

product or feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](#) (/products/#product-launch-stages).

Service Level Objectives (SLOs) are a core tool in the Google service monitoring toolkit. SLOs can give you a concise and low-noise signal as to the overall health of your services. Anthos Service Mesh lets you set SLOs for your services, and monitor and alert on your services in terms of those SLOs.

To monitor the health of a service, you need to understand which behaviors matter for that service and how to measure and evaluate those behaviors. A service level indicator (SLI) is a quantitative measure of some aspect of the service. Typical SLIs are:

- **Latency:** How long it takes to return a response to a request, usually measured in milliseconds (ms). Latency is typically presented as an aggregate. That is, the raw data is collected over a period of time and calculated as percentiles. The Service Mesh Dashboard displays a Latency graph on the **Metrics** page for each of your services. The Latency graph shows you the latency over time, which can help you determine a latency threshold or upper bound for a service.
- **Availability:** The fraction of the time that a service responds successfully. This is typically presented as a ratio of the number of successful responses over the total number of responses. The Error rate graph on the **Metrics** page can help you determine the availability of each service.

An SLO is a target value for a service level that is measured by an SLI. An SLO can be represented as: $SLI \leq \text{upper_bound}$ or $SLI \geq \text{lower_bound}$. SLOs are measurable goals for performance over a period of time. For example, you might have requirements like the following for some of your services:

- Latency can exceed 300ms in only 5 percent of the requests over a rolling 30-day period.
- The system must have 99% availability measured over a calendar week.

You can set and view SLOs for your services based on their telemetry data on the **Health** page of the Service Mesh Dashboard. You can then create alerts in [Stackdriver Monitoring](#) (/monitoring/alerts/) to warn you if a service isn't performing as expected.

- Learn more about SLOs from Site Reliability Engineering at Google:
 - [Site Reliability Engineering](https://landing.google.com/sre/sre-book/chapters/service-level-objectives/)
(https://landing.google.com/sre/sre-book/chapters/service-level-objectives/)
 - [The Site Reliability Workbook](https://landing.google.com/sre/workbook/chapters/implementing-slos/)
(https://landing.google.com/sre/workbook/chapters/implementing-slos/)
- [Designing SLOs](/service-mesh/docs/observability/design-slo) (/service-mesh/docs/observability/design-slo)
- [Creating SLOs](/service-mesh/docs/observability/create-slo) (/service-mesh/docs/observability/create-slo)
- [Monitoring SLOs](/service-mesh/docs/observability/monitor-slo) (/service-mesh/docs/observability/monitor-slo)