

Holding PBL to its Promise – the Analytical Challenge

TFD Global

Man has almost constant occasion for the help of his brethren, and it is in vain for him to expect it from their benevolence only. He will be more likely to prevail if he can interest their self-love in his favour, and show them that it is for their own advantage to do for him what he requires of them. Whoever offers to another a bargain of any kind, proposes to do this. Give me what I want, and you shall have this which you want, is the meaning of every such offer; and it is the manner that we obtain from one another the far greater part of those good offices which we stand in need of. It is not from the benevolence of the butcher, the brewer, or the baker that we expect our dinner, but from their regard to their own interest.

Adam Smith

An Inquiry into the Nature and Causes of the Wealth of Nations

Abstract

If Performance Based Logistics has an Achilles heel, it is that contractually-agreed systems of incentives may fail to induce support providers to invest enough in resources and take all the actions necessary to ensure support recipients get the things they need, where and when they need them. Incentive systems are subject to various potential failure causes, which include:

- performance metrics not reflecting operational needs closely enough,
- inappropriately weighted metrics,
- inappropriate scoring thresholds,
- inadequate analytical capability to enable satisfactory prediction of the readiness impacts of resource allocations, and
- too little money on the table.

A striking illustration of what can go wrong has been derived through use of an advanced readiness-based sparing (RBS) technique to evaluate the likely performance outcomes of applying the actual incentive provisions of an existing PBL contract to the spares support of a notional weapon system. Unsurprisingly, the lesson is that the level of investment at which a support provider's profit is maximized is unlikely to yield performance acceptable to the recipient unless the right amount is offered for attainment of the right target values of the right performance metrics.

The US DoD has approved a list of performance metrics considered supremely important to its warfighters. However, although a necessary condition for successful PBL, the mere existence of such a list is insufficient. It is also necessary to have a viable method, applicable across the spectrum of operational platforms and scenarios, for determining the relative importance of relevant metrics, specifying appropriate performance targets, and evaluating the quanta of incentives needed to induce attainment of the targets. For such a method to make a significant difference it needs to be accessible to program managers and prospective support providers

alike when engaged in early, high-leverage, context-setting activities such as business case analysis, bidding, source selection and contract negotiation.

The essential components of a viable method of this kind are plots of predicted values of performance metrics against levels of investment in support resources, optimally allocated. For the admittedly restrictive yet pervasive spares-only case, there are numerous RBS tools capable of generating suitable plots of this kind. And although multi-resource optimization is in its infancy, at least one COTS tool is now available to perform what might reasonably be termed Performance Based Resourcing (PBR).

The author has been engaged for several years in developing methods for interpreting RBS or PBR outputs not only to fine-tune PBL arrangements in the cut and thrust of negotiation but also to facilitate profitable fulfillment of ensuing contractual requirements. These methods have found expression in powerful new software tools already making their presence felt in significant PBL initiatives on both sides of the Atlantic.

A Good Idea...

Performance Based Logistics is a good idea. What makes it good is not just its focus on affordable performance outcomes, indispensable though this is. The real strength of PBL is its behaviorist stance: don't just rely on honorable intentions (even if you think you can safely trust everybody to have them) – also set up a system of sticks and carrots to make support providers unhappy when recipients are unhappy and happy when recipients are happy.

Yet good ideas can disappoint in practice. If the sticks are flimsy or the carrots have no crunch, then performance is optional from the support provider's point of view. Unless he stands to suffer consequences or reap benefits in proportion to the impact of his performance on the recipient's operational performance, a question mark must hang over his willingness and ability to play his proper part in sustaining capability.

Accordingly, performance measurement and economic incentives are central features of PBL. Support providers are scored against agreed sets of metrics, and their scores attract penalties or rewards in accordance with agreed formulae. Support outcomes below recipient expectations result in low scores, cutting returns to providers, in turn coercing increased resource flows. Support outcomes in excess of recipient requirements don't yield returns commensurate with costs, inducing resource economies. Mediated by self-interest in the spirit of Adam Smith, the cybernetic process to which these mechanisms give expression seems elegant and fail-safe. What could possibly go wrong?

