

What's in an SLA?

A service-level agreement (SLA) is simply a document describing the level of service expected by a customer from a supplier, laying out the metrics by which that service is measured, and the remedies or penalties, if any, should the agreed-upon levels not be achieved. Usually, SLAs are between companies and external suppliers, but they may also be between two departments within a company.

A telecom company's SLA, for example, may promise network availability of 99.999 percent (for the mathematically disinclined, that works out to about five and a quarter minutes of downtime per year, which, believe it or not, can still be too long for some businesses), and allow the customer to reduce their payment by a given percentage if that is not achieved, usually on a sliding scale based on the magnitude of the breach.

The SLA should not only include a description of the services to be provided and their expected service levels, but also metrics by which the services are measured, the duties and responsibilities of each party, and the remedies and/or penalties for breach.

Metrics should be designed so bad behavior by either party is not rewarded. For example, if a service level is breached because the client did not provide information in a timely manner, the supplier should not be penalized.

What Are Key Components of an SLA?

The SLA should include components in two areas: services and management.

Service elements include specifics of services provided (and what's excluded, if there's room for doubt), conditions of service availability, standards such as time window for each level of service (prime time and non-prime time may have different service levels, for example), responsibilities of each party, escalation procedures, and cost/service tradeoffs.

Management elements should include definitions of measurement standards and methods, reporting process, contents and frequency, a dispute resolution process, an indemnification clause protecting the customer from third-party litigation resulting from service level breaches (this should already be covered in the contract, however), and a mechanism for updating the agreement as required.

This last item is critical; service requirements and vendor capabilities change, so there must be a way to make sure the SLA is kept up-to-date.

What Kind of Metrics Should be Monitored?

Many items can be monitored as part of an SLA, but the scheme should be kept as simple as possible to avoid confusion and excessive cost on either side. In choosing metrics, examine your operation and decide what is most important. The more complex the monitoring (and associated remedy) scheme, the less likely it is to be effective, since no-one will have time to properly analyze the data. When in doubt, opt for ease of collection of metric data; automated systems are best, since it is unlikely that costly manual collection of metrics will be reliable.

Depending on the service, the types of metric to monitor may include:

Service availability: the amount of time the service is available for use. This may be measured by time slot, with, for example, 99.5 percent availability required between the hours of 8 am and 6 pm, and more or less availability specified during other times. E-commerce operations typically have extremely aggressive SLAs at all times; 99.999 percent uptime is a not uncommon requirement for a site that generates millions of dollars an hour.

Defect rates: Counts or percentages of errors in major deliverables. Production failures such as incomplete backups and restores, coding errors/rework, and missed deadlines may be included in this category.

Technical quality: in outsourced application development, measurement of technical quality by commercial analysis tools that examine factors such as program size and coding defects.

Security: In these hyper-regulated times, application and network security breaches can be costly. Measuring controllable security measures such as anti-virus updates and patching is key in proving all reasonable preventive measures were taken, in the event of an incident.

What about indemnification?

The SLA should include a provision in which the service provider agrees to indemnify the customer company for any breaches of its warranties. Indemnification means that the provider will have to pay the customer for any third-party litigation costs resulting from its breach of the warranties. If you use a standard SLA provided by the service provider, it is likely this provision will be absent; ask your in-house counsel to draft a simple provision to include it, although the service provider may want further negotiation of this point.

What should I consider when selecting metrics for my SLA?

Choose measurements that motivate the right behavior. The first goal of any metric is to motivate the appropriate behavior on behalf of the client and the service provider. Each side of the relationship will attempt to optimize its actions to meet the performance objectives defined by the metrics. First, focus on the behavior that you want to motivate. Then, test your metrics by putting yourself in the place of the other side. How would you optimize your performance? Does that optimization support the originally desired results?

Ensure that metrics reflect factors within the service provider's control. To motivate the right behavior, SLA metrics have to reflect factors within the outsourcer's control. A typical mistake is to penalize the service provider for delays caused by the client's lack of performance. For example, if the client provides change specifications for application code several weeks late, it is unfair and demotivating to hold the service provider to a prespecified delivery date. Making the SLA two-sided by measuring the client's performance on mutually dependent actions is a good way to focus on the intended results.

Choose measurements that are easily collected. Balance the power of a desired metric against its ease of collection. Ideally, the SLA metrics will be captured automatically, in the background, with minimal overhead, but this objective may not be possible for all desired metrics. When in doubt, compromise in favor of easy collection; no one is going to invest the effort to collect metrics manually.

Less is more. Despite the temptation to control as many factors as possible, avoid choosing an excessive number of metrics or metrics that produce a voluminous amount of data that no one will have time to analyze.

Set a proper baseline. Defining the right metrics is only half of the battle. To be useful, the metrics must be set to reasonable, attainable performance levels. Unless strong historical measurement data is available, be prepared to revisit and readjust the settings at a future date through a predefined process specified in the SLA.