

Notes on Managing Suppliers in a PBL Environment

Robert A. Butler
TFD Group

Introduction

Performance based agreements (PBAs), while simple in concept, have led to generally poor outcomes due to significant misunderstanding of how they should be implemented, both in contractual and operational terms. The confusion is largely cultural in that the entire PBL movement was only necessary because selling to the government had been conducted in so un-business like a way (cost reimbursable contracts)¹.

Managers on both sides of the negotiating table got to their positions of power and influence by being very good at cost reimbursable contracting. Everyone finds it difficult to learn new skills and we are all reluctant to abandon hard-won knowledge for uncharted territory. Those among us who have been very successful in learning and exploiting those skills are even more reluctant to abandon them. Hence, many errors have been made in the first wave of attempts to implement the fundamental ideas of performance-based logistic agreements.

In places like Japan, where the concept is quite new, there is an opportunity to avoid those errors and quickly gain the full benefits that can and should flow from this sound concept. But the key players must understand the concept and understand how mutual benefits – increased operational effectiveness at lower cost and increased profit with a more predictable revenue stream – can be guaranteed through intelligent application of the idea.

Cultural reluctance is not the only problem, of course. There are some technical issues that are of great importance. If, instead of buying a tangible commodity – say, spare parts – I wish to buy the service of immediate access to those parts, I have at least three formidable problems. First, how do I state clearly, unambiguously and objectively what service I wish to purchase? Second, how can I detect that the offer price of the service I desire is reasonable? Third, how will I be able to measure whether I have received what I paid for? The inability of managers trained in a cost-reimbursable world to answer the first question explains many of the failures during the first decade of attempts to use the PBL idea. The second question leads inevitably to an interesting paradox. Whereas, when governments buy things under cost-reimbursable contracts, forecasts of operating and support costs are both produced and consumed with a grain of salt, fixed-price or PBL contracts make accurate O&S cost forecasts crucial for both buyers and sellers. The third question is far more important than it is usually understood to be. Even when the owner of a system fleet is clear about his objective under PBL, he can make the mistake of choosing inappropriate metrics to measure it. This problem becomes even more complex when the prime sustainment integrator, the PSI, must work out how to manage the efforts of the group of suppliers on whom he depends to meet his obligations. This paper is a collection of notes prepared by the author from time to time addressing various aspects of the problems inherent in managing suppliers in the PBL setting².

¹ Those who had been good at the cost reimbursable business naturally became the decision makers by the time the PBL idea was introduced; they were good at the wrong thing. This is only worth noting because, in reality, the PBL idea is simple, straightforward and quite natural to business people in most commercial sectors.

² At this stage, this paper is specifically written about the role of avionics PSI for the new Japanese Maritime Self Defense Force maritime patrol aircraft, the P-1.

The Stages of a PBL Effort

There are roughly four stages through which a PBL contract, or PBA, will naturally evolve. The first stage is the decision, by both buyer and seller that a PBA is mutually advantageous.³ The second stage is for the prospective supplier to prepare a proposal and the buyer to evaluate the proposal. The proposal is not the contract, although one hears, from time to time of a buyer who actually makes it so. The third stage is the preparation of an agreement that accurately reflects the desires of both parties. The fourth stage is to execute the contract effectively (living up to the buyer's expectations) and efficiently (making a profit for the seller).

Stage 1: Business Case Analysis

Most owners/operators of system fleets – at least those working in government settings over the past several years – have either been encouraged or commanded to explore the use of performance based agreements to secure sustainment services for their fleets. In every case, these directives have been motivated by the desire to lower the cost of support services, increase their efficiency or both. Exploration of the efficacy of these arrangements invariably ends in the need for a business case analysis, or BCA. In short-hand, the BCA is an analytical comparison of the costs and benefits to each party of doing business (supporting the fleet) under different methods.

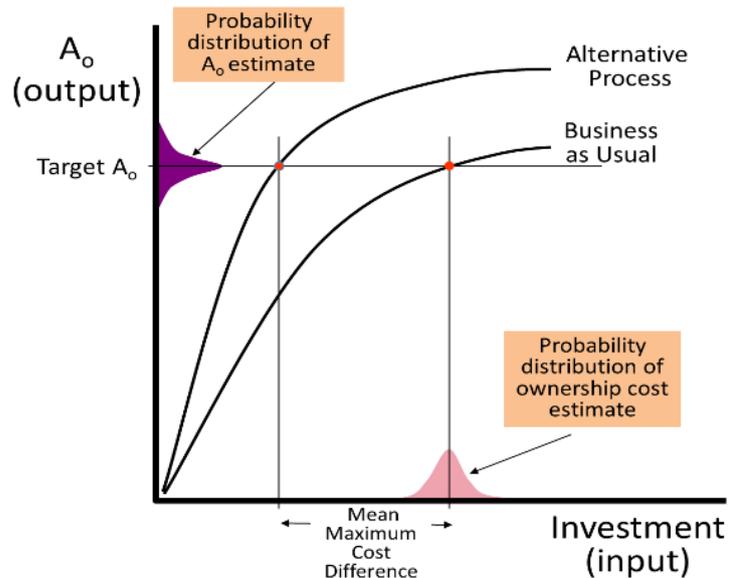


Figure 1: Alternative Production Function

Figure 1 illustrates several aspects of a BCA. The axes define inputs (investment) and outputs (achieved performance). The two curves are, effectively, production functions showing how inputs are converted into outputs. The underlying rationale for PBL is that an alternative production function can be reached, which lies above the current or "business as usual" production function. A modified form of life cycle cost analysis can help estimate a few of the points making up each of the production functions. No matter how good the method, however, there will be uncertainty about the exact value of cost and achieved performance at the intersections illustrated in the diagram. The two dimensions of uncertainty are illustrated by the probability distributions shown along the two axes of the diagram. Despite this uncertainty, it should always be true that a substantial potential savings is available at the mean value of the cost estimates, holding performance constant.⁴

The curves in Figure 1 indicate that continuing to do things the way they are currently done (business as usual or BAU) – is costly and holds little promise of real improvement. Changing the underlying approach by looking for an alternative, more efficient process, however, can move you to a higher production function. The point of the higher production function is that gains can be made for smaller resource expenditure and the absolute limits encountered at the asymptote of the BAU function are deferred to both higher cost and a higher potential support effectiveness.

