Lead scoring best practices with Mautic

YouTube recording of a talk

Steve Robinson   Brilliant Metrics

Lead scoring is an effective method to gauge the behavior of the contacts in your database. It involves adding and subtracting points based on different activities that signify a contact's intent to purchase or take action. The key point is that these contacts accumulate points over time. Once they cross a certain threshold, this is an indication that they have a strong intent and should be contacted by your team.

Managing lead scoring is a delicate balancing act, and it's important to walk a fine line between quantity and quality. If we liberally assign points for every minor action, we might end up with a lot of leads, but their quality would be poor. Conversely, if we only pass on the leads that have engaged in numerous activities and filled in all their details, we might end up with only a handful of highly qualified leads each year. Striking the right balance is crucial.

There are four keys to successfully implementing a lead scoring model:

1. **Mapping trust and intent:** This involves identifying and capturing the right signals that indicate a contact's intent.

2. **Decay model:** This is a best practice that isn't applied by every marketer. It involves reducing the score of a lead over time if they don't show any recent activity.

3. **Creating a feedback loop:** Understanding what's working and what's not is critical in refining your lead scoring process.

4. **Analyzing and improving:** Regularly reviewing the process and making necessary adjustments will ensure that your lead scoring model stays effective.

The following sections will delve into each of these steps in more detail.

Mapping trust and intent

Mapping trust and intent in lead scoring can be broken down into two types of scoring activities: low value and high value.

Low value

Low value scoring activities are where we aim to understand trust. These are actions that do not necessarily show a clear intent to purchase but rather represent early steps in a buyer's journey. Examples include submitting a non-revenue generating form, signing up for events, clicking on emails, visiting random webpages, or subscribing to an email newsletter. These activities are not explicit indicators of purchase intent, but they help us gauge a contact's engagement level. For such activities, it's recommended to add points that make up less than 10% of your scoring threshold.

High value

High value scoring activities, on the other hand, help us assess intent. These are actions that show a strong likelihood of a purchase, such as submitting a revenue-generating form, visiting pages that indicate high intent, downloading product brochures, attending webinars, or expressing interest at a trade show

booth. Depending on the activity, it's recommended to apply points between 20% and 100% of your scoring threshold. For instance, if a contact submits a form explicitly requesting to be contacted for a purchase, that's a clear indicator of intent and you can assign a score that reaches up to 100% of the threshold.

Decreasing value

Beyond these, there are also negative scoring activities which help in maintaining the accuracy of your scoring system. These activities should cause the lead's score to decrease. Examples include visiting careers pages, vendor portals, employee portals, and investor portals. If a contact unsubscribes, this is a sign of a drop in trust or relationship, and their score should be reduced accordingly.

Using a suppression segment

To ensure that irrelevant contacts do not trigger your lead scoring system and are not sent over to your sales team, it's recommended to set up a suppression segment. This is a segment that should include your employees, vendors, contacts with .edu email addresses (unless you sell to the education market), competitors, lower level job titles, and anyone who has been marked as 'marketing suspended' in your database.

Finally, to manage this suppression segment properly, you should also set up a suppression campaign. This campaign should be configured to constantly look for anyone on your suppression list who still has a score greater than zero. It should then reset their score to zero. This will ensure that anyone who shouldn't be scored doesn't inadvertently influence your lead scoring model.

In conclusion, a well-structured lead scoring model will accurately track intent, filter out irrelevant contacts, and ensure that your sales team is working with the most promising leads.

Decay model

Let's now delve into the concept of the decay model in lead scoring. This might appear complex at first glance, but it's a crucial component to consider.

What is a decay model?

The premise of a decay model is simple. If a contact's score isn't increasing, it should gradually decrease. This is based on the understanding that a contact's intent, interest, and trust in your brand will diminish over time if there is no active engagement. This is akin to saying that actions taken six months ago aren't as valuable today. Hence, it's crucial to act on current signals.

