

Outcomes-Focused SLAs

December 2006

In this article we discuss 8 key principles for preparing an outcomes-focused SLA, provide some tips on preparing better service descriptions and service levels (two of the key components of a SLA) and identify some of the common problems to avoid when preparing a SLA.

While the article is written primarily for a customer entering into a SLA, most of the points raised are readily applicable to the service provider.

8 Key Principles

Begin with the end in mind – sage advice for almost every endeavour. However, one often finds that the SLA sitting on their desk fails in this respect. It doesn't address the customer's business needs and may be difficult, if not impossible, to apply in practice.

"If only", you may lament, "they had spent time thinking about how this was going to work in the real world".

There are, of course, many reasons why a SLA ends up being ineffective. One critical reason is that the SLA was never prepared in a way that focused on the outcomes that the business really needed. Instead, there was a rush to pull out the favourite availability, response or restore service levels without standing back to think about what would make the SLA work in practice.

With this in mind, consider the following 8 points to help create a SLA that will deliver the "real world" outcomes that matter.

1. Develop the SLA from the top down

The starting point for developing the SLA is the business objectives that the service provider must meet, or at least contribute to. These will be business, and not technology, focused. For example, an objective may be ensuring that all communications are promptly and accurately sent and received (as opposed to merely ensuring that the servers and routers are available).

Fail to grasp these objectives and your services and service levels may well end up

being deficient. For example, the provider may be getting 100% on the service levels but abysmal results on the customer satisfaction survey. So, before you leap to the usual suspects such as service levels around availability, consider what will actually be meaningful for the business.

These objectives need to filter down to the service description and service levels, and are the criteria by which performance is assessed. What may only be a paragraph or a few bullet points explaining the business drivers will be an essential measuring stick in the development of the SLA.

Ideally these objectives would be included in the SLA to provide the business context for the deal and to assist with interpretation. However, service providers naturally resist the prospect of committing to business outcomes that are beyond their control. Consequently, careful drafting may be required to document a shared understanding of the business in a commercially and legally appropriate way.

2. Focus on what's important – the key services and service levels

It's possible to get drowned in too much detail. To be effective, it's important to focus on the key services and service levels. In general, it is better to opt for fewer, more targeted service levels rather than a huge variety of service levels that are unmanageable, distracting or in which poor performance can be hidden. 10 to 15 key service levels tend to be more effective in practice than 100.

This requires an examination of what really matters to the business, the risks at hand, and

which service elements are truly “mission critical”.

3. Concentrate on the “what” more than the “how”

SLAs should focus more on “what” services the service provider is to deliver than “how” those services are to be delivered. This means that it is not normally necessary to get down into a lot of technical detail. Instead, the focus should be on the outcomes that the service provider must achieve. As long as the outcome is achieved, does it matter how the service provider does it?

Sometimes it may be necessary to specify some of the “how” details to fit with your particular requirements. For example, it may be necessary to specify certain processes that the service provider must adopt so you can obtain consistency across your various providers. The challenge is to strike the appropriate balance.

Consider the following:

- The customer is paying for the service provider’s expertise and an approach that unduly focuses on the “how” does not let the service provider excel at what they do best. It can also end up in inefficiencies as the service provider is prevented from making use of, for example, its standard methodology.
- As changes in technology and circumstances are inevitable, there is a real risk that too much “how” detail will make for an irrelevant SLA down the track.
- The customer is unlikely to be the expert in the area and usually won’t have the information or expertise to properly determine how the services should be provided.

4. Think about the end-to-end service experience

In practice, the end-to-end services required by the business are often provided by a variety of service providers and even the business itself. It’s easy for something to fall between the cracks, or for “finger pointing” between service providers to undermine the service. Excellent performance from one provider of the overall service may be completely undermined by poor performance by another provider, or by sloppy handover between providers.

It’s too easy to address a SLA in a vacuum without considering what other SLAs or business processes it must integrate with to deliver the required end-to-end service. Consequently, considering business outcomes means considering the end-to-end performance that is experienced by the end user and considering what could prevent this from being “seamless”.