³ In the insipid parlance of the day, a "win-win" situation. The very first course in economics will make it clear that no transaction is possible unless both parties are better off (i.e., win). Hence, if a proposed transaction is *not* a win-win situation, there is no rational reason to expect it to happen anyway.

⁴ Note that operational availability A_o is used as the example of the performance metric of measure of effectiveness (MOE) adopted by the customer. This is, unfortunately, not always (or even usually) true. The faint success achieved by most US PBL programs is partly due to the insistence of the government on using contentious MOEs in their PBL contracts. More on this later.

TFD White Paper

The existence of two dimensions of error in a diagram like Figure 1 can be illustrated as "spherical error," which is illustrated in Figure 2. The red ellipse can be visualized as a hill climbing off the paper, which, when viewed from any side, would have the same probability distribution shape as those shown in Figure 1. The ellipse

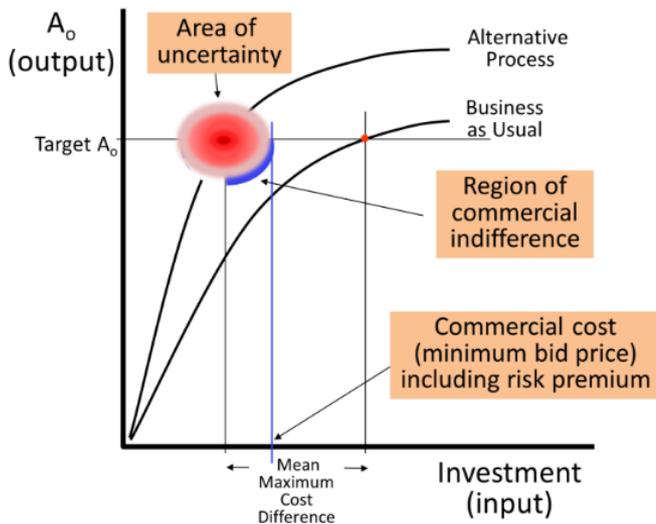


Figure 2: PBL Pricing Considerations

indicates the pricing decision problem faced by the candidate PSI. It is centered over the mean estimate of A_o and investment represented by the intersection of the optimal resource mix at the target value for support effectiveness. The mean values of A_o and cost are at the middle and the area remaining under the ellipse as you move away from the mean in any direction diminishes.

As with any probability density function, the mean is not a trustworthy indicator of what will actually happen. In fact, our sole reason for interest in the mean is to determine how far away from it we want to operate to feel secure in our choice. Therefore, we draw a vertical line at the right edge indicating the price at which the PSI candidate can be relatively sure he will not lose money in

delivering the target value of A_o . The distance to the right of the mean is a function of the variance of the distribution and how certain the candidate PSI wants to be that he will make a profit. The region in which he should be willing to operate is shown as the wider blue band in the lower-right quadrant at the outside of the uncertainty ellipse. Its curvature indicates he should be willing to trade away fee for a lower performance target. If the target performance is raised, the whole ellipse must slide up along the assumed production function, likely growing larger in doing so.

To make a reasonable estimate of costs, the candidate PSI must have some idea of what his costs will be as well as the costs his suppliers will pass on to him. To study this problem requires first that he determine how to task his suppliers so as to insure that they will provide him the support he needs to earn the incentive fees tied to the target MOE adopted by his customer. The most basic decision for the PSI will be whether to offer his suppliers PBL contracts that mirror the form of contract he anticipates between himself and the buyer. As with the prime contract, the central subject in a supplier performance based contract will be the metric used to determine the level of performance provided. If the prime contract is for operational availability, the appropriate pass-down metric is backorder hours per unit of time. If the prime contract is not for operational availability, more complex and far less efficient measures would be required. In view of this, we will assume throughout this paper that the PSI will be able to negotiate a real operational availability-based performance agreement.

To help suppliers carry out their own business case analyses, the PSI must first devise a backorder hour budget appropriate for the target value of A_o he anticipates. The specifics of this calculation and an example are presented below. Based on preliminary target values for backorder hour budgets to be offered to each supplier, the suppliers must then develop their own life cycle cost estimates, looking at alternative methods of doing business. The result will be that each supplier is either willing to sign up to a PBA or not. In all cases, having studied the costs of the preferred contracting method, each supplier should now be able to propose his own price estimates to the PSI and the PSI, in turn, will then have the inputs he needs to complete his own BCA.

Stage 2: Proposing A Performance Based Logistics Contract

Some of the hard work in developing a proposal has already been done in the business case analysis. The BCA equips the candidate PSI with estimates of the cost of performance – including costs passed up from suppliers – and the associated performance target. Unfortunately, this is only

the baseline estimate. A great deal of analytical agility will be required before the process comes to a close and buyer and seller can agree on the combination of fee and performance target.

The need for agility arises from the fact that the buyer will, as soon as he hears the first estimate of the fee required to achieve his performance goal, decide that a lower goal is acceptable and will ask for a new bid based on the lower goal. When he hears that price, he will ask for another based on either a higher or lower performance objective, depending on the results of the second iteration. These changes can go on for many, many iterations, even in a competitive situation in which the buyer has not seen any of the competitive proposals⁵. In a competitive situation, it is not the buyer changing his mind about the performance goal, it is the suppliers continually trying to improve their proposals before they are submitted.

The proposal will consist of both technical and pricing issues. We assume here that the reader knows far more about the technical issues addressed by the proposal than the author, so we address only the issue of proposal pricing, both for the suppliers and for the candidate PSI.

Managing Risk and Pricing Considerations

The analysis underlying the price/performance match can be illustrated by the graph in Figure 3. The horizontal axis tracks the incremental investment decisions in support resources required to deliver any level of support effectiveness. The vertical axis measures the monetary values of investment cost, the incentive fees earned as a result and the difference between costs and incentive revenue, which is the profit. Maximum profit occurs at the point where further incentive earnings at the margin can only be gained by investing more than what will be earned (i.e., where marginal revenue equals marginal cost). These graphs will look different for each supplier, depicting the inherent efficiencies of each supplier's ability to support his hardware as well as the effectiveness of his negotiation for backorder hour budget. The graph for the PSI will look different as well – deriving from the solution to his negotiations with each of the suppliers.