One common practice, which is not recommended, is to reset the score to zero once a lead has been dealt with. The issue with this approach is that it doesn't reflect the natural relationship of a contact with your brand. For instance, consider a scenario where someone visits your website, engages enough to hit the threshold, and a member of the team reaches out to them. However, they are not ready to make a purchase and therefore the score is reset to zero. Three weeks later, when they are ready to buy, they won't be visiting as many pages or triggering as many scoring opportunities. Therefore, a reset model might miss out on capturing this renewed intent, which is where a decay model proves beneficial.

How a decay model works

Here's how a decay model works. As a person accrues score, the model drops their score by a fixed amount every week once they reach the peak. This means if a contact wasn't ready to buy when they triggered an MQL alert, they go away for a bit and when they return with renewed intent, they engage in more scoring opportunities which in turn provides a chance for a second alert. This model, therefore, captures those contacts who disengage for a while and then return.

Implementing a decay model with Mautic

Firstly, create a segment of anyone who has a score, and set a waiting period of seven days before reducing their score by 10 points. If a contact has less than 10 points, just reset them to zero. If they have more than 10 points, reduce the score by 10 points. Then remove them from the campaign, and they will start again from the top as they still have a score. This leads to a systematic score reduction of 10 points every week.

However, ensure you also have a limit campaign in place. Without a limit, a contact who engages heavily with your website could accrue an enormous score that would take an unrealistic amount of time to decay. A recommended limit could be 70 points. If a contact's points exceed this limit, reset the points to 70. This segment should look for anyone whose score is outside your set limits. The same approach applies to negative points to prevent a contact from getting too far into the negative.

In summary, a decay model coupled with a limit campaign ensures you capture intent accurately, take note of when trust decays over time, and prevent contacts from skewing the scoring system. This approach will help you keep track of your lead scoring process and make necessary adjustments to stay on the right course.

Feedback loops

Creating a feedback loop is a crucial next step in your lead scoring model. The purpose of a feedback loop is not only to measure the quantity of leads but also the quality of those leads.

What is a feedback loop?

This systematic process allows marketing to incrementally improve by capturing the quality of leads.

To visualize the process, let's consider a scenario:

- Your marketing activities lead to website engagement.
- This engagement leads to a form fill or a scoring event.
- This event triggers a marketing qualified lead (MQL) alert to sales.

Now, this is where the feedback loop starts. Three days after the MQL alert, send a feedback request to sales. This request should inquire about the quality of the lead: Was it a good lead? Was it a bad lead? This ensures that no lead goes unchecked.

If feedback is received, move forward with your process. If not, continue to request feedback every three days. This persistence is key to ensuring that all leads are evaluated, which helps to avoid bias in reporting. Although it may seem tedious, it is crucial to get feedback on every single lead.

The feedback you receive should help determine where a lead fits into four quadrants based on two criteria: Contact and Company Fit, and Immediate Need.

Good Lead: If the contact and the company are a good fit and they have an immediate need, then it's a good lead. This means the contact is right, the company is right, and the timing is perfect.

Bad Timing: If the contact and the company are a good fit, but there's no immediate need, this is a case of bad timing. It suggests that you may have jumped the gun and interpreted the signals wrongly. In this case, you want to continue to nurture that contact and drive more leads like that.

Continuing with the feedback loop process, there are two additional types of leads that the sales team might encounter. These include bad fits and junk leads.

Bad Fit: These leads are not the right kind of company or contact, although they may have an immediate need. These contacts could be too small to derive value from your product or service or they might be from an unrelated industry using your product in an unconventional way. You want to avoid encouraging these types of leads. In fact, it's preferable if this type of person doesn't resurface in your lead scoring process again.

Junk: These are the kind of leads that ideally shouldn't even be in your database. These might include students who bypassed your EDU filter, people pitching or selling you something, or job applicants who somehow made it through your scoring model. These are not just irrelevant but detrimental to your lead scoring system. They need to be removed from your database.

