

INTERNET DOCUMENT INFORMATION FORM

A. Report Title: Demystifying RBL

B. DATE Report Downloaded From the Internet: OCTOBER 18, 1999

**C. Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): DEPARTMENT OF THE AIR FORCE,
AIR FORCE JOURNAL OF LOGISTICS
501 WARD STREET
GUNTER ANNEX
MAXWELL AFB, AL 36114-3236**

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

**F. The foregoing information was compiled and provided by:
DTIC-OCA Initials: __PM__ Preparation Date OCTOBER 18, 1999**

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.

19991022 114

Demystifying RBL

Major David A. Fulk, PhD, USAF

Whether you call it Readiness-based Leveling (RBL), the D035E system, the Air Force leveling system, or an upgraded version of the D028, RBL is the cornerstone of the supply system for setting recoverable parts levels in the supply system. Developed from the ashes of the old D028 system, RBL was implemented in April 1997 to allocate the worldwide requirement to all bases. In the previous system, Repair Cycle Demand Level (RCDL), levels were computed locally with a relatively simple formula. To know how the level was computed, all one had to do was peruse AFM 23-110. In RBL, levels are computed centrally by Air Force Materiel Command and pushed to the users. The RBL model that calculates the levels has often been viewed as a *black box* where data goes in and levels come out and only a very few people know what goes on inside. Because of this, users often feel that RBL is more of a *shove* system than a push system.

Stock Levels

What Is a Level?

The concept of stock levels for recoverable parts predates RBL, but there seems to be a lack of understanding of the fundamental concept of a level. Since the purpose of RBL is to determine levels, leveling is discussed first.

First and foremost, a level is *not* an asset. However, there should be an asset in some form in the system to cover that level. So a level can be thought of as the number of assets desired.¹ Second, a level is permission for a base to order a part.² If the number of assets on hand plus the number of requisitions in the system is less than the level, additional requisitions are authorized. A level is also a cap on the number of assets a user should have on hand. If by chance the number of assets is larger than a user's level, then the user could (and should) be forced to redistribute the assets to someone else who needs it more. Finally, a level is a method for effectively allocating a scarce resource.

A Level Does Not Equal Assets on Hand

Levels account for requisitions in the system. In fact, levels can generally be thought of as having a *pipeline* portion and a *safety* portion. The old RC DL formula made that easy to understand. The pipeline segment was computed, then a multiple of that segment was added to cover the safety portion. In RBL, these two parts of the level are not computed or reported separately, so it is easy to forget they both exist.

The pipeline portion of the level is designed to cover requisitions in the system. An unserviceable asset is placed in one end of the repair *pipeline*. It flows to the necessary on-base repair shops or depot repair facility and eventually flows back to the base where it comes out the other end of the pipeline as a serviceable asset. The part of the level for the pipeline should be expected to cover assets that are *not* serviceable on hand in stock. The remaining portion of the level, the safety level, is designed to cover some variability in the process. This includes such things as variability in demand, repair time, and order and shipping time. It is not designed and cannot be designed to cover all the variability in the system. Unusual events can occur.

A major misunderstanding concerning levels is that *a level should equate to an on-hand asset*. This is simply not true. On average, only

the safety level should be on hand, and that presupposes all the assumptions made in the pipeline model are true. Serviceable assets on hand will always be less than or equal to the level and many times less than the level.

Is Happiness a Level?

One of the most frequently asked questions is, "Why not let people just order as much as they like? The depot will send out only what they have money to fix." There are numerous problems with this idea. First, this makes requisitions meaningless. Under RBL, users only get a level if they really need one (based on projected demands or an adjusted stock level). A requisition covers past demands and needs to be filled because there is a reasonable chance it will be needed to cover a future demand. Second, the Air Force has only limited resources to buy and repair recoverable assets. These limited resources, along with user needs, are all factored into creating the worldwide requirement. To ignore the requirement (by having levels too high or too low) is an inefficient use of resources. Third, if a user has levels that are too high when compared to others and the actual requirement, they can wind up with assets that are needed more elsewhere. This potential misallocation of assets could cause more back orders (BOs) and nonmission capable aircraft at other bases than it saves at the base with the level that is too high.