Practical Cybernetics

There is presently a contract between the US DoD and a major Aerospace and Defense company for wholesale supply support to a smallish fleet of weapon platforms. It provides for an award fee in the form of a percentage of the value of spare parts supplied over the course of a 6-month performance period. A score is calculated each month and the monthly scores are averaged at the end of the performance period to determine the proportion of the maximum award fee due to the contractor.

The total monthly score is the weighted sum of scores achieved against nine separate performance metrics. Each of the nine scores is an integer in the range 0 to 4 identified by comparing the measured value of the applicable metric against a set of scoring thresholds associated with it. For example, the contractor scores 4 for the metric “Not Mission Capable, Supply” (NMCS) if its value for the month is less than 5.7%, 3 if the value is greater than 5.7% but less than 6%, and so on. The score falls to 0 once the value exceeds 8%. Table 1 is a representation of the scoring regime in its entirety. Since the weight factors sum to 100, the contractor can score a maximum of 400.

Weight Factor	Metric	Score				
		4	3	2	1	0
10	NMCS (%) ¹	< 5.7%	5.7 – 6%	6 – 7%	7 – 8%	> 8%
10	MICAPS (hrs/mth) ²	< 1000	1000 - 1300	1300 - 1600	1600 - 1900	> 1900
10	PMICAPS (hrs/mth)	< 15000	15000 - 18000	18000 - 21000	21000 - 24000	>24000
24	Response Time (Urg 1) ³ % < 36 hrs	100 – 85%	84 – 75%	74 – 55%	54 – 45%	< 45%
18	Response Time (Urg 2) % < 96 hrs	100 – 85%	84 – 75%	74 – 55%	54 – 45%	< 45%
12	Response Time (Urg 3) % < 60 days	100 – 85%	84 – 75%	74 – 55%	54 – 45%	< 45%
6	Response Time (Urg 4) % < 120 days	100 – 85%	84 – 75%	74 – 55%	54 – 45%	< 45%
5	Stockage Effectiveness (%) ⁴	> 90%	85 – 90%	80 – 85%	75 – 80%	< 75%
5	Issue Effectiveness (%)	> 80%	75 – 80%	70 – 75%	65 – 70%	< 65%

¹ The object of concern here is platform availability, often expressed in terms of a complementary measure, the Not Mission Capable (NMC) rate. To facilitate separate evaluation of sub-elements of the logistic system, NMC is conventionally subdivided into supply (NMCS) and maintenance (NMCM) components. Since in this case the contract scope is limited to supply support, NMCM is inapplicable.

² MICAPS stands for “Mission Capable, Supply” (and PMICAPS stands for “Partially MICAPS”). A MICAPS demand comes into existence at the moment a part required for the purpose of restoring a platform to mission capable status is requested. The score is based on the cumulative time (in hours) taken to satisfy all MICAPS demands over the performance period. PMICAPS is scored in an analogous way.

³ The four “Response Time” metrics apply to requisitions at different urgency levels (MICAPS demands at level 1, routine replenishment at level 4). A requisition at a given urgency level is deemed to have been satisfied if filled within the applicable time allowance (e.g. 60 days for level 3).

⁴ Stockage Effectiveness is the percentage of requisitions filled off the shelf (a.k.a. fill rate). The distinction between Stockage Effectiveness and Issue Effectiveness is that the former is measured in respect of

“stocked items” only, while the latter relates to all items.

Table 1 - Scoring Regime

Having signed up to this incentive scheme, the contractor sought a practical means of orchestrating the day-to-day actions of his team of Item Managers (IMs)⁵ to ensure customer satisfaction yet prevent erosion of profits. The TFD was engaged to provide assistance and the author was closely involved in development of a suite of support chain optimization software, chiefly intended to inform IMs as to the likely cost and performance impacts of choosing particular options over others available.

