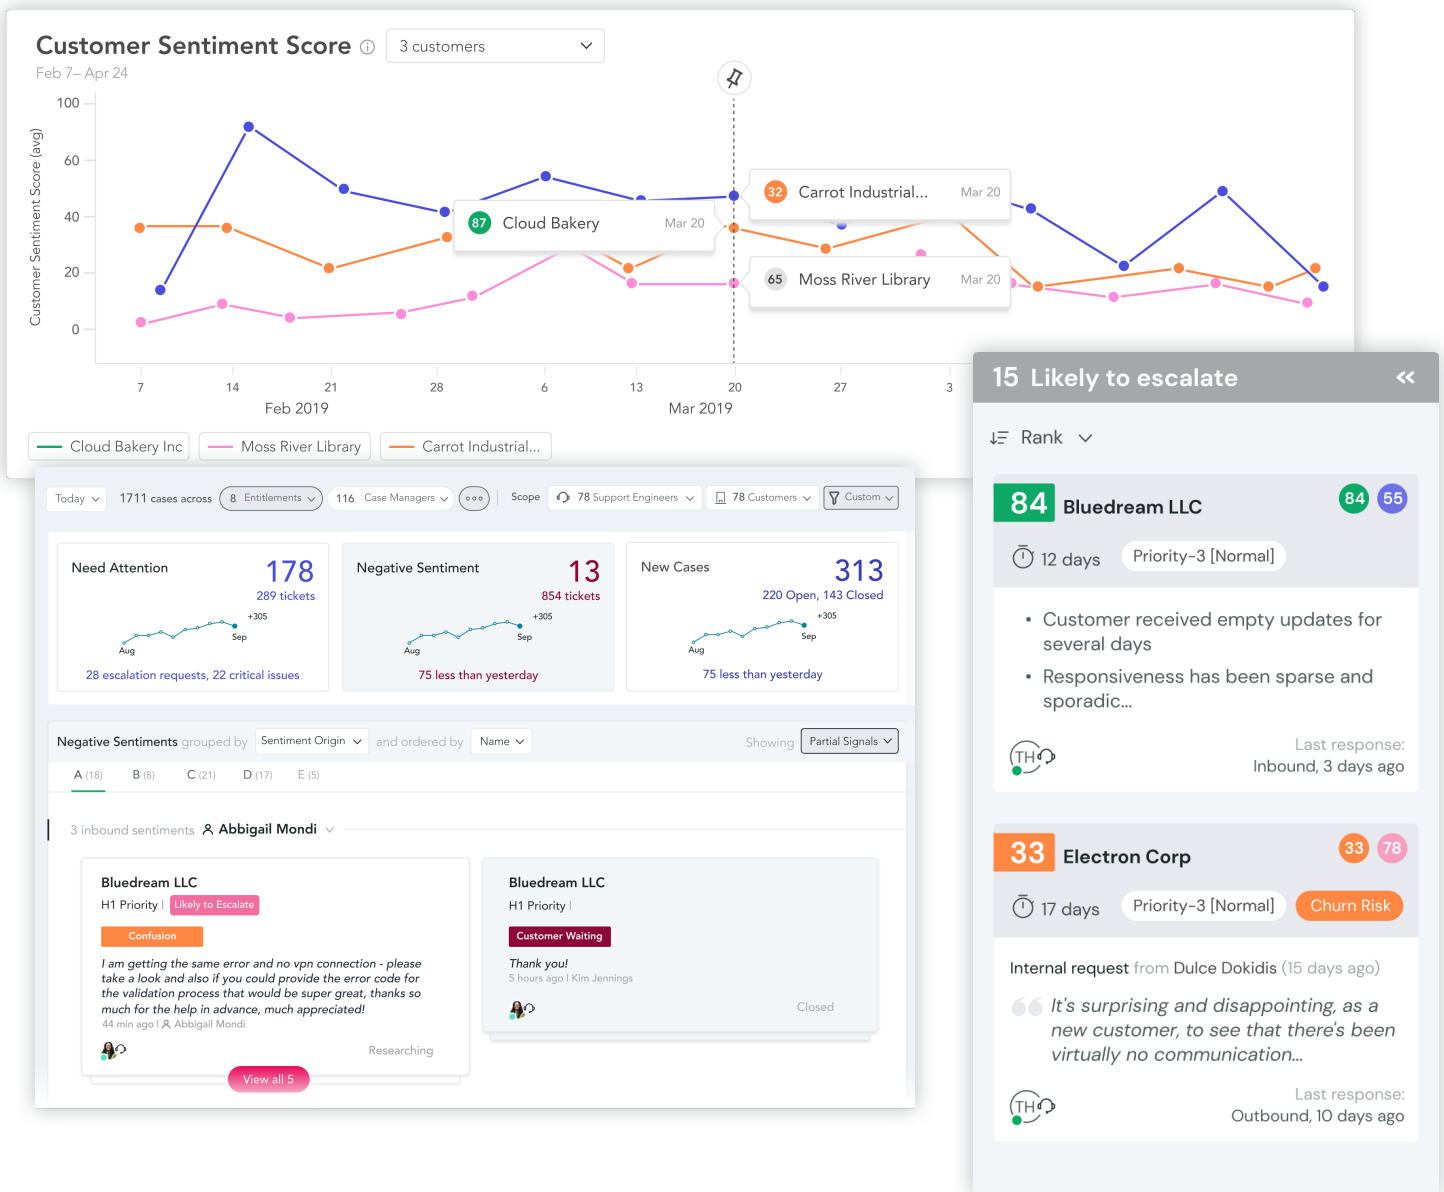
## INSIDE THE MODELS

In every aspect of support, there's value in using AI to make predictive insights, speed up workflow, and boost the efficiency (and therefore value) of your team.

40 different customer signal categories and thousands of domain-specific keywords are extracted from your past and current support cases. These customer sentiment signals pull the true voice of the customer from your case queue. This section breaks down five SupportLogic models that use those signals.

SupportLogic models learn as your support team accepts or rejects predictions and guidance. The ML then uses those inputs to quickly adapt and become more accurate. With every support interaction, the models and data science evolve to stay out of the way while helping you more easily prevent escalations, identify at-risk customers, reduce churn, spot product issues, and even discover ways to drive more customer-focused product road maps.

A sample of customer insights: Sentiment compared across customers, a list of the cases in your queue most likely to escalate, and the exact case language leading to a predicted escalation request. These are just some of the indicators from SupportLogic SX.





## ACTIONABLE CUSTOMER SENTIMENT

NLP applies labels to various words and phrases which can then be inferred as different sentiment types. Signals and sentiments tell you how a customer is feeling without having to read and monitor every case. The sentiment labeler relies on patterns to identify spans of text that relate to each label, continuously improving by using examples from customer data. For example, recognizing “the server is broken” as a negative sentiment triggers a process that captures a more broad set of phrases, such as “the system is broken”, “the server is down”, or “the system isn’t working”. NLP experts research similar usages of the phrase, constructing patterns based on part of speech, vocabulary and root words (“lemmas”) to create efficient patterns that capture a wide variety of related phrases.

The sentiment labeler categorizes sentiment into six categories articulated in the language of support:

- Negative
- Positive
- Need Attention
- Product Feedback
- Customer Request
- Escalation Activity

The signals are then compiled and converted into **Case Sentiment Score** and **Needs Attention Score**—patented indicators that reflect the overall health of every case (and customer). These scores make the NLP usable within your ticketing queue.

**Like a speedometer, scores give you the quick read on every case’s health. When case sentiment turns negative or urgent, the model immediately reduces the Case Sentiment Score and increases the Needs Attention Score, pushing the case to the top of the queue.**

These leading indicators help you quickly address the most pressing cases, letting you swarm resources, reassign a case, or check in with your team on status.