Expected Back Orders

Conceptual Example

How does RBL determine levels? It does so by finding a mix of levels for all users (base and depot) that minimizes the user time weighted expected back orders (EBOs). Before discussing how RBL makes this allocation, it will be useful to get a basic understanding of EBOs through a conceptual example.³

In the example in Table 1, on the first day, there are no back orders. On the second day, two back orders occurred, lasting 5 and 7 days respectively. On the third through fifth days, no new BOs were noted, just the existing two BOs. On the sixth day, a third BO was noted and it lasted 3 days. The last new BO was noted on day 9, and it still exists at the end of the 10-day period.

Expected Back Orders: A Definition

The number of BOs in existence on each day is: 0, 2, 3, 2, 1. Taking a daily average gives: $(0+2+2+2+3+2+2+1+1)/10 = 1.7$. This is expected back orders or EBOs. The most common interpretation is EBOs are the average number of BOs in the system at any moment in time. So in the example, on average, there are 1.7

	Day									
	1	2	3	4	5	6	7	8	9	10
1 st Back Order		X	X	X	X	X				
2 nd Back Order		X	X	X	X					
3 rd Back Order						X	X	X		
4 th Back Order									X	X

Table 1. EBO Example

(Continued on page 36)

back orders at any point in time. However, this definition is not always clear or useful, but there are other equivalent interpretations, such as BOs existing each day per day, BO days per day, or days back ordered per day.

A More Useful Form

To obtain a more useful form for EBO, start by counting the number of back orders per day and then rearrange the terms:

$$EBO = \frac{0+2+2+2+2+3+2+2+1+1}{10} = \frac{0+1+1+2+2+2+2+2+3}{10}$$

$$= \frac{0*1+1*2+2*6+3*1}{10}$$

In this last equation, 0 occurred one time, 1 occurred two times, 2 occurred six times, and 3 occurred once. If each of the terms in the numerator is divided by the 10 in the denominator, the result is:

$$EBO = \frac{0*1}{10} + \frac{1*2}{10} + \frac{2*6}{10} + \frac{3*1}{10} = 0*(\frac{1}{10}) + 1*(\frac{2}{10}) + 2*(\frac{6}{10}) + 3*(\frac{1}{10})$$

What are the numbers in parentheses? Probabilities. Consider the BOs in existence each day: 0, 2, 2, 2, 2, 3, 2, 2, 1, 1. 0 occurs one out of ten times, 1 occurs two out of ten times, 2 occurs six out of ten times, and 3 occurs one out of ten times. So the probability of getting a 0 based on this sample is 1/10 and so on. Now going back to the EBO computation, one sees: EBO = 0*Prob of 0 + 1*Prob of 1 + 2*Prob of 2 + 3*Prob of 3. Noting from the sample that the probability of getting a 4 is 0/10, then the probability of getting a 5 is 0/10, etc. EBOs can now be written as:

$$EBO = \sum_{i=1}^{\infty} i * (\text{Probability of } i \text{ backorders})$$

Mathematical Foundation

At the center of the RBL model is the Multi-Echelon Technique for Recoverable Item Control algorithm. This method was first developed by Sherbrooke in the 1960s, but it is still applicable today.⁴ A short synopsis of the mathematical foundation of the method follows.

Probability Functions

The demand pattern, length of the repair cycle, and so on are not constant values but vary (sometimes tremendously) over time. However, theory has shown they follow patterns that can be relatively accurately modeled using standard probability functions. In this case, RBL uses the negative binomial probability.

$$P(x) = P(x; \mu) = \frac{(k+x-1)!}{(k-1)! x!} \left(\frac{q-1}{q}\right)^x \left(\frac{1}{q}\right)^k$$

$x = 0, 1, 2, \dots$
 $q > 1$ is the variance - to - mean ratio
 $k = \frac{\mu}{q-1} > 0$