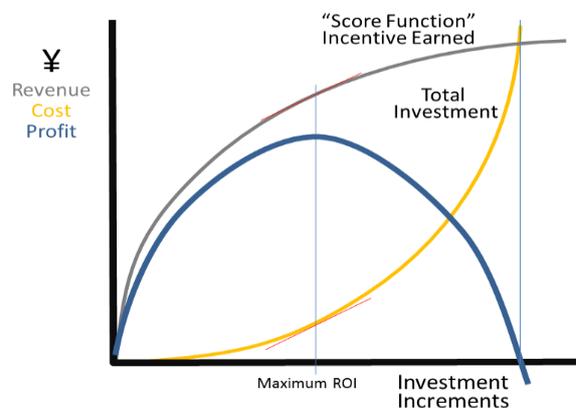


Figure 3: Return on PBL Investment

The situation depicted in Figure 3 is, however, the *ideal* situation. As with any investment, both up-side and down-side risks must be looked at and allowed to influence the investment. In the case of Performance Based Logistics, risk is often "asymmetric." That is, upside risk is limited and downside risk is always real. The asymmetry is illustrated in the Figure 4. While runs of luck will make even perfectly planned support investments under- and over-perform from time to time,

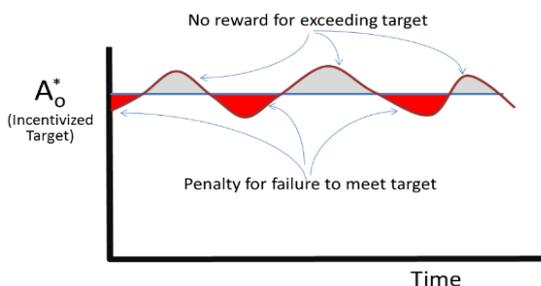


Figure 4: Asymmetric Incentive Risk

incentive fee contracts are generally written in such a way as to produce asymmetric risk. That is, penalties will be charged whenever support effectiveness falls below the agreed level (red areas in the graph), but over achievement (gray areas) is usually not rewarded. To counter this, a certain degree of over-investment may be warranted, raising the average achieved A_o above the level associated with theoretical maximum profit. An even better counter – at least from the viewpoint of the PSI – is to reach an agreement with the customer that rewards the

average performance delivered. That would remove the asymmetric risk completely and make over-investment in availability less costly. The fleet owner, however, would be making a mistake.

⁵ The author speaks from hard experience. On the LPD 17 proposal effort, for example, some 56 different life cycle cost estimates had to be produced during the last month of the proposal effort, for only four different ship systems.

TFD White Paper

The rationale for choosing a target rate of availability has to do with the size fleet he requires to accomplish some worthy goal. This should mean that the utility he derives from a larger fleet rapidly diminishes.

An additional concern for the PSI is the possibility that his suppliers will, from time to time, fail to perform according to expectations, despite the dictates of profit maximization. Excessive backorder time from any single supplier could cause disproportionate loss of earnings from the system level incentive fee. Accordingly, asymmetric incentives must be passed down to the suppliers as well in an attempt to mitigate this effect.

These considerations lead us to recommend that a small over-investment be made above the optimal or profit maximizing level in support resources required to achieve the agreed-upon support effectiveness metrics. The degree to which the profit maximizing solution is exceeded will depend on a number of factors. These include the level of statistical confidence desired, the cost of increasing the score function beyond the profit maximizing point and the rate at which incentive revenue drops past the maximum point. The latter two points can be summarized as the steepness of decline in the profit function past the maximum point. Figure 5 shows the costs and returns to an overinvestment in achieving the maximum score.

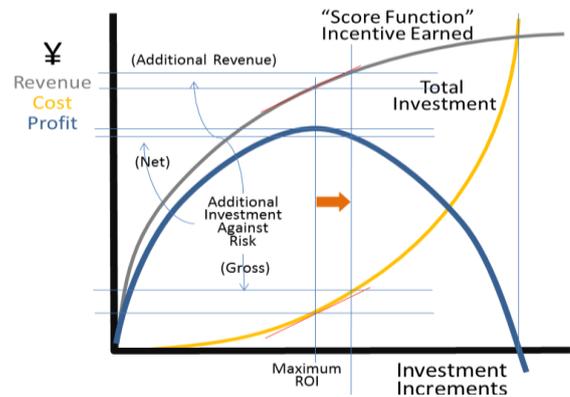


Figure 5: Hedge Against Risk

The additional investment to buffer against risk is shown with the solid orange arrow. The three pairs of horizontal lines show the difference, starting at the top pair, in a) additional revenue earned from incentives, b) the net cost of the risk-avoidance investment and c) the gross investment required. The net effect, in terms of Figure 4, would be to raise the achieved availability (the wavy line) enough to decrease the red areas to an acceptable level – at the expense of more gray, uncompensated, over-delivery of availability. The solution to this problem is a subjective one and can only be determined as a matter of business policy for the PSI.

Although ultimately, profit maximization is the goal, risk must be mitigated, future profits associated with the general happiness of the customer must be accounted for and finally, the PSI's business relationship with each of the suppliers must be taken into account. These are not analytical issues for the logistician, but the logistic analysis must be able to measure objective phenomena associated with those risks. In essence, the theoretical curves shown in the diagram above must be estimated on the basis of the character of the hardware system that will be supported, the manner of deployment and extent and pace of use, the architecture of the support system and the specifics of the incentive program to be agreed. If any of these attributes remain in flux (and usually *all* of them are still in flux) then the modeling process must be capable of making adjustments quickly and providing new answers as each component of the solution is subject to change.

Stage 3: Writing a Sound Performance Based Agreement

Contracts are prepared by lawyers. Lawyers, however, must be tasked correctly if we expect them to carry out their work correctly. Many performance based agreements written with the United States Department of Defense have had poor outcomes that were predetermined by the way in which the agreements were structured. The legal framework in which contracting is done is often freighted with a legacy of cost-reimbursable practice, which is not only inappropriate for PBAs, but toxic when forced into a performance based context. That is, those practices, adopted for a completely different form of contracting, must be abandoned and new practices adopted. In the case of the Japanese SDF, cost-reimbursable contracts have been the rule, but so has the use of trading company intermediaries. It is unclear if the latter represent an additional complication to the development of clean, useful performance based agreements.