5. Treat the SLA as both a legal and an operational document

The SLA is both a legal agreement (or part of a legal agreement) and an operational document that will be used on a day-to-day basis. This means that for the SLA to be effective it needs to be both robust from a legal point of view as well as practical and usable. A SLA that is too long, too complex, or full of legalese, becomes unusable and counterproductive. An appropriate balance must be struck. This doesn’t mean that the legal issues don’t need to be covered off, just that they need to be addressed in a manner that is not going to compromise or complicate the day-to-day operations.

6. Consider the view from the other side of the fence

To get the outcomes you want, you need to walk a mile in the service provider’s shoes and be realistic about what it can and cannot do.

Service providers need to make a reasonable profit and manage the risks that they face. They will do what they can to optimise their performance and make a buck. If you’re too hard on the service provider or unrealistic, you can shoot yourself in the foot as the service provider becomes uncooperative, is unwilling to consider changes and does not act “in the spirit” of the agreement. The last thing you want is a service provider that’s “gaming” the SLA.

This means the customer needs to understand the vendor – what drives them, and what their objectives and limitations are. Consider the sort of behaviour that your desired SLA may encourage – it’s possible to get too good a deal.

7. Cater for change

Change is inevitable and, if not appropriately handled, can quickly make a SLA irrelevant and unworkable. Consequently, it’s important to anticipate the types of changes that are likely to occur and how those changes should be addressed.

This will usually require the inclusion of a change control process that provides a structured approach to requesting, considering the impact of, and approving changes. Also, it may be valuable to consider categories of “standard changes” that can be priced into the SLA (for example, if call volumes increase by X% an additional Y% is payable).

8. Lay the foundations for good governance and management

While a good SLA is a necessary component in achieving the desired business outcomes, it is by itself insufficient. No amount of nicely worded service levels are going to make up for a dysfunctional relationship. Without a relationship of trust and effective governance and management by both parties the engagement is likely to be on shaky ground.

Consequently, while you can’t contract for good relationships, it’s important to do what you can to lay a foundation of practices, processes and principles that will assist (and not hamper) an effective working relationship. Such practices and processes will typically include regular senior management meetings, regular and targeted reporting, and prompt escalation of problems to the appropriate levels of seniority.

Service Description

The service description section of the SLA sets out the services that the service provider must perform. The key principle here is that each service element must be described as clearly and concisely as possible. In short, if you haven’t specified it don’t expect to receive it – an oral agreement’s not worth the paper it’s written on.

The service descriptions must be clear, concise and consistent. Any ambiguities, inconsistencies or omissions can come back to bite you with disputes, service failures or the need to pay more money.

This is easier said than done. Significant work is required to identify each of the main categories of services that need to be provided and then each of the corresponding service elements. This will, in many cases, require an organisation to spend a period of months analysing its current operations to reveal its current and desired service requirements.

As noted above, the SLA needs to be a practical document and so it’s helpful to group the services into meaningful categories and make it clear as to who has to do what, when and where.

The possibility of diverging interpretations is high and so it pays to work through the service description with the provider and ask a number of “how will this work in practice” and “what if” type questions.

1. Clarity of scope

Clarity of scope is vital for both the customer and the service provider. The SLA must be very precise as to what is in and out of scope. Consider including matrices that specify for each process or task who is responsible for its completion, who must assist or be consulted with, and who must approve the outputs. Process diagrams are also a good way to clarify scope and avoid something “falling between the cracks”.

Clarity of scope becomes even more important if there are multiple service providers that all need to work together. If two providers are responsible for the same thing, the potential for “finger pointing” and disputes is high. Consequently, it’s necessary for there to be a consistent and integrated approach between the service descriptions of each of the service providers. Consider handover points, inputs and outputs, and the need to exchange information.

2. How much detail?

As noted above, the SLA needs to be useable and so too much detail can be counterproductive. It can also increase the chance of inconsistencies and conflicts throughout the document. In most cases simple and concise is better than complex and long. However, the nature of the service (including its importance to the customer) may mean that you need to go into some detail. If so, do everything you can to make it understandable and easy to apply.