The screenshot displays a support ticketing system interface. On the left, a 'My Backlog (24)' shows a list of tickets with status tags: 'Missed SLA' for North Star LLC, 'Churn risk' for Dynamico, 'Approaching SLA' for Rosewood Co., 'New sentiment' for Sapling Industries, 'Waiting for customer' for Drift Systems, and 'New case' for Riddleaid LLC. The main view shows a detailed ticket for North Star LLC (ID 35) with the subject 'gnome-shell: st\_bin\_dispose(): gnome-shell killing the memory'. It includes 'Key Insights' (e.g., 'Follow-Up Response Time SLA Missed! 2 hours ago'), 'Case Comments' (e.g., 'Today, 4:44 PM by A Darren Waller'), and a 'Follow-Up Response Time SLA Missed!' alert. A floating window shows sentiment and attention scores for Blue Dream LLC (Sentiment Score 25, Needs Attention Score 85) with various tags like 'Escalation Request', 'Customer Waiting', and 'Urgency'. A text box on the right explains that these scores provide a quick read on case health to help manage the backlog.

## INTELLIGENT ESCALATION PREDICTION

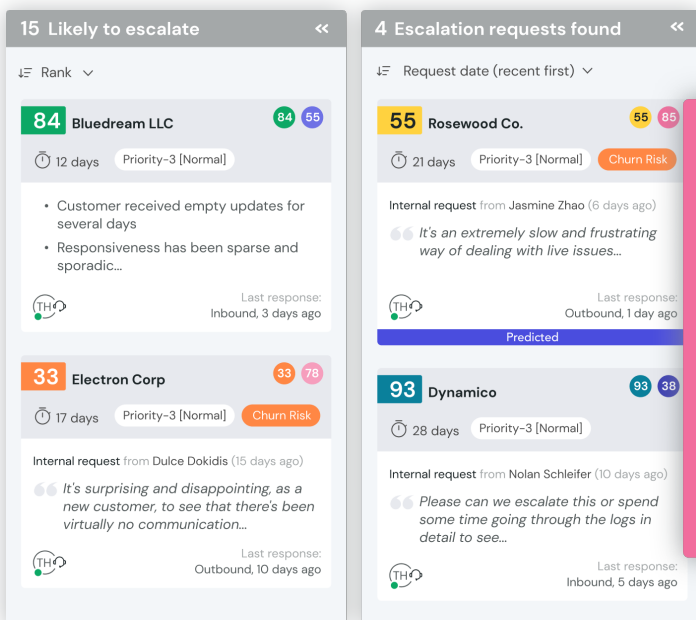
Escalations can cost two- to three-times more to address and resolve than a regular case. The SupportLogic escalation model accurately predicts which cases are likely to escalate in the next 72 hours so your teams can swarm to manage and prevent them.

Decision trees are at the basis of the escalation model, specifically a trained version of a **Gradient Boosted Classifier**. Decision trees use predicted probabilities to classify observations and are a common basis for many machine learning models, preferred for their ability to capture complex relationships between variables in order to predict a final target (in this case, the probability of escalation within the next 72 hours). The Gradient Boosted Classifier averages the predictions of hundreds of decision trees to form a stronger predictive model which can be further tuned for even more accurate predictions. Think of this like performing the same experiment hundreds of times to settle on a predictable outcome.

The escalation model considers key case information along with the NLP-driven Case Sentiment Score and Needs Attention Score. For new cases, the model can predict an escalation in 24 hours—once enough case information has been collected. The model uses a weighted combination of five factors to show you the likelihood of an escalation occurring.

- **Customer history:** This factor uses both the customer account and the individual’s case history as indicators of future escalation behavior.
- **Case activity:** Cases with a high level of activity or even a sudden burst of activity tend to escalate. This factor monitors this frequency and uses it in the prediction.
- **Case age:** This factor tracks how some customers are likely to escalate earlier in the life of a case while other customers escalate later, referencing case history to predict escalation points in time.
- **Poor support responsiveness:** ML tracks support agent responsiveness against historical data to predict when a customer’s impatience could lead to an escalation.
- **High customer urgency:** Cases with a high level of activity or even a sudden burst of activity tend to escalate. This part of the model monitors this frequency and uses it to predict escalations.