This is interpreted as the probability of exactly x demands given mean m where m is the average number of demands during a repair cycle (base and depot repair, order and shipping time, and so forth.). So to use this probability function, two parameters are needed: mean, m , and the variance to mean ratio, q . In RBL, q is obtained through an empirical formula instead of using the data.⁵ There are both base and depot means to consider. The formulas used in RBL for these are:

$$\text{Depot Mean} = \text{Depot DDR} * \text{Depot RCT}$$

$$\text{Base Mean} = \text{RTS} + \text{NRTS} = [\text{Base DDR} * \text{PBR} * \text{Base RCT}] + [\text{Base DDR} * (1 - \text{PBR}) * (\text{OST} + \text{NCT} + \text{ADDD})]$$

Where:

- RTS = Repaired This Station
- NRTS = Not Repaired This Station
- DDR = Daily Demand Rate
- RCT = Repair Cycle Time
- PBR = Percent Base Repair
- OST = Order and Ship Time
- NCT = NRTS/Condemned Time
- ADDD = Average Depot Delay Per Demand

Look at each of these means. The depot mean takes the depot daily demands (sum of based NRTS demands) times the average number of days to repair in the repair cycle to give the number of demands during an average depot repair cycle. The base mean involves a similar computation; however, it just splits the demands between those repaired locally (RTS) and those sent to the depot for repair (NRTS).

Expected Back Orders

Unlike the conceptual example, in RBL, there is no sample of back orders to compute the probabilities. As a result, an assumption must be made as to the distribution. As can be guessed, the assumption is made that demands are distributed based on the negative binomial just discussed. In addition, it must also be realized that in RBL the number of back orders is going to be dependent on the number of levels allocated to satisfying demands and the pipeline. With this in mind,

$$O(s) = \sum_{i=1}^{\infty} i * (\text{Probability of } i \text{ backorders given } s \text{ levels})$$

given s levels. Remember that $P(x)$ is the probability of x is the general representation of the expected back orders demands and that a level can be thought of as a unit on hand or due in during the repair cycle period. With s levels and x demands, then $x-s$ is the number of demands for which there is no asset, therefore a back order. If $i=x-s$ and is substituted into the above equation:

$$EBO(s) = \sum_{i=1}^{\infty} i * (\text{Probability of } i \text{ backorders given } s \text{ levels})$$

Expanding this equation for $s-1$ levels and s levels respectively:

$$\begin{aligned} \text{EBO}(s-1) &= 1 \cdot P(s) + 2 \cdot P(s+1) + 3 \cdot P(s+2) + \dots \\ \text{EBO}(s) &= 1 \cdot P(s+1) + 2 \cdot P(s+2) + 3 \cdot P(s+3) + \dots \end{aligned}$$

Subtracting results in:

$$\begin{aligned} \text{EBO}(s-1) - \text{EBO}(s) &= 1 \cdot P(s) + (2-1) \cdot P(s+1) + (3-2) \cdot P(s+2) + \dots \\ &= P(s) + P(s+1) + P(s+2) + \dots \\ &= \sum_{x=s}^{\infty} P(x) = \sum_{x=0}^{\infty} P(x) - \sum_{x=0}^{s-1} P(x) \\ &= 1 - \sum_{x=0}^{s-1} P(x) \end{aligned}$$

But the last sum is just the cumulative probability function at $s-1$. Since this is a positive value less than 1, the change in expected back orders by adding the *next* level is positive. That is, when a level is added, the expected back orders decrease by 1—cumulative probability ($s-1$).

What are the EBOs if no levels are given? The EBOs will equal the mean (depot or base as appropriate). Why? If there are no levels, then each demand will cause a back order. Since there are no levels to account for parts in the repair cycle, it cannot be assumed that there are any parts in the pipeline due into the base. So it will take the entire repair cycle for the back order to be satisfied. Or the number of back order days per demand will equal the repair cycle period. That multiplied by the number of demands per day (DDR) equals the number of back order days per day. But that is just EBOs. Therefore, $\text{EBO}(0) = \text{DDR} \cdot (\text{Repair Cycle Period}) = \text{Mean}$.