TFD White Paper

What is completely clear is that the first principle of a good PBA is the adoption of a clear, simple performance metric. The only way to do that is to convince the buyer that he must state clearly what he wants – and *not* how to achieve it. The scenario that leads to an untenable contract goes something like this: the buyer wants more use of his systems (higher A_o) and asks for a target value of X. Rather than leave it up to the PSI to perform, however, the buyer then suggests that to achieve X, he will need a fill rate of Y and a time to fill of T, maximum delay times of such and such and so on. The end result will produce an impossible situation in which the buyer never gets what he wants and the PSI not only cannot make a profit, but can't produce anything like what he's been asked for and risks future business with that customer as a consequence of this failure. Another way to say this is to note that every additional incentivized target represents another constraint on the optimization problem faced by the supplier, removing more degrees of freedom from his solution to a complex problem.

Lest these remarks appear extreme, here is a case in point. Without naming the parties to the agreement, this example – a real one on which the author has worked – is for the supply support of an aircraft fleet in the United States, in which the USAF is the buyer and the seller was the original manufacturer of the aircraft.

The following analysis, portrayed in Figure 6, is based on a simulation using the exact scoring regime agreed to by buyer and seller for the aircraft in question. Rather than use data for the aircraft subject to this agreement, we substituted data from commercial aircraft that look similar (in the modeling sense). The problem also involved the use of other data elements that portray the way the real aircraft is deployed and operated.

Using this data construct, we developed an optimal inventory investment curve based on incrementally maximizing the reduction in expected backorders divided by the unit price of the part added to inventory (maximum bang-for-buck ratio).⁶ That curve is then translated into each of the elements that entered into the incentive score and plotted the scoring function as well.

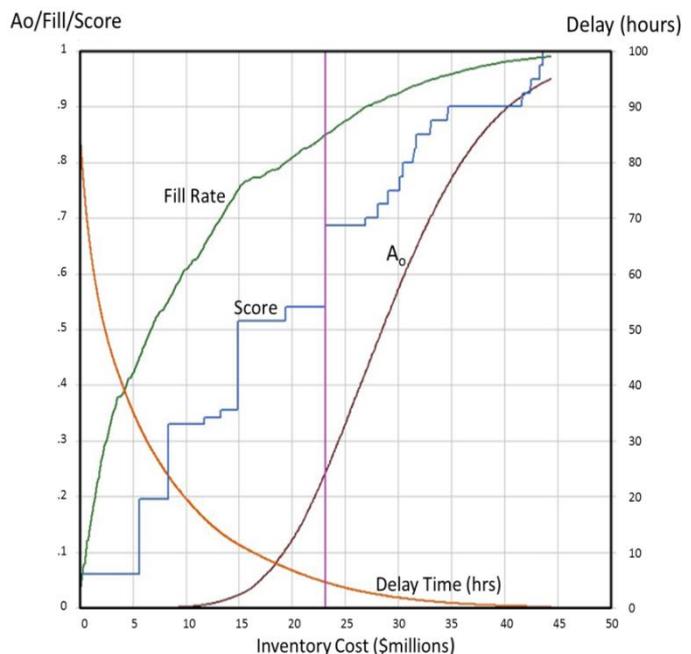


Figure 6: MOEs and Score Function for Spares

In this particular contract, fill rate, delay time and operational availability were all the subject of different levels of incentive fees. The various fee (and penalty) levels were set by the negotiators, who felt that they knew enough about supply chain management to set these levels in a sensible fashion.

The score function results from combining the amounts paid for achieving higher and higher levels of MOE values (including lower values of delay time). It is a step function because of the assignment of a giving fee level over ranges of performance achievement in each of the metrics. Its position on the graph indicates the relative value placed on each of the metrics in terms of incentive fees offered. It is useful to note that, if a single metric were selected and used by itself in the incentive program, then the score function would follow that metric's curve more or less exactly. There would still be steps because of the grouping of parts, but each award step would be coincident with that metric's graph.

The problem with the graphs in Figure 6 arises when you compare – as you must to determine the business case – the cost of inventory with the returns from incentives. In this case, the cost of

⁶ This analysis was based on the use of the VMetric spares optimization model, which uses the marginal optimization process developed by Craig Sherbrooke.

TFD White Paper

additional inventory exceeded the increase in incentive fees earned at about \$25 million (the vertical line intersecting the score function at about .53). The implications of this incentive regime can be read directly from the graph. At this optimal stopping point (the profit maximizing investment limit for the PSI, where marginal cost is equal to marginal revenue) the achieved operational availability would only be 25%.

But the contract stipulated that *the A_o target was 94%*. While all parties to the contract negotiation agreed that this was the correct and reasonable goal, none of them had the technical knowledge to understand that, in paying so much for an irrelevant target (fill rate) they used up precious budget money that should have been spent on the only target value they really cared about, fleet availability.

One might conclude that no harm is done – buyer and seller agreed on a scoring regime and an incentive program and the results were somewhat different than anticipated. So what? The trouble comes from the consequences. If the PSI actually delivers an A_o of only 25%, he will soon find out that the buyer really *did* want what he said he wanted – 94%. The buyer will, at a minimum, become frustrated and angry at the outcome, despite the fact that he shares responsibility for causing it. The buyer, who also reports to his own management, will soon discover that they too expected him to write a sensible contract that at least had some chance of delivering what they really wanted. No one will be happy and the chances that the PSI's company will be invited to bid for future PBL contracts is rather low.

The only expedient – and the one followed in the case in point – is for the PSI to put his own money into the additional support resources necessary to deliver a higher A_o. The rationale for doing this and foregoing significant profit is the potential for loss of profits in the future if the PSI fails to deliver performance closer to the level expected by the buyer.

The lesson suggested by this example is simple and straight-forward: extend every effort to *encourage the buyer to state clearly what he wants and not how to deliver it*. Almost invariably, the buyer really wants more use of the fleet for which he is buying support services. That translates, no matter how he measures that greater use, into some reduction in expected backorder time, which in turns translates into any one of a number of such measures, A_o, on-time departure rate, dispatch reliability and so on. As soon as additional measures of performance are introduced, the technical result is the departure of the score function from achievement of availability and the practical result is the loss by the PSI of degrees of freedom in the problem of how best to achieve the single important goal.