Finally, queue depth is a key factor in tuning the escalation model to your organization’s capacity. Your organization can handle only so many escalations, so a threshold maintains a “target queue depth” with the highest precision and recall possible. Since the escalation model is constantly being refreshed (about every 15 minutes) and your queue is fluid, developments in any case can change its escalation likelihood.



LIKELY TO BE ESCALATED

25

**Emergence Inc.** 25 85

🕒 12 days Priority-3 [Normal] Churn Risk

Escalation factors

Poor support responsiveness

Customer has many escalations

Case activity

🔄 Unassigned Last response: Inbound, 21 days ago

Putting out the smoke before the fire: The reasons behind every escalation prediction display right on the case. The escalation model can predict an escalation up to 72 hours in advance, allowing you time to swarm resources.

## INTELLIGENT CASE ASSIGNMENT

Several models combine to recommend the most capable and available support agent to every case. Similar to how NLP is used to understand sentiment, case assignment reads every case and constantly evaluates every agent's case outcomes and communications to predict a compatibility with their skill level. In all, the model takes five distinct factors into account:

- **Time overlap:** This factor does so much more than look at the agent's current shift—it also calculates the estimated time it will take to solve the case and ensures that it falls within the agent's working window. After all, there's no use in assigning an agent to a case if they're off duty 10 minutes after the assignment occurs.

- **Skills match:** The NLP engine is used to scan the case for required skills, matching key phrases to the system's knowledge of every agent's skill set.
- **Customer experience:** This factor draws from an agent's success with previous similar cases. Support experience shows that agents with prior experience with a particular customer tend to resolve cases faster and with higher customer satisfaction.
- **Bandwidth:** This factor adjusts with every assignment, ensuring the best agents aren't overloaded. Cases that are predicted to take longer to resolve weigh more heavily on the bandwidth part of the model.

The case assignment recommendation menu on the right compares every available agent and shows you their ability to handle a case, based on five composite factors shown below.

**Overall match (95%)** Customer interactions

**Time overlap** 96%

9 AM | 5 PM | 10 AM | 6 PM

⚠ Rachel will be Out of Office from 05/10 to 05/15

**Case complexity match** 75%

This case: Low Medium High

Rachel's cases

**Bandwidth**

Active case load: L (+3 new cases assigned today, 2 escalations)

9 cases in the backlog

- Waiting on engineer
- Open (4)
- Closure required
- In progress (1)
- Solution saved

**Skills match** 70%

Experience working on cases with similar issues:

- Ubuntu Linux | i386
- Red Hat | Installation

**Company experience** 95%

Worked on 8 cases with this account.

Average Sentiment: Positive

**Recommended engineers for this case:**

- Angela Love** 95%
  - Active case load is Low
  - Time overlap 9 AM - 5 PM
  - Handling 3 cases from this customer
- Adam Simpson** 80%
  - Active case load is Low (+3 new cases assigned today)
  - Time overlap 9 AM - 5 PM
  - Will be Out of Office tomorrow
  - Assign**
- Amanda Richardson** 80%
  - New agent. No backlog yet
  - Time overlap 8 AM - 4 PM
- Anna Smith** 75%
  - Active case load is Medium (+3 new cases assigned today)
  - Time overlap 9 AM - 5 PM
  - Will be Out of Office tomorrow