Depot Impacts

In the previous example, the levels allocated and the probability function determine EBOs. Earlier, it was seen that the probability function used the base mean (m), but how does the depot mean and depot levels come into play? In the base mean formula, the term ADDD (average depot delay per demand) appears. It is in this term that the depot levels, demands, and delays are accounted for.

$$\text{Average Depot Delay Per Demand} = (\text{Average Depot Delay Per Day}) / (\text{Average Demand Per Day})$$

But, Depot Delay is just another way of saying Depot Back Order Days. So:

$$\text{ADDD} = (\text{Average Depot Back order Days Per Day}) / (\text{Average Demand Per Day})$$

Earlier it was shown that back order days per day is the definition of EBOs. Also, average demands per day is the definition of DDR, so:

$$\text{ADDD} = (\text{Depot EBOs}) / (\text{Depot DDR})$$

To compute the depot EBO, exactly the same approach is used; however, the depot mean and depot levels are used instead of base means and levels. The longer the average depot delay per demand, the longer it takes to return an asset to the base, which increases base back orders. RBL measures the effect on the base EBOs of the level allocated to the depot.

RBL Allocation

Detailed Algorithm

Increasing the levels changes EBOs, but how does the system choose which base to give a level to? RBL uses what is often referred to as a marginal analysis approach. The basic allocation rules are as follows:

1. Obtain the input variables (DDR, RCT, OST, NCT, PBR) for all bases and depot for one subgroup master (SGM) national stock number (NSN). If more than one SGM is present for the family, read all the SGM NSNs for the family. Sum the base DDR and weight average the other quantities. This gives the input variables at the family master level—where RBL computes levels.
2. Set the depot level to 0. From that and the input variables, depot EBOs can be computed, but more important, the ADDD can be computed.
3. Set all base levels to 0. From that, ADDD, and the input variables, the base mean and $\text{EBO}(0)$ can be computed for each base.
4. Using the change in EBO formulas given earlier, compute the reduction in EBOs for all bases going from 0 to 1 level [$\text{EBO}(0) - \text{EBO}(1)$]. The base that has the biggest reduction in EBOs is selected to receive the level. Only one level is allocated, and all other bases do not receive a level (yet).
5. Repeat the previous step determining the reduction in EBOs for all bases getting their next level. For one base, that will be going from 1 to 2 levels; the others will be going from 0 to 1 level. Once again, the one with the largest decrease in EBOs is selected.
6. Keep allocating levels to the bases one at a time until the number of levels allocated to the depot and bases equals the requirement, then stop.
7. Sum the base EBOs to obtain the system EBOs. That is the best allocation given 0 depot levels.
8. Now, try a depot level of 1 and repeat steps two through seven. If the resulting system EBOs are less than the previous allocation, keep the new allocation. Otherwise, keep the allocation with 0 depot levels.
9. Keep trying to increase depot levels until the entire requirement is given to the depot. For each depot allocation, an optimal base allocation is obtained and system EBOs compared in order to keep the smallest one.
10. When done, there is an allocation to the depot and bases with the smallest system EBOs. This is reported, and the whole process is repeated for the next NSN.

Example. In order to further understand the allocation process, an NSN has a requirement of five and three stock record account numbers (SRANs) (A, B, and C). The results of each pass through the algorithm are shown in Table 2.

In this example, there are six passes, one each for the depot levels 0 to 5. The detailed middle steps are eliminated, and just the results at the end of each pass are shown. In all six passes, the entire requirement of five was allocated, some to the depot and some to SRANs. Looking at the system EBOs, the smallest value is for the third pass.⁶ So that is the allocation that will be used.

	1st pass		2d pass		3d pass		4th pass		5th pass		6th pass	
	Level	EBO	Level	EBO	Level	EBO	Level	EBO	Level	EBO	Level	EBO
Depot	0		1		2		3		4		5	
SRAN												
A	2	0.90	2	0.70	1	0.80	1	0.75	1	0.70	0	1.20
B	2	0.80	1	0.70	1	0.65	1	0.60	0	1.00	0	0.90
C	1	0.50	1	0.40	1	0.30	0	0.80	0	0.70	0	0.60
System EBOs		2.20		1.80		1.75		2.15		2.40		2.70

Table 2. RBL Allocation Sample

There are, however, some interesting outcomes seen in Table 2. For example, looking at SRAN B going from the first to second passes, the base levels decreased, but so did the EBOs. How can that happen? Because of the multi-echelon nature of RBL, this is frequently the case. The depot level increased, which would cause the depot delay to decrease (ADDD decreases). Since ADDD is used in the base mean, the base mean would also decrease. A different base mean results in a slightly different probability function (remember the mean is used in the probability function) and different EBO value. Outside of the mathematics, if the SRAN sends most of its parts to the depot, having a larger depot level will help the base. At the same time, having fewer levels at the base hurts the base. It is a matter of one helping the base more than the other hurting the base (in terms of EBOs). Later on for SRAN B going from the fourth to fifth pass, the depot receives another level, and SRAN B receives one less, but this time EBOs increase since there are diminishing returns. Adding more levels to the depot will always help, but the amount it helps becomes less and less.