You may also need to go into some detail to cover off specific processes and approaches. In such cases consider capturing this detail in a separate “Policies and Procedures Manual”. This can have contractual effect (by referencing it in the contract), but can also be more easily varied without going through a formal change control process (unless, of course, that needs to happen).

3. Consistency with the head agreement

As the SLA is often a schedule or component of a larger agreement, it needs to be developed with this in mind. This is particularly the case with both the service description and the service levels. The various documents need to mesh together in the way anticipated by the main agreement. For example:

- Make sure that any definitions that are used in the main terms and conditions are used consistently in the SLA.
- Read through the terms and conditions for any points that are supposed to be included in the SLA. Prepare a “shopping list” of points to cover off in the SLA.
- Remember the priority regime in the main agreement. This will specify that in the case of a conflict the SLA either will or will not take priority over the main terms and conditions. If you don’t draft the SLA with this priority regime in mind you may inadvertently vary some of the key contractual clauses (if the SLA takes priority) or find that some of the services and service levels are negated (if the main agreement takes priority).

4. Assumptions and exclusions

The service provider will rarely be able to assume complete responsibility for the services. Consequently, it will usually insist on there being exceptions in the SLA that clarify when it is not responsible. For example, it may be dependent on the customer or another service provider fulfilling certain responsibilities. The service provider will also require that it not be held responsible for performance in the case of a force majeure event.

Further, the service provider will often make a number of valid assumptions in setting out and pricing the services. For example, that the call volumes to the help desk will not exceed a certain number.

If you are the customer these assumptions and exclusions need to be treated with care. The prospect of the service provider “getting off the hook” in an unexpected way is high.

Often the consequences of an assumption proving to be invalid are left up for grabs. This is shaky ground. For example, if the call volumes do exceed the specified limit, does that mean that the service provider does not

have to provide any services to those calls above that limit? Should there be a more managed approach? Should there be a temporary measure in which the customer pays (at an agreed rate) for extra personnel on the helpdesk until either volumes subside or new volumes can be agreed? Should there just be a proportionate increase in the charges?

Work through the detail of each assumption or exclusion and consider questions such as:

- If they are excused from performing, to what extent are they excused?
- What limitations should be placed on the assumption or exclusion?
- When should the assumptions and exclusions not apply?
- What contingency plans need to be in place?
- What should the service provider do to minimise impact and manage the consequences of the exception or assumption?

Service Levels

Once the service descriptions are clear the next step is to consider which of the services are critical and/or need specificity in relation to the actual level of service to be provided. It’s these services that will need “service levels”.

The aim is to prepare a manageable set of service levels that adequately reveal good and bad performance in a meaningful and useful way.

It pays to remember here that one size does not fit all. Choose the wrong metrics (such as metrics that are difficult to enforce, or ones which encourage the wrong sort of behaviour) and you could have a disaster on your hands. Each service level needs to be tailored to the services at hand (so don’t rely too heavily on the precedent material).

1. Top down approach

As noted above, it’s helpful to take a top down approach when developing the service levels.

One text helpfully suggests that these high level questions be addressed:¹

- Who are the stakeholders?
- What do they care about? What does good service look like to the stakeholder?
- How will they know if they got what they wanted? (Identify the measurements that, if answered, will let the stakeholders know they got what they wanted.)

2. Best endeavours and reasonable endeavours

In some cases a service provider may resist agreeing to absolute or guaranteed service levels and instead push for a “reasonable endeavours” or “best endeavours” approach to meeting service levels. The problem with this approach is that it’s unclear what exercising best or reasonable endeavours will actually mean in practice. Consequently, to avoid uncertainty in relation to key service levels a customer can consider:

- pushing for absolute service levels but with more clear exclusions as to when the service provider is not required to meet the service level; or
- if it needs to stick with a reasonable or best endeavours approach inserting some minimum performance levels that make it clear when the service provider will be in material breach of the agreement (giving rise to remedies such as termination).