Success in this task requires, above all, a means of quantifying the change in overall score likely to result from choosing A, B or C over the do-nothing option. This, in turn, predicates an ability to predict the value each metric will take as a consequence of the

adoption of any given inventory state (i.e. the array of stock levels s_1, s_2, \dots, s_n across the item population) – accepting that one inventory state exists prior to an IM's exercise of choice and another will exist after it. Although this may seem problematic, because all the metrics in the negotiated scoring regime turn out to be simple functions of quantities typically output by Readiness Based Sparing (RBS) models⁶, prediction of scores associated with particular inventory states is well within the analytical state of the art.

The development agenda was based broadly on the following chain of logic:

- a given inventory state and a settled scoring regime together imply a definite score and hence a definite award fee value;
- a given inventory state also implies a definite inventory ownership cost;
- for any given score, at least one inventory state is optimal and can be identified through application of RBS-inspired techniques;
- hence the maximum achievable profit (taken to be award fee minus cost) can be calculated for any score;
- at some definite score the maximum achievable profit will itself be maximized;
- hence there exists an identifiable inventory state which maximizes profit; and so
- in principle, options available to IMs can be ranked in order of cost-effectiveness in steering the inventory from its current state towards the profit maximizing state.

⁵ A.k.a. Resource Managers or Supply Chain Managers. Each IM has responsibility for a manageable subset of the total inventory, typically parts in the same technology group (e.g. electrical components).

⁶ RBS models solve for optimal inventory states. An optimal inventory state is one that promises a targeted level of performance (usually operational availability - A_o) at the lowest cost. State of the art RBS implementations employ the VARI-METRIC algorithm, including various government models (e.g. AAM, ACIM, SESAME) and two notable COTS implementations (VMetric, OPUS).

In practice, the complexity of the IM guidance problem is considerably greater than these glib statements would suggest. One additional consideration is that the inputs to which an RBS procedure is most sensitive (apparent demand rate and apparent lead time) tend to change from one day to the next, making the optimal inventory state a shifting target. Another is that the very short scoring horizon is inconsistent with the long-term, steady-state premise of conventional RBS tools.


Be Careful What You Ask For – You May Get It

It seemed prudent to assess the topography of the scoring landscape from a strategic distance before settling on ways to steer IMs around bumps in the road. No new analytical tools were called for at this stage. The initial approach, admittedly laborious, was to employ suitable RBS software (VMetric) to identify the set of optimal inventory states for a suitably sized, deployed and operated fleet of representative platforms⁷, next use a spreadsheet to convert the measure of effectiveness (MOE) values associated with each state to a score⁸, and finally show how the profit function behaves by plotting award fee entitlements against the ownership costs (assumed equal to the inventory holding costs) that would have to be borne to earn them.

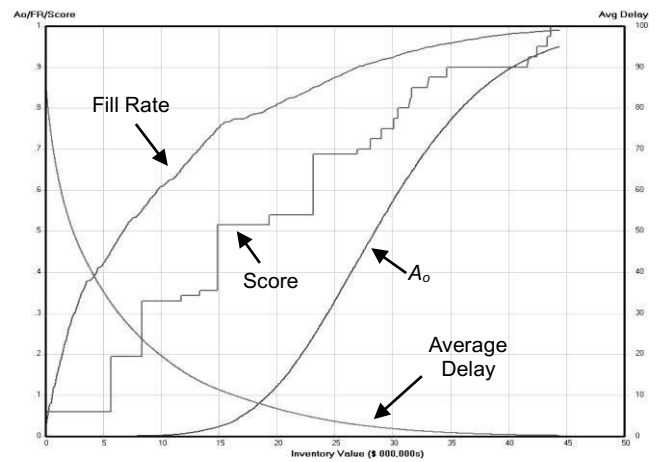
⁷ To avoid compromise of classified data, TFD “mocked up” a notional platform from several hundred line replaceable units in the bill of materials of a modern commercial airliner. This ensured the set of parts eligible for inclusion in the spares inventory would have price and demand rate statistics representative of high technology hardware. A fleet of 50 was assumed to be operating at a rate of 70 hours per platform per week from a single base, supported by an onbase warehouse.