You can probably see now why RBL is often viewed as a black box. With several input variables, many SRANs (both organizational intermediate maintenance and depot level maintenance), two echelons (base and depot), uncertainty (probability functions), predicted results (EBOs), and so forth, it is a very complicated process. However, the *real* Air Force supply system is often an even more complicated, intertwined, and uncertain system. Compared to the RCDL, RBL makes great strides in taking many of these factors into account.

Other Issues

We have talked about expected back orders several times here, but the everyday Air Force talks in terms of back orders and not expected back orders. What is the relationship between the two?

Back Orders Versus Expected Back Orders

Many people want to examine BOs and not EBOs. This is because the data on BOs (mission capabilities, due-outs, and so forth) is collected. BOs are event-driven transactions that occurred in the past, whereas EBOs are statistically predicted future averages.

From the conceptual example given in Table 1, there were three BOs. A common mistake is to think the RBL EBO number is the number of back orders or the number of back orders divided by number of days. Neither is true. BO/day would be interpreted as the average number of new back order *occurrences* per day, which is different from the EBO definition. We can also see in the example that $4/10 = 0.4$ is not the same as 1.7 given earlier as the EBO.

Can one be converted to the other? Yes and no. In this simple example, yes. In RBL, it is not as easy. In the conceptual example, take EBOs, multiply by the number of days and divide by the average length of a back order to get number of BOs. In the

example, the BOs lasted 5, 7, 3, and 2 days, so the average length was 4.25 days. Therefore, number of BOs = $EBO * (10 \text{ days}) / (4.25 \text{ days per BO}) = 1.7 * 10 / 4.25 = 4$ BOs. That was simple enough. But what about RBL? In RBL, the average length of a back order is not known. Approximations have been attempted using various methods that included fixed numbers, average order and ship period, and the average repair cycle period. It is hard to verify these approximations because of many other factors. Therefore, this is not done.

Are EBOs Good Numbers?

Yes. Minimizing EBOs means that either the number of BOs is minimized, or the length of the BO is minimized, or both. Eliminating BOs, or at least reducing the number of them, is a primary goal. However, that cannot always be done, but the user can still be helped by reducing the time waiting for the part. So reducing the number or length of BOs is a good goal, and using EBOs allows both to be minimized.

What are some of the *other factors* that keep EBOs from RBL (even after conversion/approximation to BOs) from being closer to *real world* BOs?

1. **Back Order Length.** The average length of a back order is not completely known, so there are inaccuracies in the conversion from EBOs to BOs.
2. **Time Frames.** BOs are discrete events from the past given such things as existing levels, assets, and funding. RBL uses past data to predict future data and determine the optimal levels, which may be different than currently exist. So EBOs are forward-looking predicted values that are based on several assumptions.
3. **Changing Demand Pattern.** A major assumption in RBL is that past demands are good predictors of future demands (also true of RCDL, economic order quantity, and almost any supply system). If this assumption does not hold for some part, then EBOs were computed on the wrong values. Of course, this would not be known until after the fact.
4. **RBL ignores some of the *real world*.** For example, it does not consider parts in a Readiness Spares Package (RSP) in determining levels, yet a base with an RSP will use those parts to avoid BOs. Similarly, RBL does not consider High-priority Mission Support Kit parts, cannibalization, and lateral support. These are not mistakes but deliberate choices by the supply community to not consider those parts and concepts when leveling.
5. **Assets.** Readiness-based Leveling deals with levels, not assets. It assumes that assets are available (or will be made available) if levels are available. In general this is true, but for a few thousand parts it is not true.
6. **Funding and Priorities.** RBL has to assume that a part will get fixed based on a repair pipeline. In reality, some parts are never fixed because of funding and priorities or get fixed and sent to places other than the base that is next in the queue based on priorities.
7. **Bottom Line.** The RBL EBO number is good based on what it is asked to do. RBL is given some reasonable assumptions and told to ignore certain things and then come up with base and depot levels that should provide the best overall support.

