

DON'T LET YOUR SERVICE LEVELS "DO A NUMBER" ON YOU

This article first appeared in *FSO Magazine*, First Quarter 2008.

by Joshua B. Konvisser and Vipul N. Nishawala



Joshua B. Konvisser

Partner, Global Sourcing

+1.212.858.1027

joshua.konvisser@pillsburylaw.com



Vipul Natwarlal Nishawala

Partner, Global Sourcing

+1.212.858.1021

vipul.nishawala@pillsburylaw.com

Joshua B. Konvisser and Vipul N. Nishawala are New York-based partners in Pillsbury's Global Sourcing practice, advising clients on sourcing strategies and transactions spanning information technology, business process outsourcing, technology transfer, and related commercial matters.

Businesses outsource work for reasons ranging from cost savings, to exiting non-core businesses, to converting fixed assets into variable costs. Regardless of motivation, outsourcing is characterized by a customer transferring certain performance responsibilities to a third-party supplier.

In addition to transferring responsibility to perform the functions, effective outsourcing arrangements also transfer control over the means by which the supplier will perform the functions. This accompanying transfer makes sense: if the supplier is to bear the risk of performance, it also should control how it will achieve that performance. This transfer of control gives the supplier freedom to perform in a manner that is economically advantageous to both parties.

The supplier's freedom typically is not, however, absolute. Most outsourcing agreements specify certain rules that govern the supplier's behavior. For example, the supplier may be obligated to meet certain technology requirements, deadlines, and data security and privacy standards. Despite these rules, customers inherently relinquish a significant amount of day-to-day control over the supplier's performance. How then, does a customer

keep the supplier accountable for delivering day in and day out? The answer is service levels (also known as "key performance indicators" or KPIs) and service level credits.

A service level is a contractual commitment by the supplier that it will perform certain of the in-scope functions within a specified period of time (e.g., time to respond to a problem) or at a certain level of performance (e.g., minimum level of system availability). Effective service levels share a number of characteristics. They are objectively measurable, genuinely achievable, and measured on a periodic basis. They focus on those elements of service delivery over which the supplier has control. Most important, they reflect the customer's key business drivers with regard to the in-scope functions. In essence, the most meaningful service levels are tied to business outcomes rather than how a supplier performs the function.

But what is the supplier's incentive to meet the service levels? If the supplier fails to meet a service level, it is technically a breach of the contract. But the contract remedies available to the customer for breach (e.g., terminating the contract, suing for damages, transitioning functions back in-house, etc.) are usually

disproportionate to the failure and are hardly ever viable options. A better solution is to include a system of service level credits in the contract. A service level credit is specified, pre-determined amount the supplier must pay if it fails to meet a service level. Service level credits provide the customer with a contractual remedy by putting “teeth” into the service levels in a manner that does not upend the entire outsourcing relationship. The following hypothetical illustrates how service levels work when a customer transfers control of a function to a supplier.

Hypothetical: BIG Bank Outsources to DO-IT

Best In Gotham Bank (“BIG Bank”) is the largest retail bank in New York, with over 1,000 ATMs in Manhattan alone. BIG Bank has determined that running its ATM network is not a core function for the bank and has decided to outsource it to Distributed Operations and Infrastructure Technologies, Inc. (“DO-IT”), a supplier that specializes in managing and maintaining distributed systems. In order to achieve best practices and maximize cost savings, it does not make sense for BIG Bank to define with specificity how DO-IT will run BIG Bank’s ATM network (e.g., specifying exactly how many people DO-IT will have on staff to fix broken ATMs). DO-IT will be fully responsible for ensuring that the network is up and operational. The parties, however, are negotiating the service levels that BIG Bank will use to evaluate whether DO-IT is in fact properly meeting its responsibilities.

Setting the Right Service Levels

BIG Bank understands that service levels are the key to controlling DO-IT’s performance, so it sets out to establish the right metrics. After some consultation, BIG Bank determines that, within the financial services industry, the best-in-class network availability service level is “five 9s” availability—that is, DO-IT must operate the ATM network so that each ATM is available to BIG Bank’s customers 99.999% of the time each month. DO-IT objects vehemently to the proposed service level, arguing that that it is impossible to meet and that, even if it were, the price of the outsourcing would go up substantially. What happened?

We first need to understand what “five 9s” availability means. There are 2,592,000 seconds in a 30-day month. “Five 9s” availability, or .001% of permitted down time, amounts to just under 26 seconds of downtime for the entire month.

There are, no doubt, systems that cannot and should not be down for more than an extremely short period of time each month (e.g., an air traffic control system at a busy international airport). But it requires a significant investment to build a technical solution that meets such a high standard. For many systems, this may be technically impossible or cost-prohibitive. In particular, while a financial institution may require “five 9s” availability for electronic funds transfers (which are subject to strict Federal Reserve requirements), many of its other systems may not require this level of availability.

Back to the hypothetical: If BIG Bank requires each ATM to have “five 9s” availability, no individual ATM can be out of service for more than 26 seconds a month. Because ATMs at some point do go out of service, this is an impossible standard and DO-IT would be foolish to agree to it. Alternatively, “five 9s” availability at each bank location is a bit more feasible: if one ATM were out of service, the others in that location would be available. But this too is likely to be cost-prohibitive—a location with five ATMs would get only 130 seconds of permitted down time a month.