⁸ Supply availability (A_s), fill rate and average delay MOEs provide all the data needed for a satisfactory prediction of the score a given inventory state is likely to achieve. For example, $NMCS = 1 - A_s$ and $MICAPS = \text{average delay time} \times \text{monthly frequency of events causing loss of mission capability}$. Note that because demand for parts is a stochastic phenomenon, measured achievement over a given period of performance is exceedingly unlikely to match the prediction precisely.

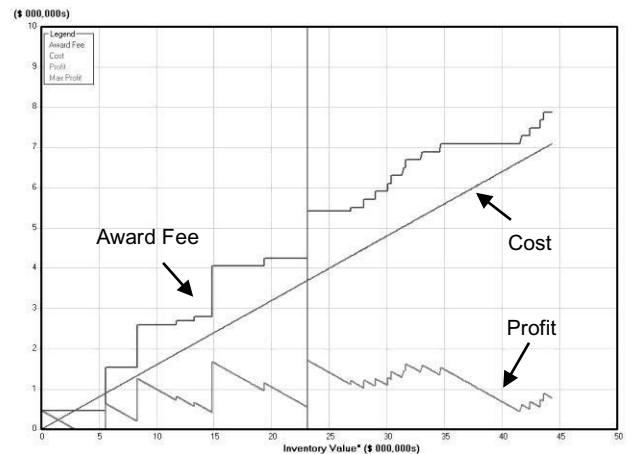
It would be counterproductive to divert too much attention to RBS concepts. However, readers unfamiliar with the technique will benefit from an appreciation of the characteristic shapes of curves produced by plotting the MOE values associated with optimal inventory states sequenced in ascending order of inventory value. The figure at right, for example, is a plot of A_o values produced by RBS analysis of the notional fleet described above. The sigmoid shape is entirely typical, showing the cost per increment of performance growing at an ever accelerating rate as operators demand higher and higher levels: a classic illustration of the “law of diminishing returns”. A curve of this kind is commonly called a *production function* (a plot of output levels against the inputs needed for their attainment). Curves of average delay time and fill rate versus inventory value are also correctly characterized as production functions, as indeed would be a plot of the expected score. This time, the figure at right shows all relevant RBS outputs for the case under consideration, plotted alongside the score derived from them through careful application of the provisions of the scoring regime. The stepped nature of the score function is a consequence of the small number of discrete score thresholds. Note that the score values have been converted, for convenience, from numbers in the range 0-400 to simple proportions of the maximum attainable.



Converting the score to an award fee is entirely straightforward. If we take the maximum fee to be 10% of the value of parts supplied, then (since the annual demand predicted for the notional fleet happens to be priced at about \$80m) the amount on the table is a little less than \$8m per year. Multiplying this by the score proportion gives the amount due to the support provider. Naturally the award fee plot has a shape reminiscent of the score function.



Equating the cost of attaining a given score to the inventory holding cost causes the cost function to be linear, with a slope determined by the holding cost rate. The figure at right portrays the financial outcomes of the modeled scenario, assuming a holding cost rate of 16%.

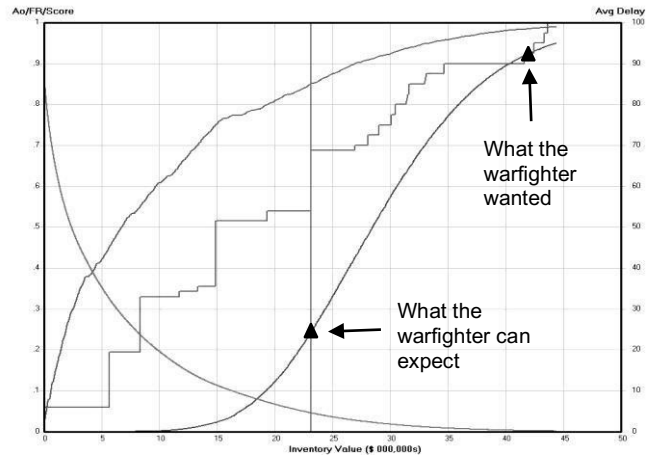


The profit function, of course, is obtained by subtracting cost from award fee values. Its resemblance to a mountain range (rather than a single peak) is inevitable, given the shape of the score function. Notwithstanding, there is a distinct profit maximizing point (marked by the

vertical line). The profit maximizing optimal inventory state evidently requires an investment of about \$23m.

