White Paper
ITIL Reporting Pitfalls (pt 2): Metric Tunnel Vision (Or…The SLA Gap)

WP0187 | May 2015

Jason Dove

Jason Dove is an ISEB accredited Business Analyst, Developer and Professional Writer.

He consults for multiple leading businesses across various industries – from marketing to counter-terrorism.

Jason specialises in Business Intelligence related disciplines, with a strong emphasis on ITIL systems - a commonly overlooked opportunity for organisations to get the most from their IT investment.

With over 15 years of experience in the industry, Jason has leveraged his knowledge into that of author, blogger and is a contributor to print and online publications.

In the previous instalment we looked at how measuring elapsed time across reporting periods can cause heavily skewed results that can lead to the complete omission of key activities.

This is, by far, the single most damaging ITIL reporting gotcha across the board, but the subject of this instalment can be a close second for certain organizations.

Which organizations?

The big ones with large, sprawling ITIL systems. Those ones which can often provide an environment conducive to processes evolution beyond the originally mapped activities, or are so complex from the start that analysts fall back on standard SLA suites rather than fully investigating the real world scenarios.

For example: a Ticket getting bounced around multiple Teams before being assigned to the responsible Resolver Group due to increased enterprise complexity and redistribution of Service Support.

Another example: a Triage Team is inserted into the organization to deal with increased Ticket volume and/or diversity. This additional step, and the time it takes to carry out, can impact numerous SLAs.

A third example: KPI focus is not extended to incorporate the complexities of the true ITIL Service and only base metrics are used (usually 'Response', 'Resolution', 'Close' and occasionally 'Re-Open'). This is normally due to an inexperienced analyst implementing 'by-the-numbers' without actually understanding the processes or how to measure them accurately.

Once again, I am going to focus on Incidents as examples because Service Support is normally where this Pitfall manifests, but this can apply across the entire ITIL Service to varying degrees.

The Pitfall

At first glance, this does not look like a reporting issue and while this is strictly true, any KPI based report that is not aligned to the reality of the service is skewed….usually in a way that omits work being done and is detrimental to the Resolver Group.

This pitfall arises when ITIL reporting focuses purely on SLA metrics and does not include the full picture required for context (at the very least) or to quantify the real life process (ideally).

Continuous improvement being compromised may not seem a big deal as it is actually easy to miss in the short to midterm, however, there is a wider impact to the reporting accuracy.

As an example, we can assume a NETWORK Resolver Group will have a minimum of two SLAs, which run concurrently:

- Response Time

- Resolution Time

Side Point (Not Quite a Pitfall!)

The Resolution SLA measures the elapsed time from when the Stakeholder logs a Ticket. The Response SLA runs concurrently to the Resolution SLA.

From a Resolver Group perspective, this may seem fine (maybe even fair!):

Timeline

Resolution SLA   Resolution SLA
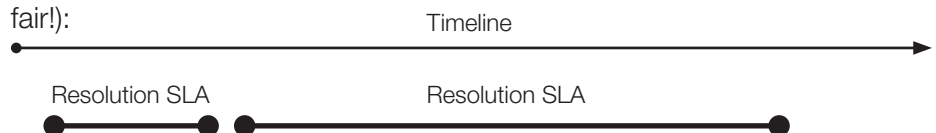
However, if the Response SLA overruns, the Resolution SLA is pushed back in its entirety. In theory, the Resolver Group could spend two months responding to a Priority 1 Ticket and then still meet the Resolution SLA:

Resolution SLA   Resolution SLA

Side Summary

SLAs are all about the service that the Stakeholder is experiencing, so the Resolution SLA is an indicator of how long the Stakeholder has been waiting. Whether or not it is 'fair' the available Resolution SLA time is reduced by the time spent waiting for the Response, when this should be irrelevant.

With the Response SLA and the Resolution SLA running concurrently, the Stakeholder has a definitive date/time when they can expect the Ticket to be resolved by:

Resolution SLA

Resolution SLA

And any overrun on the Response SLA only limits the time available for Resolving the Ticket:

Resolution SLA   Resolution SLA

Between these two SLA metrics we can identify which of the two of the Resolver Group's responsibilities has led to the SLAs being breached.

The difficulties do not start until the Resolver Group routinely fails one (or both) SLAs while working flat out, and is writing the same Report Commentary every month about tasks/obligations which are not covered by SLAs.

At this point, it becomes obvious that fine-tuning is required; or as ITIL 3 would call it: "Continuous Improvement".

If not, the SLAs will eventually be discounted. Which is understandable as, based on these SLAs, we can offer no guidance beyond "be quicker", which is pretty useless.