The key for BIG Bank is to have service levels that are tailored to meet its business needs. “Five 9s” availability across the board in this context is overkill. A more reasonable approach is for BIG Bank to have “five 9s” availability on the core backbone network and the systems that operate the ATM network as a whole, with a reduced standard for each ATM. This would prevent a major system-wide outage, but would allow individual ATMs or locations to go out of service for a more reasonable period of time. For example, if the service level for an individual ATM is 99.9% availability, the permitted downtime would 2592 seconds, or just over 43 minutes, each month—a more reasonable standard especially where a physical dispatch may be required to perform a repair. (Note that DO-IT may find even 99.9% availability too high for each ATM, but likely would accept this standard if it applied to all ATMs on average, with a longer permitted repair time for any individual ATM.)

Regardless of the service level, it is always important to "do the math" to compute what performance the service level actually requires and then to align that result with the business needs. Although there is often a temptation to seek the highest possible service level, there is almost always a higher cost associated with higher service levels. Customers therefore should strike a commercial balance between service levels and cost on an informed basis.

Setting the Right Service Level Credits

As described above, customers in outsourcing arrangements can manage the performance of their suppliers by setting the right service levels. Just as important are the service level credits that enforce the most critical performance standards of the arrangement. But service level credits are an effective management tool for the customer only if they are carefully calculated. As we will explore, they should not be too big or too small.

For the customer, there is often a temptation to seek a service level credit that will "make the customer whole" following a missed service level—effectively allowing the customer to recover damages resulting from the contractual breach. The primary goal of service level credits, however, should not be about recovering damages, getting money out of the supplier, or "punishing" the supplier. Rather, service level credits should be designed to give the supplier a

financial incentive to perform to the agreed upon performance standard. In order to achieve this, a service level credit cannot be so big that it would put the supplier out of business or eliminate the supplier's profit margin. On the other hand, a service level credit that is too small could create the wrong incentive if it is cheaper for the supplier to miss the service level than it is to meet it. If the service level credit does not cause the supplier some financial pain, it loses all value.

If we return to our hypothetical, let us assume that BIG Bank and DO-IT have agreed to a much more lenient "four 9s" availability standard (99.99%) for the core ATM network (i.e., 4.3 minutes' permitted downtime per month). Assume also the parties have agreed that, if DO-IT fails to meet this standard, DO-IT will pay BIG Bank a service level credit equal to BIG Bank's resulting losses. Under this scenario, both parties are setting themselves up for a fall.

If the core ATM network goes out of service for 30 minutes, all of BIG Bank's 1,000-plus ATMs in Manhattan would be unavailable and BIG Bank would incur huge losses due to lost transaction fees. If DO-IT has to issue BIG Bank a service level credit equal to these losses, DO-IT's revenue under the contract could be wiped out or severely diminished. DO-IT, in essence, may have to provide services to BIG Bank for "free" for a period of time. While this may seem appealing to BIG Bank, there is no way DO-IT can

enter into contracts that have it losing money—DO-IT cannot survive under that model and BIG Bank will ultimately suffer.

In practice, the amount that a customer can recover through service level credits each month is commonly capped as a percentage of monthly contract spend (i.e., the "at-risk pool"). While at-risk pools will vary from contract to contract based on the parties' respective negotiating leverage, they usually fall within the range of 5%-15% of monthly spend, and more commonly within the 8%-12% range. A 5%-15% hit would affect a supplier's profit margin, but would likely not eliminate it (which would put the supplier in a loss position). This is the "sweet spot" for effective supplier management.

But a healthy at-risk pool does not itself ensure that individual service level credits will be meaningful. This is because the parties usually divvy up the at-risk pool among individual service levels. The challenge for the customer is that if the at-risk pool is divided among too many service levels, the service level credit for any particular service level can become so small that it ceases to be a meaningful incentive for the supplier to meet the service level. More often than not, a customer is better off having fewer service levels with meaningful credits rather than many service levels with diluted credits. The key point, again, is that the customer should always do the math by calculating the absolute dollar value of each potential service level credit.

There are various strategies commonly used to magnify the at-risk pool so that it can be applied to more service levels and still remain meaningful. For example, there is often an application of weighting factors that allows allocation of multiples of the amount at risk, but which remains subject to the absolute at-risk cap in any monthly period. In addition, where there are a larger number of service levels at issue, the customer may negotiate for the ability to shift allocations of the at-risk pool among service levels during the term so that the service level credits can be reallocated to the "pain points" that may develop during the outsourcing. Further discussion of these strategies is beyond the scope of this article, but they should be considered whenever the math is not working.

Going into a service level negotiation, the customer often has a bias to ask for the most and the best service levels. Paradoxically, this frequently diminishes the effectiveness of sourcing arrangements by placing both customers and suppliers at an immediate disadvantage, while overlooking the power of financial incentives behind superior service. By doing the math and calculating the true impact of each requested service level and appropriate corresponding service level credits, the customer can ensure that it has optimized the commercial benefits of the outsourcing relationship.

Pillsbury Winthrop Shaw Pittman LLP
1540 Broadway | New York, NY 10036 | 1.877.323.4171.

© 2008 Pillsbury Winthrop Shaw Pittman LLP.
All rights reserved.