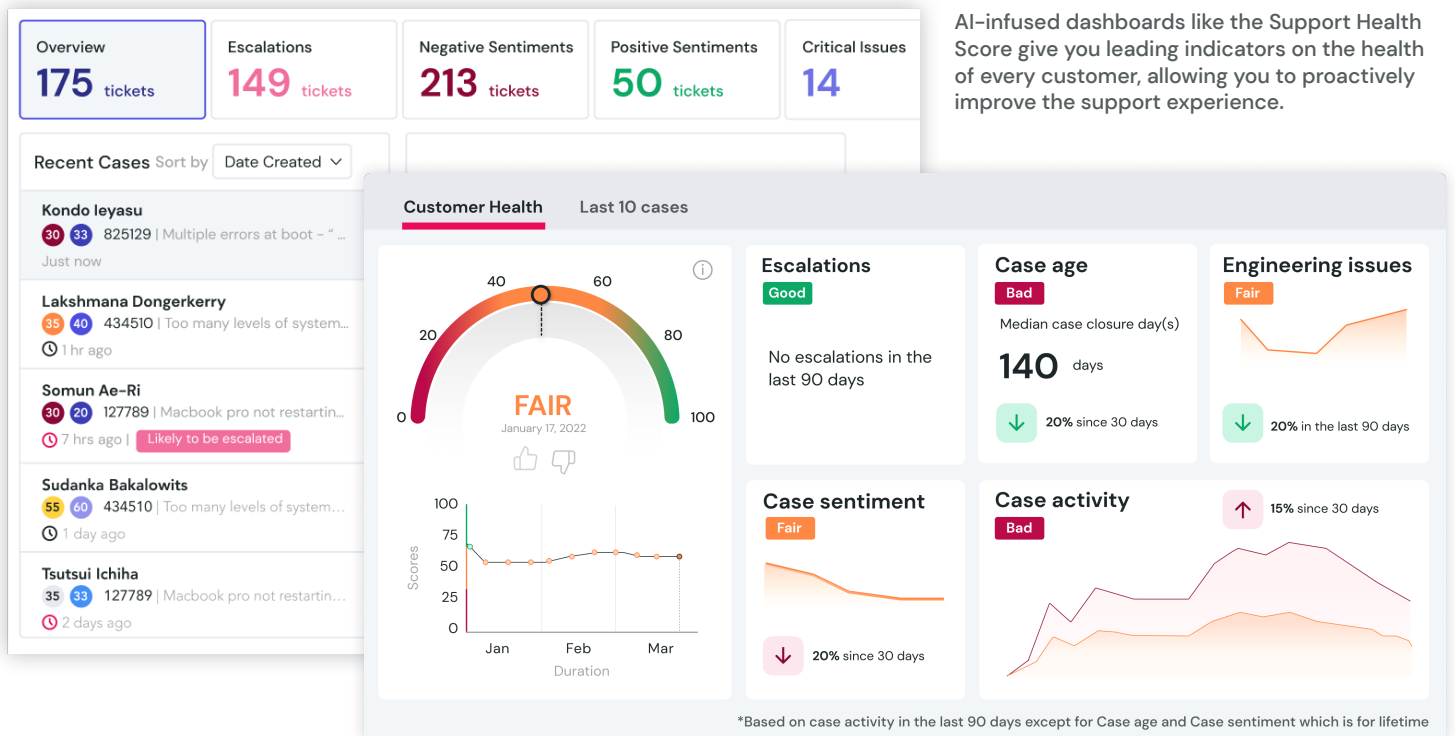
## PROACTIVE SUPPORT HEALTH SCORING

The support health score model combines several insights to provide a holistic view of every customer's support experience. The score takes several factors into account such as life of the customer, recent events, sentiment, support history, and more to create a composite support satisfaction score. Obviously, a declining customer health score could eventually lead to churn, but getting a clear view of every customer's support experience still enables you to take action to reverse trends. This model is another example of AI creating a unique leading indicator, letting you in on your customer's temperature well before CSAT or NPS (both lagging indicators).

From examining thousands of cases, trial and error built this model from the combination of factors that accurately predict a customer's support health. The model is a weighted sum of the following factors, with recent activity weighing more heavily:

- Escalations
- Case sentiment
- Support quality
- Account activity
- Case composition
- Engineering issues

A data science approach is used to integrate the information produced by other models, which is then blended into a holistic score reflecting a customer's health. These algorithms are constantly refined by adding in new events and tuning the reaction to those events individually. User feedback on the accuracy of the score gives the score a label, which can then be analyzed by ML to improve the model. Evaluating the outcomes of the support health score allows the model to identify trends in accuracy and improve the weights of the separate factors. The more feedback the model gets on its predictions, the more the model can learn, adjust the factors that make up the model, and improve.