When the same vertical line is superimposed on the production functions, a striking, if not entirely unexpected, phenomenon comes to light. Profit is maximized at a point where platform availability is about 24%. Yet the warfighter evidently wants more than 90% (judging from the NMCS thresholds he set).

There are significant implications in this glaring disconnect for design of a system to influence day-to-day activities. Even if the support provider, motivated by a somewhat more enlightened form of self-interest than Adam Smith strictly envisioned, is willing to shave his profit and slash his internal rate of return, how do you guide myopic IMs down from the highest peak, across the broken ground, to some far distant inventory state thought to represent a satisfactory compromise between recipient and provider concerns?



The answer is that you probably can't and shouldn't have to. And you wouldn't have to if the scoring regime and incentive provisions could be fine-tuned before contract signature to ensure congruence between performance and profit.

This realization set TFD on the path to developing a new strategic decision tool – a tool for testing PBL options ahead of committing to them. Before there can be a scoring regime against which to design a tactical decision support process there has to be a concept, a business case (persuasive both to support recipients and would-be providers), a bidding phase and a negotiation. The players in the early, high-leverage, context-setting stages don't have to be troubled by the volatility in demand or short decision horizons that will be the curse of participants in the support arrangements ultimately agreed to. Their responsibility is simply to frame rules and set threshold values that will ensure the game can have a mutually acceptable outcome in all contingencies⁹. But without some sort of prognostic device, few who play this formative role will be able to foresee with sufficient clarity the outcomes to which their settings will give rise.

What the Warfighter Really Wants

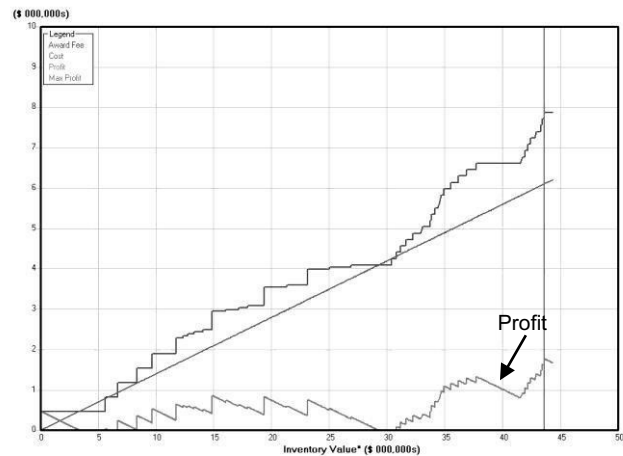
At first glance, the functional specification of a minimally satisfactory prognostic tool seems straightforward: it simply has to display, in near real time around the negotiating table, the optimal profit curve associated with any combination of performance metrics, weights, threshold

⁹ Analogous to the way carbon-based life forms are said to owe their existence to the very narrow limits within which the values of certain universal physical constants fall – e.g. “If...nuclear forces were slightly weaker than they are in our Universe, no complex nuclei could form at all. The entire Universe would be composed of hydrogen...”, from John Gribbin and Martin Rees, *Cosmic Coincidences, Dark Matter, Mankind and Anthropic Cosmology*, Bantam Books, 1989.

values, weapon system characteristics, operating and support scenarios and incentive pool size that any participant might see fit to propose.

The author demonstrated a prototype of such a device in March 2005 at a PBL-focused workshop presented to an audience from Defense, industry and academia. At one point discussion turned to ways and means of ensuring profit and warfighter satisfaction would be maximized together. Dialing in

increased weights for the first three metrics (NMCS, MICAPS and PMICAPS) and doubling the number of thresholds produced the result at right, establishing at the first attempt that the desired condition could certainly be achieved. But this drew an impassioned objection from a former colonel, who felt strongly that support providers should not be able to profit at all until quite close to meeting the warfighter's primary goals.