This is bad in its own right, but it also carries the risk of undermining the audience's faith in this report even though it is correct for the SLA related data it is showing.

The issue is the data it is NOT showing. Every Resolver Group deserves a true picture of their work.

If a Resolver Group is writing the same commentary to supplement their monthly SLA report(s), the SLAs and/or the reports themselves are sadly lacking to a point of hampering improvement within the organization.

Processes and Reality and Mitigation

This pitfall tends to originate in one of two ways (or two of two ways!):

- Poorly mapped processes
- Fixation on Metrics

Poorly Mapped Processes

Processes tend to be mapped as part of a new software implementation or creation/restructuring of resolver groups.
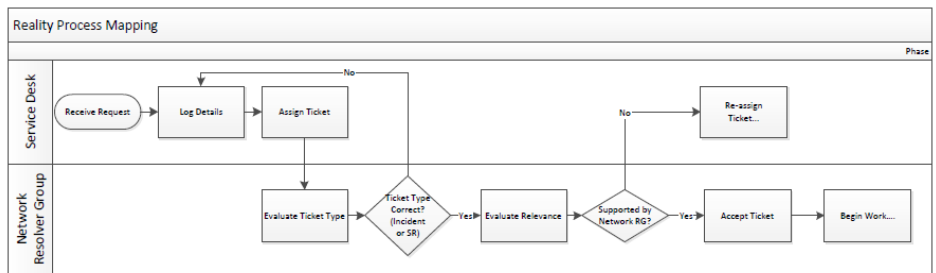
This process view tends to be very rose tinted and it is not until work

begins for real, with real world pressures, that the shortcomings become apparent.

It is the difference between this:

Simple Process Mapping

And this:
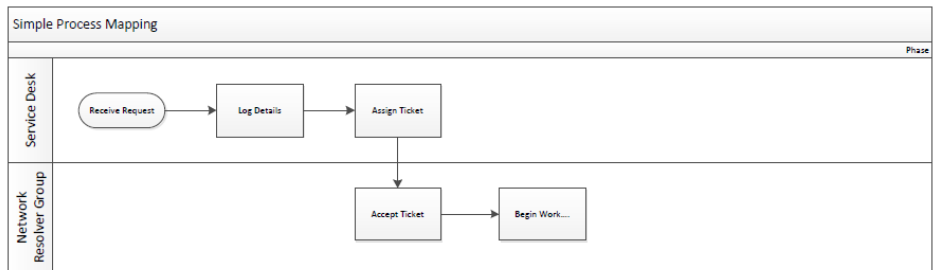
Reality Process Mapping

These are the same process map snippets, with the same Response SLA. The only difference is that the second mapping is more 'real world' and makes it clear that the Resolver Group can be working flat out on handling Ticket assignments with only a small percentage of said work being measured.

Note: Processing Mapping should be a key stage in Business Intelligence requirements gathering. Unfortunately, it seldom is for most BI implementations, and is particularly rare for ITIL reporting. This is mainly because BI is an afterthought, and who needs processes when you have metrics?

Fixation on Metrics

Metrics are important - they are the thermometer that shows the health of the organization's ITIL Service.

But like a thermometer, the supplementary information is equally important to avoid getting the wrong impression from the pure number being displayed (eg: a hot drink or bath).

One (poor) solution to this is to apply SLAs to everything. But with this approach the importance of the SLA (think KEY Performance Indicator) is soon diluted.

SLAs are important and should be important. They are the standards by which the provided Service is judged.

OLAs Do Not Fill SLA Gaps

SLAs and OLAs should never be presented in the same report. Never.

I appreciate that is hardly a new nugget of wisdom, but when gaps

appear it can seem that throwing in an OLA will solve the problem.

Don't!

Not least, because a well-developed set of OLAs can potentially fail, while the SLA is still met.

The service level agreed (SLA) with the stakeholder is all the stakeholder cares about. They don't need to know how it works under the hood and doing so can lead to a skewed view of the Service and general confusion.

The Solution

The solution for this issue is a lot fuzzier than with the previous pitfall in that it manifests itself in different ways specific to each organization and even each Resolver Group.

The good news is that the way to get to the solution is identical regardless of organization or Resolver Group:

*Accurate Process Mapping of a **Mature** ITIL System.*

'Mature' is emboldened because the difference between Process Mapping the 'dream scenario' for a proposed system is seldom the real life result once the system is established.

And that, for now, is what we are interested in. We need to know exactly what is happening in its entirety. The SLAs and KPIs will all still be in there somewhere and they will leap off the page, but the trick here is to treat them as being no more important than any other aspect of the process.

Indeed, it is these 'other things' that we have to focus on in order to fill in the blanks and get past the mind set of being purely KPI driven.