Determining the Backorder Hour Budget Offer to Suppliers

Assuming that the PSI has been successful in his pursuit of an availability-based performance contract, he must next turn his attention to how to insure that he can deliver that level of availability with reasonable confidence. He depends, in this, on the suppliers who will provide both repairs and replenishment. Typical practice has been to negotiate agreements with the suppliers that demand fill rate stock performance or repair turn-around-times. Neither of these practices is particularly useful in meeting an availability performance goal and both will surely cost more than necessary.

A word is in order about why neither fill rates, nor delay times should be used as metrics in a performance based agreement. Fill rates are measured at the base – where the systems are operated. Any stock held at a depot has no impact on fill rate. The result is that fill rate-based calculations ignore the obvious advantages of multi-echelon stock policies. This is especially damaging in the case of very expensive assemblies. More to the point, time is not accounted for in a fill rate calculation – only the probability of a stock-out at the base. The result is that a backorder that lasts for a year is given the same weight as a backorder that lasts for a day. It should be clear that the difference is profound, especially if a system is forced down waiting for the part.

Delay times, which are crucial to the limitation of backorder time are not useful subjects of independently stated metrics for a different reason. Although they are related to a useful measure (backorder hours), to stipulate a specific delay time is self-defeating and it unnecessarily removes a degree of freedom from the supplier's ability to meet the true objective. Some suppliers will find it

TFD White Paper

easy to reduce delay times far below any threshold value, while others may find it impossible to meet even what seems to be a reasonable target. In both cases, freedom to trade off delay time for investment in stocks, increased administrative efficiency or repair capacity and other variables is the essence of optimization.

An alternative is to negotiate performance-based agreements with suppliers that properly support the PSI's contract with the fleet operator. Just as with the PSI's contract with the fleet operator, how the supplier satisfies his target metric should be left up to him. In the case of a PSI faced with an operational availability target, the proper metric for his suppliers is a maximum of backorder hours per unit of time. The sum of these "backorder hour budgets" maps directly onto the value of achieved operational availability for the system.

What follows is an example of how a set of target backorder hour budgets can be determined as the starting place for negotiations with suppliers. The PSI must prepare to negotiate backorder hour budgets with each supplier that are consistent with the A_0 target for the system as a whole. To do this, the PSI must determine how big a budget to offer each supplier and he must also develop some idea of a reasonable price to be paid for that level of performance. The first datum – the budget offered to each supplier by the PSI – is relatively simple to establish.

To determine individual budgets, it is first necessary to see what the overall budget must be and then work out how each supplier's ship-set contributes to the reliability/maintainability of the system. Based on the different levels of exposure to expected backorder (EBO) risk represented by each supplier's ship-set, the rough budget can be set. To illustrate how this might be done, we start with a system whose overall attributes are as follows:

	A_0 Requirement	95%
	Fleet Size	100
	Fleet Hours per Month (720 hours per month)	72,000
	Required Available Fleet Hours per Month (95% of 72,000)	68,400
	Total Backorder Hours allowed per month	3,600

Each supplier's ship-set contributes a certain amount of failure exposure (expressed by the maintenance replacement rate per million hours or MRR6). Further, each supplier can be characterized as having a typical turn-around-time for the maintenance of his parts (or the re-supply of consumable parts in his ship-set). These values are shown in Table 1, which also works through the preliminary determination of an equitable budget offer to each supplier. The reliability of the entire group of supplier components shown in the table is an MTBF of 159.4 hours.

Supplier	TAT	MRR6	Raw Weight	Normalized Weight	Preliminary Budget
1	235	800	188,000	24.311%	875.2
2	100	70	70,000	9.052%	325.9
3	35	100	3,500	0.453%	16.3
4	200	,000	200,000	25.862%	931.0
5	45	200	9,000	1.164%	41.9
6	67	150	10,050	1.300%	46.8
7	92	600	55,200	7.138%	257.0
8	75	450	33,750	4.364%	157.1
9	88	100	8,800	1.138%	41.0
10	125	1,000	125,000	16.164%	581.9
11	97	200	19,400	2.509%	90.3
12	215	150	32,250	4.170%	150.1
13	10	600	6,000	0.776%	27.9
14	55	225	12,375	1.600%	57.6
Totals		6,275	773,325	100.000%	3,600.0

Table 1: Backorder Hour Budget Calculation

TFD White Paper

The preliminary monthly backorder hour budgets (the right-most column) would serve the purpose of guidelines in preliminary negotiations with suppliers. Note however, that 3,600 is the maximum number of backorder hours that provide the agreed availability rate (95% in this fictitious example). If exactly these values are negotiated, then all 14 suppliers will have had to satisfy exactly their requirement in every period in order to earn the profit-maximizing incentive. The hedge against risk depicted in Figure 5 would suggest that the end of the negotiation should produce a significantly smaller total of allowed backorder hours per month than would exactly satisfy the requirement. Moreover, target values for a negotiation might be best set considerably below the desired ending point, simply because negotiators for the suppliers might be very good at their trade.

Stage 4: Delivering Performance at a Profit

Having won the competition and signed the contract, the fourth stage is to fulfill your performance responsibilities at least cost, thereby maximizing profit. The first concern is to determine how much performance is the right amount, given all the considerations mentioned above. If the profit maximizing availability is 80%, you definitely don't want to deliver 90%⁷. No one will thank you and you will waste a large amount of the company's money. To insure that you are doing this correctly requires you to arm yourself with the correct data in a timely way and that you then process the data intelligently. Having done so, the results of the processing (i.e., the analyses) must be disseminated to the actors who will carry out the work.

Information – timely, pertinent and delivered to the right recipients – is essential to meeting the target and not wasting resources in the process. Sharing the information and analyses among team members, including not just those working in the PSI's company, but among all the employees of all the suppliers involved in the program, and often the fleet operator (the customer), is absolutely essential. In a competitive landscape this is one of the most difficult changes required in corporate culture, but it is necessary. Here's an illustration.

