

A PRIMER ON SERVICE LEVEL INDICATOR (SLI) METRICS



Keeping track of customer satisfaction is crucial for any company aiming to provide exceptional service now and even better service in the future. Service Level Indicator (SLI) metrics provide the key to evaluating the performance of a company and its services by creating a representation of the customer's experience.

Once SLIs are in place, they can help reduce ongoing system issues as well as drive quick and effective response to sudden outages. First, you'll need to understand what SLI metrics are and how they relate to Service Level Objectives (SLOs) and the Service Level Agreement (SLA). From there, you can take a look at the various types of SLIs and which ones typically provide optimal results.

The SLI, SLO, SLA Relationship

SLIs form the lowest level of the contractual SLA hierarchy within a company. That's not to say that they're the least important—on the contrary, they can be thought of as the foundation. The SLIs are key measurements of the performance of the system, while the SLOs are the objectives or goals that the company aims to achieve with regards to performance. The SLA is then the reigning authority, stating the consequences for not meeting the SLOs.

An SLI should be measured as a percentage such that 0% constitutes horrible (nonexistent) performance and 100% represents a perfect performance. This percentage is ultimately a representation of what your customers are experiencing. A higher percentage means that they are generally happier due to proper site functionality.

Some examples of common SLI metrics include:

- Latency
- Availability
- Error rate
- Throughput
- Response time
- Durability

The SLO is then the minimum target percentage that you wish to achieve for the SLI. For example, reaching 100% availability may be unreasonable, even if you are constantly striving for it. So the target level or SLO might be set to around 90%. But remember, the SLO is the level of customer satisfaction with the performance, meaning if a high SLO occurs for something like an error rate, then SLI corresponds to a low error rate. That is, a 1% error rate would result in an SLO at 99%, indicating that the customer is having a relatively error-free positive experience.

Optimizing SLI Contributions

SLIs serve various functions both within and outside of a company. Self-evaluation of SLI metrics allows you to determine the capabilities of your systems. Meanwhile, engineers from companies you work with may use the SLI data when determining the extent to which they depend on your services. The indicators are also useful for helping the larger organization make informed decisions regarding investment levels to balance reliable work against the velocity of engineering.

Choosing SLI Metrics

However, having too many SLI metrics can be overwhelming for engineers, preventing them from focusing on the most important performance indicators. Modern software platforms have hundreds or even thousands of unique components from databases and service nodes to message queues and load balancers.

Instead of trying to establish SLIs and SLOs for all of them, it's best to focus on system boundaries. System boundaries are the points where components expose capabilities to customers, such as a login point which exposes the capability for a user to authenticate their credentials and access a site. Focusing on the boundary will inherently capture the performance of the various components involved in exposing the capabilities to customers.

Customer Happiness Indicators

Because SLI metrics measure the customer experience, they should correlate with measurements of customer happiness. Optimally, a good SLI should rise when customers are happier, fall when they are less happy, and correlate to known outages. When the customers' disposition is unchanging and no outage exists, then the metrics should only oscillate within a narrow low-variance band.

There are a few measurements that you can use to help determine this correlation with customer happiness. Customer input, such as through support center calls, support forum posts, or maybe even mentions on social media, give a good indication of how happy or unhappy customers are at any given moment. These signals can even be used as SLI metrics, although it is not advisable

because signals which rely on human actions in this manner inevitably introduce lag. SLI metrics measured instantaneously via software avoid such lag as well as human errors in reporting, making them worth the investment.

However, the above happiness signals can still provide a good starting place for calibrating new SLIs. The process is as simple as lining up the SLI data with something like support calls in a spreadsheet to see if spikes in the SLI correlate to swells in customer complaints.

Calibrating SLI Data with Customer Happiness

As you compare SLI metrics with happiness indicators, you might notice certain patterns. Sometimes, the SLI data will show spikes and dips where there were no apparent changes in customer happiness. Other times, the customers will express discontent but the SLI shows no indications of significant changes.

There are several reasons that the SLI data may not align with customer happiness signals. If SLI data indicates unhappiness but customers appear satisfied, then the SLI data is likely “polluted.” For example, the system may be sending batches of multiple errors all at once rather than gradually over the time, making it appear more concentrated and dramatic than it really is. Alternatively, the system might lack a distinction between errors in the internal systems vs. errors on the customer’s side.

Be sure to focus on the errors that affect customer experience while excluding other errors from SLI analysis. This can be accomplished by creating category tags, conducting post-processing to remove batch queries, or measuring the SLI at a different place in the architecture.

On the other hand, if an event causes customers to be unhappy but SLI data doesn’t show any problems, then there is likely a gap in the SLI (assuming the measurement of customer happiness accurately reflects satisfaction). For the event in question, check for existing non-SLI monitoring signals that *did* show negative user impact. If possible, derive a new SLI metric from these signals. If not, look for other measurements you can make which would detect the problem.

Steps to Implement SLIs

Both SLIs and SLOs are vital for making informed decisions and changes to a company. Practically, any SLI is better than having no SLI, and getting started is a fairly straightforward task. SLI metrics align concrete numbers and goals with performance concerning customer satisfaction. A company’s first metric doesn’t need to be part of complex tasks like page oncallers or freeze releases—it just needs to focus the conversation on areas that require improvement.

The basic steps for selecting and implementing SLI metrics are:

1. Identify a useful system boundary on your platform.
2. Identify the associated customer-facing capabilities at the boundary.
3. Determine what it means for these capabilities to be available for the customer.
4. Use the definition of availability to define one or more SLI metrics.
5. Start measuring the SLI metrics to get a baseline performance percentage.
6. Based on the baseline performance, define an SLO for each capability.
 - a. Each logical instance of the platform should get its own SLO.
 - b. Multiple SLIs for a single capability should be combined into a single SLO for that

capability.

7. Track how the SLI performs against the SLO over time.
8. Track SLI correlation with customer happiness indicators.
9. Fine-tune the SLI until it matches both customer happiness, meeting the SLO.

The final and on-going step is to stay engaged—remembering that both SLIs and SLOs will likely change over time. Even if SLIs measured against happiness signals indicate that everything is running smoothly, you should continue conducting reviews every few months or so. You just might catch something slipping through the cracks.