3. SMART service levels

Elizabeth Sparrow has helpfully suggested in her book “*Successful IT Outsourcing*” that the following SMART criteria can be used to assess each performance measure:²

- Specific: is the measure clear and focused to avoid misinterpretation?
- Measurable: can it be quantified and used in meaningful statistical analysis?
- Attainable: is the measure achievable, reasonable and credible?

¹ Sara Cullen and Leslie Wilcocks *Intelligent IT Outsourcing: Eight Building Blocks to Success* (Butterworth-Heinemann, Oxford, 2003) 79.

² Elizabeth Sparrow *Successful IT Outsourcing* (Springer-Verlag, London, 2003) 138.

- Relevant: does the measure reflect your organisation’s objectives and is it cost effective?
- Timely: is the measure collectable within a given framework?

As with the service description, the service levels must be unambiguous and precise. The problems almost always reside in the detail. There would rarely be a SLA that we encounter that would not have a myriad of problems when you start to scratch below the surface and work through if and how it could be applied in practice.

4. Objective and subjective measures

Given that the services levels are a contractual means of determining the service provider’s performance it is important that objective performance measures are used.

For example, when exactly would a call need to be answered if it had to be answered “in a timely manner”? It’s better to specify a specific time that can be objectively measured, such as “within 20 seconds”.

However, it is not always easy to obtain hard data in relation to the service provider’s performance. In some cases it may be appropriate to refer to subjective service levels, such as customer satisfaction. Naturally, service providers will be concerned that such surveys can be rigged. Some of these concerns can be alleviated by working through the questions that will be used for the survey, how the survey will be administered, how often the survey will be conducted etc.

The subjective nature of a service level may also mean that service providers will resist having any significant service credits tied to performance against such service levels.

5. Capturing the metrics

It’s important that the service levels are captured in a useable format that is readily understood by all involved. This will typically mean that a service level table is used to identify each of the critical elements of each service level. Set out below are some of the key elements to focus on

Specify the service level

State what the service level is, eg “99.99% network availability”. This can be a concise statement that is fleshed out by the use of definitions and formulas. For example, in the

service level above, what does “availability” actually include? What if the network is still working (i.e. “available”) but the performance is severely degraded; should that be considered available?

What if the customer has a number of sites; is availability measured per site, or across all sites? If it is measured across all sites you may end up with a scenario in which the service level may be met even though one site has no service at all.

Formulas

Some service levels will require a calculation to be performed to determine if the service level has been met. In the case of, for example, network availability, a common formula is:

Availability = the total number of minutes in the calendar month during which the [service] was available, divided by the total number minutes in the calendar month, multiplied by 100.

Alternatives to this formula will often carve out agreed maintenance periods from the calculation.

It’s often helpful to include an example at this point to demonstrate how the calculation should be performed.

Formula:	
Availability =	$\frac{\text{minutes [service] was available in calendar month}}{\text{minutes in calendar month}} \times 100$
Example:	
January Availability =	$\frac{44,500 \text{ minutes}}{44,640 \text{ minutes}} \times 100$ $= 99.69\%$

When does the service level apply?

Clearly state when the service level will apply (i.e. what is the “service window”). Does the service level apply 24/7 or only 9am to 5pm? If there is a service level requiring faults to be remedied within 4 hours, does a fault is logged at 4pm on Friday afternoon, need to be fixed by 8pm Friday, or 12pm Monday?

What is the measurement period?

Is the service level measured on a daily, monthly, quarterly or annual basis? If the

measurement period is too short it may not capture the results it needs to. Alternatively, long measurement periods can favour service providers because it means they can stuff up service in January but make up for it in February and March. Take availability as an example, if you measure 99.0% availability across one 24-hour day, the service provider is permitted 14 minutes of downtime before they fail the service level. If availability is measured across 12 months, the service provider is permitted over three-and-a-half days of downtime before they fail. For mission-critical services you would typically want to reduce the period of measurement.

Frequency of measurement?

You may need to state how often the service level must be measured. Will it be measured continuously, or every 5 minutes? Will each phone call to the helpdesk be measured, or just every tenth call? The approach to this facet of the service level will likely be dictated by the service provider’s systems. For example, if they can’t measure availability continuously then there’s little point requiring that in the service level.