It is a trivial exercise to show that emergence into profitable territory can be postponed for as long as desired by suppressing performance metrics that facilitate early scoring. The expected values of all the Table 1 metrics other than NMCS, MICAPS and PMICAPS are slaved to the fill rate production function¹⁰, which happens in this case to increase much more rapidly than A_o . If, for example, the weights of these metrics are reset to zero (equivalent to taking them out of the mix altogether) then the score does not become non-zero until A_o has risen to about 0.45, and there is no profit until A_o exceeds 0.9. In fact, profit functions capable of resolving any likely objections can be assembled at will through selection of the right combinations of weights and thresholds.

And as we have seen, sometimes the appropriate weight is zero. That is, some performance metrics in common use may be out of place in the PBL environment. The US DoD has been keenly aware of this. After several years of committee work the Pentagon recently approved certain "Performance Based Criteria" for application to PBL contracts, stipulating that "PBL metrics should support these desired outcomes". The list¹¹ is as follows:

¹⁰ Take, for example, Response Time for the highest urgency requisitions, where the score will be at least 1 provided the delay time is less than 36 hours in not less than 45% of cases. Since the procurement lead time is highly likely to be more than 36 hours, the proportion of highest urgency requisitions filled in less than this time is identical to the proportion filled off the shelf – a.k.a. the fill rate.

¹¹ Taken from an August 2004 directive from the office of the Under Secretary of Defense for Acquisition Technology and Logistics (USD AT&L).

1. **Operational Availability.** The percent of time that a weapon system is available for a mission or ability to sustain operations tempo.
2. **Operational Reliability.** The measure of a weapon system in meeting mission success objectives (percent of objectives met, by weapon system). Depending on the weapon system, a mission objective would be a sortie, tour, launch, destination reached, capability, etc.
3. **Cost per Unit Usage.** The total operating costs divided by the appropriate unit of measurement for a given weapon system. Depending on weapon system, the measurement unit could be flight hour, steaming hour, launch, mile driven, etc.
4. **Logistics Footprint.** The government / contractor size or “presence” of logistics support required to deploy, sustain, and move a weapon system. Measurable

elements include inventory / equipment, personnel, facilities, transportation assets, and real estate.

5. **Logistics Response Time.** This is the period of time from logistics demand signal sent to satisfaction of that logistics demand. “Logistics Demand” refers to systems, components, or resources, including labor, required for weapon system logistics support.

While it may be wishful thinking to interpret this list as the deathknell of the fill-rate culture (see footnote 10 in respect of response time), its publication certainly seems to have coincided with a decisive shift in attention to strictly operational concerns such as availability and reliability. It is also worthy of note that the cost per unit usage criterion has little to do with performance – a would-be support provider has determined the outcome in this respect as soon as he has proposed a price per operating hour.

Delivering PBL Outcomes

If we can agree that timely exposure of the kind of information contained in the graphs shown above is likely to reduce the incidence of contracts carrying the seeds of their own failure, there still remains the thorny issue of PBL delivery. Nothing can alter the fact that outcome measures like availability and reliability lack the reassuring *tangibility* of resource inputs like parts, manhours and tools. In days gone by, the support provider simply sold the inputs, leaving the buyer to shoulder the risk that the input mix might fail to produce satisfactory operational performance, due to insufficiency, imbalance, poor application, runs of bad luck, or all of these things in combination. The essential change that PBL has wrought is an irrevocable shift of performance risk from buyer to seller. And along with performance risk come risks to both revenue and reputation.

Suppose, as is becoming fashionable, that the parties to a long-term PBL contract have settled on a single performance criterion, NMCS, and that the provider stands to receive the maximum incentive payment if $NMCS \leq x\%$ and no incentive payment if $NMCS > y\%$. Because both sides of the negotiating table were suitably informed by performancebased resource analysis, the provider expects to achieve the greatest differential between incentive and cost (i.e. profit) by delivering an NMCS outcome as close as possible to $x\%$. And, of course, the recipient is hopeful the provider will indeed maximize his profit, since any other position is tantamount to an admission that $x\%$ is not strictly necessary.