AI-infused dashboards like the Support Health Score give you leading indicators on the health of every customer, allowing you to proactively improve the support experience.

### COACHING AND CASE EVALUATION

The coaching and evaluation model was designed to save support managers time—helping them continuously engage with agents, develop skills, evaluate to a standard, drive more constructive dialogue, and reduce agent burnout while increasing retention.

The model begins with NLP analyzing 100% of every agent’s queue against 19 signal types to find the ideal cases to review and a mix of positive, negative, and neutral interactions. This enables a balanced set of options for immediate coaching and case evaluation.

NLP highlights the relevant interactions that can be used to evaluate performance. To infuse workflow, the case evaluation queue is populated with cases predicted to be strong candidates for review.

In a coming version of this model, an agent’s case history will be automatically read against a defined rubric, highlighting the specific actions the agent took and didn’t take (did they provide a proper greeting? did they respond to a customer asking for an escalation? did they ask to close the case?). The rubric engine can also determine if a knowledge base article was written or applied, helping boost KCS policies with your support team.

Save time reviewing cases and level up the skills of your team with in-case coaching and evaluation. The AI-infused workflow keeps the whole process within one platform: reading every case for you, making coaching recommendations, and tracking agent development.

The screenshot displays a support case interface for 'North Star LLC'. At the top, it shows the case status as 'Waiting on customer' and identifies the reporter as Daria Waller and the case owner as Darren Smith. The case is categorized as 'Likely to escalate' with a 'Churn Risk' indicator. Key metrics include a 'Sentiment Score' of 35 and a 'Need Attention Score' of 88. The engineering issue is identified as 'GEOB-24400'. A 'Key Insights' sidebar on the right provides performance metrics: 'Was open for: 12 days', 'Total customer wait time: 8d, 6h, 30min', 'Total agent wait time: 10 hours', and 'In escalated state: 5 hours'. The main case view shows a customer message from Diana Washington: 'The reason I installed Avocado\_IT is that the HS300 interface is sooooo bright (white background) that this is REALLY annoying for my eyes...'. Below this, internal notes from Albert Flores (4 days ago) and Leah Vuong (5 days ago) discuss the hardware issue. A note from Daria Waller (6 days ago) quotes the customer's complaint. A 'Coaching feedback' panel on the right shows a message from Leah Vuong: 'Please remember to check if the message you're posting is an internal case note or an outbound message.' It also shows two replies: one from Leah Vuong ('Yes I'll triple-check, thank you') and one from Melody Lawson ('Thanks for the reminder (edited)').

## CLOSING POINTS

Confidence in AI grows when you know how it was built and when you can use it in workflows that show you exactly what goes into every prediction. Thousands of hours of research, engineering, data science, and support expertise have gone into developing SupportLogic models trained on support language and interactions that are then tuned to your specific environment. This puts data from thousands of your past support interactions as well as every newly incoming interaction at your fingertips for analysis and action—with models continually improving based on new data, training, and model adjustments. From there, workflows keep the AI usable in your day-to-day operations, supercharging the team instead of staying a novelty.

Generic AI will never be the difference-maker in customer support. The right solution is built on decades of customer support experience combining with adaptive models to answer support's enduring problems.

**By analyzing thousands of cases, iterating on analytical models, and labeling sentiment from day one, purpose-built, adaptive AI saves you time, improves your team, catches the smoke before the fire, and helps delight and retain your customers.**

The real proof is in customer feedback. Head to [supportlogic.com](https://supportlogic.com) for customer testimonials that prove how the SupportLogic SX platform reduces mean time to resolution (MTTR), improves CSAT, helps leverage the right resources, and unlocked the hidden value from support teams.

If you're ready to see it in action, schedule a demo with a SupportLogic expert or begin a proof of value for your organization.