How and where measured?

Include detail about how and where the service level will be measured. What will the data sources be? Will this measurement be done manually? Or using the service provider’s network monitoring system? Aim for getting some comfort as to the reliability of the measures.

In many cases the service provider will be responsible for monitoring their own performance and providing reports to the customer. In such cases, it’s helpful for the customer to be able to retain the right to be able to undertake periodic checks of the service provider’s performance and monitoring systems to check the accuracy and reliability of the data being collected.

Responsibility for measurement

Confirm in the SLA which party is responsible for measuring performance against the service level. Usually this will be the service provider, but it may be the customer.

6. Reporting

Regular reports on performance against the service levels should be provided. As this is where the “rubber hits the road”, it is well worth getting into some detail in the SLA as to

what reports will be provided and when. In many cases the service provider will be able to supply example reports in advance to show the type of reporting it is accustomed to providing.

The focus should be ensuring that the customer receives a comprehensive report of the performance of the service provider in a manner that is readily understood.

The reporting regime should usually require the service provider to use exception reporting to emphasise any missed service levels. Further, it can be helpful for the service provider to provide reports on the trends in relation to each service level. For example, a rolling 12-month report.

Consider also the frequency of reports (daily, weekly, monthly, quarterly, annually) and the use of colour coding in the reports to highlight serious problems (e.g. the traffic light system of red dots for missed levels, etc).

Common problems to avoid

1. Averages

Averages can be misleading. For example, if the service provider has committed to fixing faults within an average of 4 hours per fault. Make sure that they're not fixing 9 faults in 1 hour, and the tenth fault in 31 hours! Consider setting some upper and/or lower limits when averages are used. In this case, you might want to state that "faults will be fixed in an average of 4 hours, and in any case within 10 hours".

2. Adding up to 100%

If a service provider has committed to answering 80% of all phone calls to the helpdesk within 20 seconds, what's the commitment to the remaining 20% of calls? Consider using a 2-step service level in these situations, for example 80% of calls answered in 20 seconds, and 100% of calls answered within 40 seconds.

3. Failure to do the math

It helps to run a couple of scenarios through the formulas to make sure they work and meet your needs.

We often find that formulas include mistakes. Don't be fooled by percentages; 99% availability might seem great, but look at how much downtime that permits over a month, and how it's affected when the service window is less than 24/7.

Availability	Permitted downtime per month	
	24/7 service	9am–5pm service
99.999%	0.5 mins	0.2 mins
99.99%	4.5 mins	1.5 mins
99.9%	44.6 mins	14.9 mins
99%	7.4 hours	2.5 hours
98%	14.9 hours	5 hours

4. Lots of little failures

Lots of little failures can cause a customer endless grief but never quite trigger a service level failure. For example, there may be a number of small outages of a system that do not get caught by the availability service level but still have an adverse affect on the business.

The "mean time between failure" metric is a useful one in this circumstance. It basically sets out the average time allowed between failures. If this figure is set high, the service provider has less scope to permit lots of little failures without failing the service level.

5. Maintenance windows

Scheduled outages or maintenance is often carved out of service levels such as availability. If there's too many scheduled outages the business needs that lie behind the availability metric can be completely undermined. Ideally, all scheduled outages should be agreed with the customer in advance and only after sufficient notice. Also, if you're the customer consider whether you need service levels in relation to how much scheduled maintenance is permissible.

We welcome your feedback on this article and any enquiries in relation to its contents. This article is intended to provide a summary of the material covered and does not constitute legal advice. We can provide specialist legal advice on the full range of matters contained in this article.

Wigley & Company is a long established specialist law firm. Our focus includes IT, telecommunications, regulatory and competition law, procurement and media/marketing. With broad experience acting for suppliers and customers, government agencies and corporates, Wigley & Company understands the issues on “both sides of the fence”, and helps clients achieve win-win outcomes.

With a strong combination of commercial, legal, technical and strategic skills, Wigley & Company provides genuinely innovative and pragmatic solutions.

Wigley & Company, Barristers & Solicitors | E: info@wigleylaw.com | P: (04) 472 3023