Special RBL Rules

So RBL does a reasonable job with what it is asked to do; but it has the tasking to do everything for everybody. So there have to be some exceptions and other special rules for different subsets of the parts. Although these special rules can confuse matters a bit, having all recoverable parts run through RBL, even if they use a special rule, puts all the rules in one place. Otherwise, there could be many different leveling systems in many places—not a good choice. Some of the special rules in RBL follow.

1. **Adjusted Stock Levels (ASLs).** ASLs are honored in RBL as they are part of the worldwide requirement. Since a base with an ASL may or may not have demands, the regular algorithm will not work for them. Given the worldwide requirement is sufficient, ASLs are essentially allocated first, regardless of the savings in EBOs at any other base. The ASL was (theoretically) approved and included in the requirement, so RBL should allocate the requirement to the ASL. For each level allocated for an ASL, the EBOs for that base are reduced just as they were before. However, if the base had no demands, the EBO would be zero for that base, and there would be nothing to reduce. Once all the levels to support ASLs are allocated, the algorithm continues as before allocating to the user that reduces EBOs the most. If the ASL is a fixed type, that base is eliminated from consideration for any more levels once the fixed ASL is met. Similarly, if a maximum ASL is present, once the maximum is reached (including a Max 0), that base is no longer considered for a level.
2. **Initial Spares Support Lists (ISSLs).** For the most part, ISSLs are allocated like any other minimum ASL. However, there are some NSNs where the requirement is insufficient to meet all the needs. In those cases, ISSLs are given a lower filling priority. That is, regular ASLs and demand pipelines will be filled before ISSLs are considered.
3. **Contingency Spares Support Levels (CSSLs).** CSSLs are levels in support of a Contingency High-Priority Mission Support Kit (CHPMSK). CHPMSKs use peacetime stocks in support of contingency operations. CSSLs are added to regular ASLs and then allocated in RBL like ASLs. However, RBL levels pushed to the base have the CSSLs deleted so as not to double count them in the requisition objective at the base, since CHPMSK increases the requisitioning objective (by the CSSL amount).
4. **Smaller Depot Level Cap.** When considering one pass with a given depot level and the next pass, the one with the smaller system EBO is kept. But what happens if the two allocations have virtually the same system EBOs? Instead of sticking strictly to the algorithm, if the difference in system EBOs is *less than a very small number*, it is considered that they both have the same EBOs. In those cases, the allocation with fewer depot levels and more base levels is used.
5. **Requirement Cap.** The system keeps allocating until all the requirement is given, but there are diminishing returns in giving levels to either the depot or base. So if giving the next level to the best base (the one that reduces EBOs by the most) only changes system EBOs by a trivial amount, the allocation is stopped. Giving levels after this point would fill the system with unnecessary requisitions that probably would not be needed for years (encourages fixing *buggy whips*).
6. **Pipeline Cap.** In order to ensure outliers in the input variable do not overly sway the allocation, RBL caps certain variables. Base repair cycle time (RCT) is capped at 10 days. Depot RCT is capped at 210 days. CONUS Order and Ship Time (OST) is capped at 24 days, while OCONUS OST is capped at 52 days. NRTS/Condemned Time is capped at 3 days.
7. **Insurance and Nonconsumable Item Materiel Support Code (NIMSC5) NSNs.** Insurance NSNs are checked for demand usage. If two or more demands are found, the cataloging is considered suspect, and base levels are allowed. Otherwise, by policy, base levels are not allowed on insurance NSNs. Similarly, NIMSC5 NSNs are parts where the Air Force is the Single Inventory Control Activity. These parts are not allowed to have depot levels as the depot repairing the part is from another Service.
8. **Communications-Electronics Rule.** Budget program 8M parts tend to be expensive and seldom used parts. Based on Air Force Logistics Management Agency (AFLMA) studies and work at the Air Force Communications Agency, special rules were developed to handle these parts. Basically, the rules will fill ASLs first, then try and put at least two levels at the depot, then finally allocate anything remaining to bases with demands (for non-Numeric Stockage Objective parts). The main differences are the depot levels and lack of demand-based levels. The two levels at the depot are to provide a central stocking type concept on most of these parts instead of a distributed stocking concept. The lower priority for demand-based users is because of the very low demands and reliability of the parts. Past demands at one location are not necessarily a good predictor of future demands at that location. ASLs are used in locations where the assets are really needed, single-point failures, and the lack of an asset would make an essential communications system inoperable.
9. **AMC Forward-Supply Locations (FSLs).** FSLs are supply locations that handle the en route needs of the strategic airlift system. At the time this article was written, Headquarters Air Materiel Command computed the levels, and these were loaded as fixed ASLs at the FSLs. RBL then allocates these off the top like any other fixed ASL. However, a project is underway to have RBL do the computation for the FSLs. This would be a separate algorithm within RBL that takes a system approach to demands in determining the allocation.⁷
10. **Depot Working Level.** Once the best overall allocation is determined, the depot level needs to be split into components: consolidated repairable inventory (CRI), work in progress (WIP), and consolidated serviceable inventory (CSI). CSI will be the levels above the depot mean (if there are any). CRI and WIP will prorate the level (up to the depot mean) based on the retrograde portion of depot RCT and repair portion of depot RCT respectively. These pieces of depot RCT are input variables to the model. Once the depot level is split, WIP and CSI are added to form the depot working level. This number is output for use in the Execution and Prioritization of Repair Support System for depot repairs.
11. **Miscellaneous.** There are several other minor rules in RBL, such as ignoring Federal Stock Code 1300 and 1399 items and non-FB SRANs. These are all very special cases and only affect a few parts and users.
12. **Problem Item Heuristic.** The last special rule is for problem items. When the requirement is less than the heuristic pipe (demand pipeline plus ASLs), the NSN is flagged. Something must be wrong with the requirement, data, policies, and so forth for this to occur. The question here is how should RBL allocate these levels? As noted earlier, ASLs are allocated first. This implies some bases would receive less than their mean demand pipelines (for each part put into the pipeline, less than one part comes back out). This unduly harms demand users, the ones who have an established history of use, at the expense of ASLs. The problem item heuristic is a *deepest hole* type of algorithm that resolves this problem. In these cases, the RBL model is run above ignoring all ASLs. The resulting allocation is used as targets for the deepest hole. To these targets, ASLs are added. The heuristic then allocates levels one at a time based