In normal practice, a supplier performing MRO services waits for a part to land on his receiving dock with an order for repair. Only then will he order the parts necessary to repair it, adding the delay time to receive the repair parts to his own administrative waiting time and the actual time required to repair the assembly. Each increase in the total delay, called the mean logistic delay time, requires increased investment in spare stocks to hold backorder hours (and, therefore, system A_0) constant. If the MRO supplier had full information, he would have seen the failed assembly long before it reached him and been able to gather the parts and other resources needed to fix it before it arrived. The reduction in delay time translates into both reduced investment in stock and more reliable delivery of service within the backorder hour budget agreed.

Complicating this picture is that, generally, retrograde shipping time is extremely long. In fact, often the field activity with the unserviceable component will hold it until an "economically large" shipment can be assembled to return to the depot. When MRO supplier and the fleet he is supporting are in different countries, the imposition of customs officials, ITAR regulations and other forms of official barriers to international flow of material have the dual effect of increasing both the mean delay time and the variance of the delay time distribution. The resulting increase in delay time and consequent increase in spare stock investment required to overcome these delays can be shocking. For example, a study conducted by TFD Group on behalf of the US Coast Guard (that did not involve foreign suppliers) found that 8% of the entire value of inventory carried for a single aircraft fleet was due solely to one repair contractor taking longer to complete repairs than the turn-around-time he had contracted for. An additional very significant amount of USCG inventory was being held because air stations were collecting unserviceable items in large shipping containers before returning them to the depot. Availability of the fleet increased a great deal and almost immediately when the head of supply sent an email to each air station with his FedEx account number and asked them to use it for immediate return of every unserviceable reparable item. It

⁷ We assume here that the target mentioned – and the one adopted by the PSI – reflects, not only the contracted target, but consideration of the discounted present value of increased probability of winning additional contracts in the future. When this effect is added to the tableau of incentive-bearing metrics, it can have a profound effect on the PSI's behavior. For example, if the profit maximizing stopping point for investment turned out to be 80%, adding the company's internal appreciation of the value of future business might lead to a *de facto* target value of 85% or more.

TFD White Paper

turned out that each air station was just saving a very small amount of shipping money in exchange for a staggering increase in inventory requirement.

Outline of a PSI-Delivered Information Service

There is a large community of people who will have a legitimate interest in the health of the fleet being supported by the PSI. A computer system that ties all of them together over a network (WAN or VPN) is an essential tool in maintaining adequate visibility of data and information across the community of interest. For a relatively small investment of time and effort an information service such as the one depicted in Figure 7 can be developed with appropriate security precautions such as controlled access for each user according to his rights and legitimate interests.

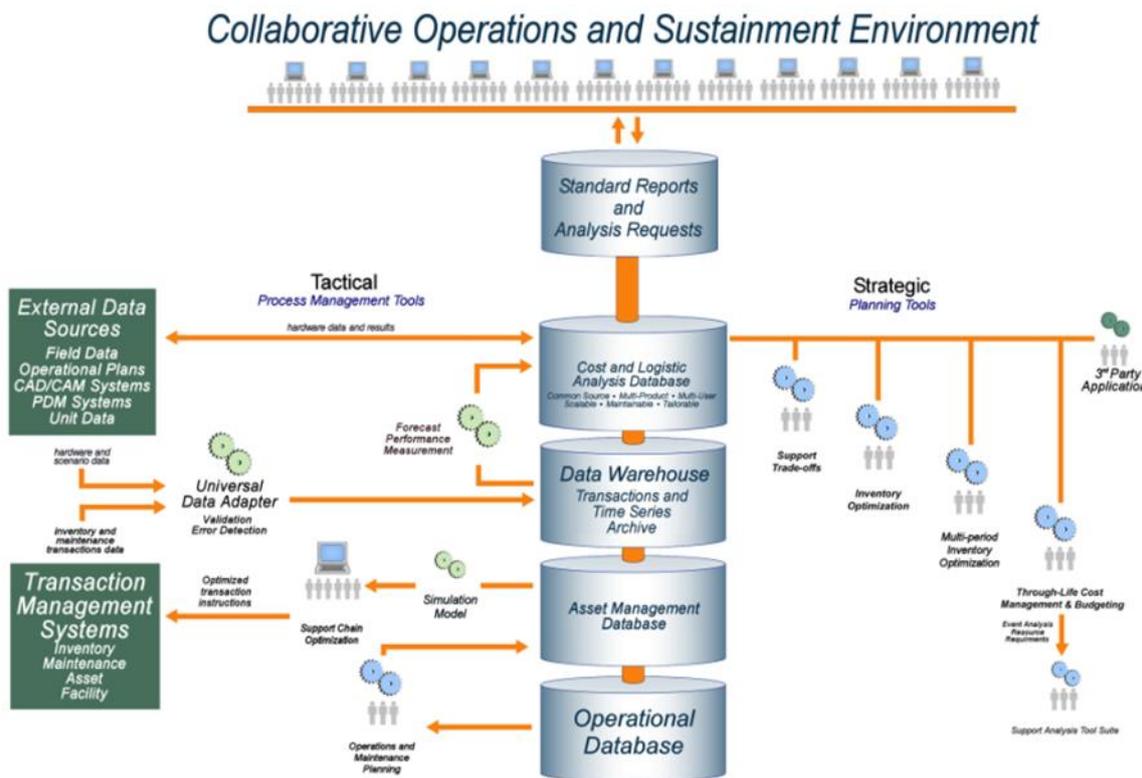


Figure 7: A Logistic Information Service Architecture

The information service should consist of the main elements of a logistic decision support system (LDSS) plus security and communication overhead functions. These include:

- Complete configuration information
- Asset visibility
 - Location
 - State
 - time due at next state/location
 - probability density function associated time due
- Fleet status
 - Achieved operational availability
 - NMCS
 - Hours of operation by tail number
 - Tail numbers in maintenance

TFD White Paper

- Type of maintenance action
- Crew status
- Operational plans
 - Changes in basing changes in operating pace
 - Anticipated deployments
 - Bedding down plans for new systems

In addition, analytical requirements should become part of the information service, including both standard analytical reports and special requests. Studies undertaken on request can also be transmitted to the requester over the service.