The feedback loop's purpose is to classify every lead into one of these four quadrants: Good Lead, Bad Timing, Bad Fit, or Junk. This categorization gives you quantitative data to assess whether or not you're providing high-quality leads to the sales team.

Using the feedback received, you can then adjust your marketing activities or change the way you're scoring to improve the quality of the leads you're sending to sales. The feedback loop is thus a continuous process of assessing the quality of leads, making necessary adjustments, and refining the lead scoring model for better results.

The goal is to ensure that leads fall into either the 'Good Lead' or 'Bad Timing' categories, represented by the left-hand side quadrants. Ideally, you want to focus on leads that fall into the 'Good Lead' quadrant - the right contact and company with an immediate need.

Implementing the feedback loop in Mautic

So, how do you go about getting this feedback from the sales team? The best approach is to ask directly.

Creating an MQL (Marketing Qualified Lead) alert campaign

To get started with the feedback loop process, we need to set up an MQL (Marketing Qualified Lead) alert campaign. This campaign will trigger when somebody crosses the threshold, meaning their score is higher than a predetermined limit. Once this happens, we alert the owner and add the lead to a static segment in Mautic, a segment we can call the 'Needs Feedback' segment. It's essential to note that this segment **doesn't have any filters**.

After adding the lead to the 'Needs Feedback' segment, they will enter a different campaign - the Feedback Campaign. In this campaign, we wait for three days before sending an email to the contact owner requesting feedback. After sending the email, we remove the lead from the campaign. If the salesperson doesn't provide feedback, the lead will remain in the 'Needs Feedback' segment. Then, the cycle starts again - the lead waits for three more days before the salesperson receives another feedback request. This process will loop indefinitely until the lead leaves the 'Needs Feedback' segment.

When the salesperson clicks on one of the four buttons in the feedback email (representing the four feedback categories: Good Lead, Bad Timing, Bad Fit, Junk), they will be directed to a simple landing page. This page will display a 'Thank You' message and an embedded form. The salesperson won't actually see this form, as it will automatically submit upon page load. Auto-submitting forms are crucial to ensure we collect good data without annoying the sales team by making them fill out forms. We can accomplish this auto-submission functionality with a script in an HTML area that waits for the document to load and then submits the form. In addition, the form should be set to 'kiosk mode' to prevent salespeople from being cookie-d as contacts.

The form will have two fields: the email address of the contact and the feedback category. As soon as the salesperson clicks the feedback button, the form auto-submits, and we're ready to process the feedback.

Creating the process feedback campaign

Now, we need to set up a 'Process Feedback' campaign. This campaign performs two crucial tasks. First, it removes the lead from the 'Needs Feedback' segment, thereby stopping the endless loop of feedback requests to the salesperson.

Secondly, it processes the feedback from the salesperson, determining how to handle the contact record in the database. If the feedback categorizes the lead as 'Junk,' we can safely delete that record to declutter our database. If it's a 'Bad Fit,' we know we don't want to market or sell to them. However, we can't guarantee they won't resurface. So, we set a 'Marketing Suspended' field for these leads. By doing this, we can suppress all email marketing or any paid media running off of CRM lists for them. Moreover, they won't reappear as an MQL.

By following these steps, we can streamline the feedback process, improve the quality of leads, and ensure that our marketing efforts are targeted at the right people.

Extracting the data from Mautic for analysis

The next step in our process is to extract the data from Mautic. We accomplish this by creating a couple of simple reports using the campaign events data source. The power of this approach is that it allows us to identify all the contacts we've sent as MQLs to the sales team. Moreover, we can set our date range based on when we alerted the sales team, giving us the ability to analyze the quality of leads we sent in a specific time frame, like March, for example.