The illustration below is a standard SLA data grid report showing the two key targets for Response and Resolution.

Priority	Brought Forward	Open in Period	Resolved in Period	Outstanding	Response within SLA	Response Breach	Response %	Resolved within SLA	Resolved Breach	Resolved %
1	2	4	3	3	3	1	75.00	4	0	100.00
2	14	23	29	8	22	3	88.00	22	3	88.00
3	23	45	51	17	26	20	56.52	32	8	80.00
4	4	8	2	10	7	1	87.50	2	0	100.00

For pure SLA reporting there is no real value for more than half the columns on that report, this is all that is needed:

Priority	Response within SLA	Response Breach	Response %	Resolved within SLA	Resolved Breach	Resolved %
1	3	1	75.00	4	0	100.00
2	22	3	88.00	22	3	88.00
3	26	20	56.52	32	8	80.00
4	7	1	87.50	2	0	100.00

But the lack of context renders the report virtually useless. In fact, with this version of the report, the Network Resolver Group can meet their SLAs endlessly while providing a terrible level of support (by only

resolving tickets that will meet the SLA and forgetting about the rest).

Our original version showed that the Network Resolver Group was handling the ongoing workload with the 'Brought Forward' number being higher than the 'Outstanding' total.

However, this is entirely possible:

Priority	Brought Forward	Open in Period	Resolved in Period	Outstanding	Response within SLA	Response Breach	Response %	Resolved within SLA	Resolved Breach	Resolved %
1	12	8	3	17	3	1	75.00	4	0	100.00
2	22	43	29	36	22	3	88.00	22	3	88.00
3	23	45	32	36	26	20	56.52	32	8	80.00
4	6	8	2	12	7	1	87.50	2	0	100.00

Exactly the same SLA results, but it is obvious that the Network Resolver Group are drowning in work. They may be finishing the jobs they start in good time, but so much is falling by the wayside that the quality of service is undeniably suffering.

As an aside, ITIL Reporting that focuses heavily on SLAs will inspire the Resolver Groups to do the same. Visibility in a review report can be enough to mould behaviours in a positive manner.

So, the above report contains some pretty standard 'context columns' which should be applicable for most Resolver Groups.

Now, let's consider that last variation for the Network Resolver Group again. While it has the same columns as before, this example clearly shows that the Network Resolver Group is struggling with the workload once more.

But is it?

A Process Mapping exercise may well uncover that a large percentage of Network Tickets need to be passed to the Server Resolver Group, but this can only be identified by the Network Resolver Group after some initial analysis.

Note: one approach to solving this particular anomaly is to simply omit any tickets that the Network Resolver Group didn't Resolve. It sounds like an easy solution, but each one of these 'allocated-on' tickets still takes time and so impacts the support available for tickets they should be concentrating on.

With a column added to the report, the full picture starts to emerge and we can see that the Network Resolver Group spends most of its resources on triaging Tickets to the Server Resolver Group.

Possible Solutions

After using the above example, I feel duty bound to offer a potential solution or three on how to handle the issue of a Resolver Group taking on duties not covered by their SLAs:

1. Create a new Triage SLAs to set targets and recognise the work being done. This may require several SLAs to cover fully, e.g. target time, triage accuracy.

Once these measures are represented within the ITIL Report Suite the Resolver Group can be presented to the organization in an accurate light.

2. Up-skill the Service Desk and/or Triage Team. Whichever team supplies the Tickets to the Resolver Group, train them to identify Tickets in the same way the Resolver Group does. With this training in place, the normal Service Desk SLAs measure their performance while the Resolver Group gets back to the day job.

3. Make it someone else's problem! Taking the example above, if the vast majority of Tickets going to the Network Resolver Group are then triaged to the Server Resolver Group then it may make sense to push all the Network Tickets to the Server Resolver Group first and let them identify the few that are for the Network Resolver Group.

4. Split the Service Desk or Resolver Group to create a Triage team to handle these ambiguous Tickets. This will certainly be overkill in most circumstances, but if the volume of Tickets justify it, it is the most reliable solution (and will be absorbed into the Service Desk at a later point).

The Bigger Picture

There is a wider point here that spans this pitfall which could be considered as much a root cause as the 'Metric Tunnel Vision' I labelled it as earlier.

And that is: one size does not fit all!

Developing one report for all Resolver Group SLA monitoring may work, but it probably won't as each Group will tend to have its own unique provisions, working practices and processes. As such, they require their own unique reporting solutions to capture what they are doing and how well they are doing it.

Warning: The process is still king and all reporting should be process centric and this white paper is intended as a call to unique process reporting, not 'reporting whatever random behaviour' which is a huge pitfall in its own right.

orbus software