Since NMCS is a function solely of the average number of supported systems out of commission for lack of parts, the provider is guaranteed to achieve the desired performance outcome if he holds the average number of shortages of parts needed for system operation below a readily-calculated value¹². There are various well-understood ways to drive the number of shortages down:

- modify the system (e.g. by substituting more reliable parts) so that there will be fewer failures to deal with;
- improve maintainability (and hence the average repair duration);
- reduce delays in the support chain; and/or
- hold so many spare parts in suitable locations that shortages will seldom occur (or persist for long when they do).

Because each of these potential remedies carries its own cost implications, risks and time to take effect (none brings instant improvement), there is ample scope for confusion as to the appropriate emphasis. The task of harvesting profit over the term of the contract boils down inescapably to getting the proportions right and keeping them right in the face of ever-changing circumstances, a feat requiring highly flexible analytical capabilities spanning both tactical and strategic decision horizons.

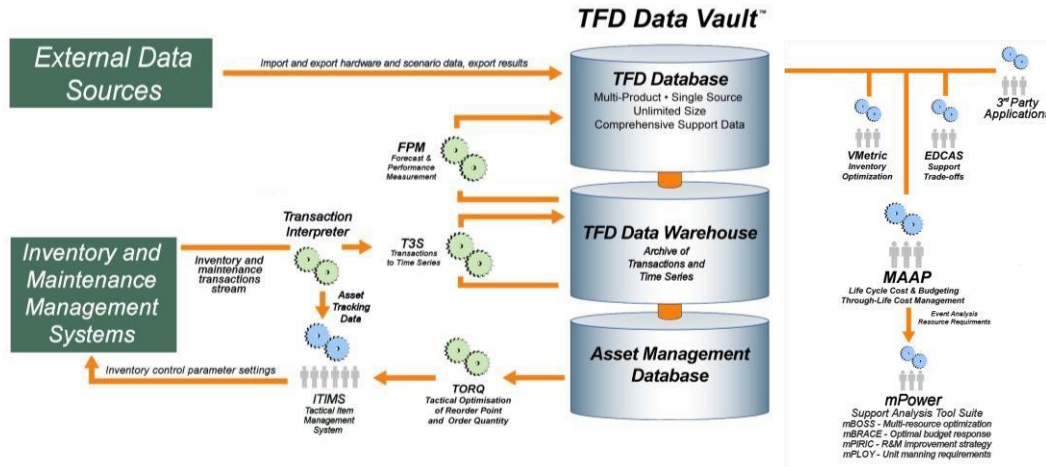
Extended Supply Chain Management

The TFD Global's response to this analytical challenge is its Extended Supply Chain Management (*XSCM*) System. The key features of *XSCM* are:

- a common database infrastructure, serving the needs of multiple analytical models in regard to definition of resource characteristics, system configurations (whether as-built or as-maintained), operating and support environments, operational scenarios and performance requirements;
- SCO (Support Chain Optimization), a tactical decision support system focused on the needs of IMs; and
- an unmatched suite of strategic decision support tools, enabling resource requirements to be projected over time in the light of fleet size changes, unfolding operational contingencies, ageing of components, technology insertion, etc.

The *XSCM* system architecture is shown below (note that computational elements are portrayed by a "meshed gears" icon, green for back-office and blue for user software).

¹² The average number of shortages is commonly known as Expected Backorders (EBO). If the number of supported systems is N, then $EBO \approx -N \ln(1 - NMCS)$