A simple architecture for the information system might look as shown in Figure 7. The end users of the service depicted above will include members from several groups, including operating personnel, program office, PSI staff and maintenance managers from all the suppliers. Data access control must be carefully determined and supervised by the data administrator to insure that proprietary information is protected and that security requirements of the operational groups are observed.

Meeting Target Metrics

The contract, we assume, has been written with an incentive fee tied to a performance measure that depends on cumulative backorder time per accounting period⁸. Generally, incentive fees increase up to an agreed maximum and penalties may be charged below some threshold. As a consequence, planning for support resource levels that can sustain operations meeting the performance requirement must account for the asymmetric risk discussed above, illustrated in Figure 4.

There is an extensive literature on risk management. Nothing in that literature, however, promises to eliminate risk. A properly equipped life cycle management team must deploy tools, not just to *manage* risk, but to *accurately manage the team's response to runs of luck* – both good and bad. While achieving the best possible starting point (the goal of strategic planning in logistics) is of great benefit, it is possibly even more crucial to execute the plan in such a way as to continually respond to a rapidly changing landscape of risk.

Consider spare parts. There are several models available to the logistician that are capable of computing optimal stock quantities and locations in a support system.⁹ Surely it is better to use such a model to equip the support structure than any non-optimal, that is, sub-optimal method. But spares optimization modeling comes with a different set of risks. Namely, that optimization in current models depends on "long-term, steady state" conditions. These are technical terms that are important to understand. First, long-term means forever – eternity. Next, steady state means that all the inputs to the model were not only true when you entered them in the model, but will remain true forever. That is a pretty heavy dose of assumptions. The argument for spares optimization and the caution about its shortcomings are illustrated in Figure 8.

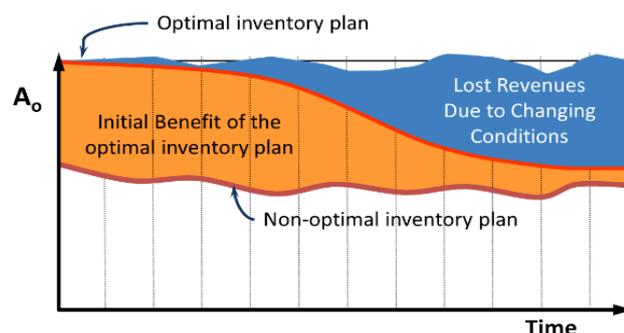


Figure 8: The "Stuff Happens Curve"

⁸ Note that there is a family of such metrics. The two most significant are operational availability and logistic delay time. There are several variants of both and all are based on expected backorders. Measures associated with fill rates do not depend on expected backorders or any other time concept.

⁹ By optimal we mean that no better solution (number or location of any unit of stock) is possible.

TFD White Paper

The starting point on the (horizontal) time axis shows the difference in operational availability achieved with an optimal versus a suboptimal plan. Both the lower and upper red curves represent the same investment in spares, but the upper curve achieves a higher A_0 because the selection and location of the spares is more advantageous. Because the optimal spares plan is based on the long-term, steady state assumptions mentioned above, however, its performance tends to decay over time. This happens for the simple reason that, as time passes, "stuff happens" and the assumptions (i.e., the input data) that drove the model, not all true in the first place, become even less so as the world moves on. New deployments, unnoticed failure propensities, vendor delay time changes, new basing and operating plans and a host of other commonplace, but significant changes all take their toll. Ultimately, if nothing is done, an optimal solution can be expected to perform no better than a conventional, non-optimal solution.

Three responses can be made to the "stuff happens" phenomenon. The first and most time-honored is to do nothing and allow the system to become as inefficient as it would have been if optimal modeling had never been used in the first place. The second is to periodically revisit the spares problem and "top-up" the spares package. This practice, most often followed in the management of fleets by sophisticated military operators, is very costly, both in terms of the cost of periodic new investments and in the sense of the lost benefits depicted by the blue region above¹⁰. The third response has only been put into practice in the past few years and only in a few places. We might call this dynamic optimization because it depends on a continually changing optimal analytical solution to the question, what should be done today? Optimized near-real time decision support for inventory management (tactical inventory optimization) is in its infancy, but shows great promise in allowing the PSI to keep the performance of his support system as close to the target level as possible – and at least cost.

Supplier Performance Over Time

The best supplier performance can only be called forth if each supplier is performing at or near his best potential. The PSI and the fleet operator should share the goal of insuring that the most efficient distribution of target values is found, thereby decreasing the total cost of the support package required to sustain the system. The problem is that the initial negotiation of backorder hour budgets introduced above will inevitably lead to target budget values that are only approximately correct. Some suppliers will be better negotiators than others and some suppliers will be victimized by their own lack of understanding of their responsibilities under the backorder hour budget regime. However, it is possible to move from an initial, imperfect state to a far better one by the use of two expedients. First is the creation of a market for backorder hour budget and the second is the introduction of a regime of continuous improvement. The first increases the likelihood that all suppliers will be able to meet their responsibilities in the initial periods. The second makes it possible to drive the costs of fleet maintenance to its theoretically lowest level by driving every supplier to achieve equivalence between marginal cost of driving down backorder hours to the marginal revenue earned from doing so.

A Market for Backorder Hour Budget Ceiling

Inevitably, some suppliers will have successfully negotiated larger budgets than they really require and others will have ended up with smaller budgets than they should have. That is, the budgets will be influenced by the subjective quality of negotiations, in addition to objective measures¹¹. The end result is bound to be suboptimal, taking the project as a whole; some suppliers will struggle while others have a free pass to behave inefficiently. This situation can be dealt with in a rational and sensible manner, by the introduction of a market for backorder hour budget among the suppliers. The existence of a market provides a place where those with excess can sell their hours until they reach a point where selling further budget ceiling would call forth greater costs (at the margin) to insure that they do not lose incentive fee. The buyer of budget hours will do so until the market

¹⁰ The periodicity of these re-optimization events is almost invariably tied to the budget cycle, rather than any period or metric having to do with optimal timing.

¹¹ The "excess" budget held by good negotiators is a form of economic rent. The following discussion indicates how the value of that rent can be capitalized by making a market in budget hours.