Creating the summary report

The first report we're going to create is a summary report. It filters on campaign events that sent the MQL email. The structure of your campaign might vary, but in this case, it pulls from a specific campaign ID and looks for the 'send a user email' step in that campaign. We then group the data by the

MQL feedback field we collected earlier—this is the field that contains the categories 'good', 'bad fit', 'bad timing', or 'junk'. The report will then count the number of contacts in each category. If you want to include 'junk' leads in this report, you might consider setting a delay on your delete step in the feedback process. If not, you'll only get records for the other three categories which is fine if that's all you're interested in.

Creating the diagnostic report

The summary report provides an excellent high-level view, but sometimes you need a more in-depth analysis. For that, we create a second, more diagnostic report. This report doesn't group by category. Instead, it grabs the contact IDs and the feedback but filters on a specific feedback category. For example, if you want to investigate why certain leads were a 'bad fit', this report will give you all the contacts classified as such. You can then click on the contact IDs to examine their activities and figure out why they ended up being classified as a 'bad fit' MQL.

Analyzing your data

The final step involves regularly reviewing your diagnostic reports and looking for patterns. This is a manual process, and it relies on human pattern recognition. As you review, you'll start to see patterns. You'll identify how your model might be off from a scoring perspective, or how your marketing activities might be driving the wrong kind of traffic that ends up being classified as MQLs. This process of reviewing and learning from the reports is key to continuously improving your lead scoring model.

Question: When do we trigger the Marketing Qualified Lead (MQL) alert?

Determining when to send a Marketing Qualified Lead (MQL) alert can be a tricky process. In Mautic, you do have the ability to set a trigger within points

that will send an alert. However, experience suggests a more effective method. It's more efficient to create a campaign that features a segment that looks out for individuals who have points over a certain threshold.

Deciding on this threshold is up to your discretion. Some clients have set the threshold in the 200-300 point range. However, those thresholds can sometimes be too high and unnecessary. Starting somewhere between 50 and 100 points is a good place to begin, and you can adjust your activities accordingly.

Keep in mind, every business is different. The number of scoring activities and the volume of activity on your website will vary. Your starting threshold might need adjustments based on the MQLs you're generating. If you're not generating any MQLs or generating too many that turn out to be junk, you'll need to adjust your threshold. It's a science, really.

Now, there are certain activities which you won't find in the points section of Mautic. There are two ways you can tackle this. One way is to have a centralized model where you have a campaign that you call from other campaigns, which goes and adds those points. This approach keeps the point value in one place.

For instance, if you want to track engagement on a particular type of form, instead of having a points action, which would require maintaining a list of forms, it's easier to add a campaign action to either add the points to the campaign processing that particular activity or call a campaign to do so.

Another approach, which has been found to be quite useful, involves setting up Tag Manager to send a phantom page view to Mautic. This phantom page view is for a page that doesn't exist but has a unique URL. You can then pick that up through a regular points action, looking for that uniquely formatted URL. This method allows you to assign a point value to any activity that you can detect through Tag Manager on your website. It is incredibly handy for things like video views or clicks on the share button, which are great

engagement signals but hard to track in Mautic.

When Tag Manager is mentioned here, to clarify, we're referring to Google Tag Manager.

Question: Who should receive the Marketing Qualified Lead (MQL)?

Another important aspect is deciding who to send an MQL to. It's strongly recommend to use the lead owner or contact owner field within Mautic. If you're syncing to a Customer Relationship Management (CRM) system, this is often set on the CRM side. You can have a campaign inside of Mautic that looks for anyone who has a score above a threshold, or you can set the lead owner as part of the lead notification process.

These campaigns can get complex. Most CRMs have a built-in tool for assigning lead ownership based on certain criteria. A more efficient way might be to ensure that the record is synced over to the CRM before alerting. Then, check to make sure the lead owner is set, and send the alert to the lead owner. While the CRM is usually the easier place to assign lead ownership, it can also be done in Mautic if necessary.

Remember, the process of fine-tuning is ongoing. Over the years, you will have to refine this process to align with your business goals and needs. The key is to continually analyze, learn from the data, and make necessary adjustments to improve the quality of your leads.