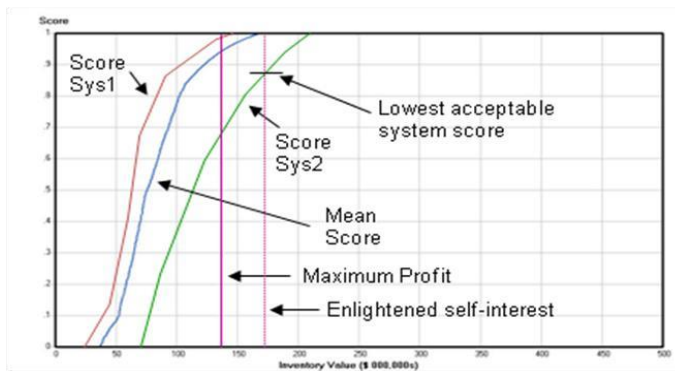

XSCM System Architecture

The SCO operating concept is that inventory and maintenance transactions are harvested from whatever sources may be available. The transactions are processed by a Transaction Interpreter, giving rise to information about changes in inventory state and locations. This information is sent directly to the *ITIMS* software for reference by IMs on a near-real time basis – or whatever the update process supports. The same information is stored in the Asset Management component of the TFD Data Vault, making it available for conversion by the *T3S* (Transactions to Time Series) program into various forms (chiefly time series, hence the name) for use in forecasting routines. At suitable intervals, the *T3S* outputs are fed to the Forecast and Performance Measurement (*FPM*) module, which produces updated mean, variance and trend data for demand rates, delay times and any other variables that can be mined from the available transactions. These updated values are stored in the TFD Database (TFDdB) and made available for use by *TORQ* (Tactical Optimization of Response Times and Quantities).

Because the long-term, steady-state pipeline assumptions of traditional Readiness

Based Sparing models are inapplicable inside the tactical planning horizon, *TORQ* uses Monte Carlo simulation in overnight runs to produce reliable estimates of backorder hours over the remainder of the current evaluation period, taking into account anticipated part deliveries, the age profiles of fitted components, and so forth. Unless tactical decision support is compatible with the assumptions on which contract negotiations were based, outcomes will disappoint the recipient, the provider or both. Accordingly, *TORQ*, unique among tools with claims of supporting tactical decisions, bases its quantification of the costs and benefits associated with plausible supply chain management interventions (e.g. buy more spares, redistribute stock, speed up maintenance or procurement, relocate parts) on a comprehensive understanding of the context: the nature of the supported system, its deployment and operating profile, the current states of buffers and repair or re-procurement queues, apparent demand, performance requirements, and the contract terms actually in force.

The process of configuring TORQ's evaluation functions draws heavily on the logic, discussed at length in an earlier part of this paper, for establishing satisfactory incentive provisions. A natural byproduct of the profit versus investment function is a plot of Expected Backorders (EBO) against profit as shown at right. Since the slope of the curve at whatever value of EBO has been achieved up to the present time translates directly into the expected monetary benefit associated with each intervention (of whatever kind) available to IMs to reduce delay, this plot serves as a means of ranking interventions in descending order of ROI or other relevant criteria.



In similar vein, a family of score functions can be generated to cater for multiple contracts. Because the applicable incentive pools will differ from one support contract to the next, the mean score function has to be a “stakes-weighted” average of the individual score functions. It is important that contracts for supply of common parts used by multiple systems be in harmony with the set of individual system objectives. In this regard, the key to retail/wholesale alignment is the

realization that a “backorder budget” consistent with profit maximization at retail can be assigned to the range of commonly managed parts. Note that comparatively low performance outcomes can be expected for some contracts (the ones with the rightmost score functions) if profit maximization is slavishly pursued. Enlightened self-interest will almost certainly dictate that the lowest individual score objective be kept above a designated minimum level, implying a somewhat lower backorder target than needed simply to maximize short-run profit.

While the SCO process described above is sufficient to support day-to-day IM decisions, strategic considerations must also be taken into account. For example, evaluation of R&M proposals is a frequent requirement, and a budget crunch partway through the year would cause an upheaval in many of the underlying decision criteria. Rather than rely solely on tactical adjustments to creep toward a new set of levels, strategic optimization tools can be used to identify appropriate one-time adjustments. Links to strategic tools are implicit in the XSCM system – data updated continuously through SCO now becomes continuously available, virtually eliminating a stupendously time-consuming and uninteresting part of the analyst’s job – data collection, updating, validating, reformatting, etc., before any analysis can actually be performed.

For more information about all TFD products please contact:

US: Robert.Nomelli@tfdg.com Tel: +(831) 649-3800

UK: Fergus.hawkins@tfdg.com Tel: +44 (0) 7425 801932