TFD White Paper

price of additional hours exceeds their value in either lowering his costs to comply with his target or the additional revenue represented by increased incentive fees. The price per budget hour will ultimately settle at a value that best reflects the marginal costs and benefits of meeting the budget constraints across all the suppliers in this market place. The essential ingredient of an efficient market is information. In this case, the PSI can provide information to the suppliers by acting as the clearing house for exchanges.

Continuous Improvement

While the initial inefficiency of budget distribution can be resolved by an open market exchange among suppliers, this is only a partial solution. All the market will have been able to do is overcome some of the differences in negotiating ability among the parties, turning economic rent into additional profit for those with the best negotiating skills. The process will not have called forth any *improvement* in the ability of suppliers to lower backorder hours. One speculates that the initial solution – the size of the backorder hour budgets and the resulting operational availability achieved – is associated with historical precedent set largely under a cost-reimbursable regime. There is no reason for anyone in the system – fleet owner, PSI or suppliers – to believe that the initial targets are efficient. One way to seek a more efficient solution is to arrange for continuous improvement. That is, lower the performance target incrementally over time.

Recall that due to asymmetric risk in the PSI's incentive program for the system as a whole, it made sense to over-invest in support infrastructure (spares and other resources). By doing so, the PSI would buffer himself against the penalties associated with runs of bad luck. As time passes and all the members of the PBL team become more familiar with the program gather more data to support more precise calculations and so on, it is feasible to obtain a return on this extra investment by a planned increase in delivered operational availability. The specific mechanism that allows this is the narrowing of the uncertainty ellipse depicted in Figure 2. The accretion of experience and knowledge about the program reduces uncertainty and makes it possible for the PSI and the suppliers to remove some of the risk premium from their pricing. So even though the price of the service is dropping, the rate of profit from it should be increasing.

With the agreement of the owner/operator of the fleet (one hopes), the goal of continuous improvement can be introduced for the incentive metric. An A_0 target, for example, that starts at 80% might be gradually increased to 90% over a period of several years. This is valuable for several reasons. First, an easy target to begin with makes all players more likely to embrace the PBL concept, seeing it as a lower risk proposition than it would be if very high results were required immediately. Second, all players learn more about the system and how to keep it running smoothly as time passes – the PBL buyer should be looking for a return on this improved knowledge. Finally, the ability to reduce backorder hours from the original budget will be distributed unevenly among the suppliers.

The last reason for continuous improvement brings us back to the value of a market for budget hours. Whereas the initial purpose of the market was to redistribute budget ceilings to overcome the effects of differences in negotiating ability, the continuing market allows those capable of improving their performance at lower cost to sell their increasingly excess budget to those who find it more costly to do so. Some suppliers will find that they can invest in additional repair capacity, for example, thereby decreasing delay time in repairs. Others will seek reduction in administrative waiting time or decide to invest in larger stocks of spare parts. Some firms will find that there is little they can do to improve their performance and will become the buyers of increased backorder hour budget from other market members. In any case, the existence of the market will induce each supplier to invest a profit-maximizing amount in reducing backorder hours as much as possible, rather than only to the level necessary to maximize the incentives he earns. Because he is able to sell the excess he can create, he is induced to improve his performance to the maximum extent.

Conclusion: Decision Support Infrastructure

By now it should be clear that the central management device in most of the initiatives discussed above is a robust Logistic Decision Support System, together with a means of communicating information, requests and results among the members of the system community. Taken together,

TFD White Paper

this is what we called a Logistic Information Service, illustrated in Figure 7. In all the stages of a PBL engagement – the business case analysis, the proposal, preparing the contracts and finally, execution – the ability to carry out analysis quickly and accurately is absolutely essential. One reason for there having been so little success in the PBL attempts undertaken to date (in the U.S.) has been the general absence of this ability in every phase of these efforts. The result has been firms awarded PBL contracts who are not capable of delivering what is required, a fact unknown to either the buyer or seller; PBL contracts written with incentive schemes that cannot possibly result in what the buyer actually wants; fleet owners purchasing PBL contracts at far higher prices than necessary, and so on.

All of these unfortunate outcomes can be avoided by the investment of sufficient time, effort and money in the early development of a decision support capability aimed at smooth management of performance based engagements. Along with development of the capability, of course, is the need to use it appropriately, namely, to share data and the results of analysis between buyer and seller; fleet owner, PSI and suppliers. The most frequently misunderstood aspect of the need for this sharing of data is how sensitive the data are. Generally, the data involved in logistic and life cycle cost analyses are rarely sensitive, with the exception of the price of individual parts. But part prices must be shared sooner or later in any case, or the parts won't be sold. Unlike a cost-reimbursable context, in a PBL environment, buyer and seller must share a consistent understanding of the pricing process in the context of operating and cost. Hiding data makes efficient contracts impossible.

While the development of an information service is not a particularly difficult matter, certain steps must be taken in a sufficiently timely way to allow the delivery of the service as soon as it is required. Since the first need for the service is during negotiation of contracts between PSI and suppliers on one side and the fleet owner on the other, life cycle cost analysis and support resource requirements analysis is of paramount importance. Other functions of the information service can be added as time goes on and the program comes closer to initial operating capability. We can think of the development of the information service as taking place in a series of steps. These steps are:

1. A model or set of models must be identified capable of fulfilling the role required – i.e. that produces the answers required and can be re-run quickly as various factors influencing the outcome are changed. TSL and TFD Group use the MAAP and VMetric models for this purpose.
2. The inputs required by the models must be identified and provision made to a) collect them from both the customer and the suppliers, b) manage the data in a secure fashion and c) maintain the currency of the data as the situation evolves. TSL and TFD Group use the TFD Data Vault for this purpose.
3. The analysis must then be carried out in as transparent and cooperative a way as possible between fleet operator, PSI and suppliers.

With this system in place, all parties will have the ability to make responsible decisions about the level of service that can be delivered and how much the service should cost. With an eye to the future, understanding that risk can only diminish as time passes, a program of continuous improvement will also make it a reasonable choice for the fleet owner to begin the program with modest, achievable goals, allowing elevated risk premia as the suppliers hedge against unknown factors. By the time the complete fleet is put into service, however, risks will have diminished, support will have become more efficient and the savings will have been passed on to the fleet owner in the form of lower costs and higher performance – and to the PSI and suppliers in the form of greater profit.

Robert A. Butler
TFD Group
rab@tdg.com