on the largest hole determined by (levels allocated)/(target). This rule provides a more equitable distribution of the shortage between demand users and ASL users. The NSNs are flagged and passed on to AFMC item managers for resolution.

Conclusions

The Air Force supply system has many constraints: resources, manning, funding, and facilities. By performing a *constrained optimization*, RBL provides the best allocation of those limited resources. It reduces back orders and provides a more complete picture of a base's need than RC DL, and it ties the base levels to the funded requirement.

Notes

1. Craig S. Sherbrooke, *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*, New York: John Wiley & Sons, 1992, 23.

2. Steven B. Reynolds, Mark A. Abramson, Douglas K. Blazer, Karl Kruse, and Wayne B. Faulkner, *Setting Recoverable Item Stock Levels*. AFLMA Report LS199500500, December 1995, 3.
3. Henry Triwush, *D028 Training Document*. HQ AFLC, Wright-Patterson AFB, Ohio, November 1978, 40.
4. Sherbrooke, 45.
5. J. J. Bartko, "Approximating the Negative Binomial," *Technometrics*, Vol 8, No 2, 1966.
6. Sherbrooke, 75.
7. James Neice, David A. Fulk, Bradley Silver, Douglas J. Blazer, and Harry A. Berry, *Forward Supply System-Forward Supply Locations (FSL) Data and Requirements Pass*. AFLMA Final Report LS9710500, September 1997.

Major Fulk is currently an analyst at Headquarters Air Combat Command, Langley AFB, Virginia. At the time of the writing of this article, he was an analyst at the Air Force Logistics Management Agency.

