DEFENSE CASE STUDY / MRO LOGISTICS

The Ten-Cent Bolt That Grounds the Aircraft

Why a uniform service level fails readiness, and what pricing a shortage in availability changes

EXECUTIVE CONCLUSION

A single service-level target is a pricing decision nobody made. It under-protects the cheap parts that ground aircraft and over-buys the expensive parts that do not.

In maintenance, repair and overhaul, the cost of a shortage is lost availability, and it differs by orders of magnitude from one part to the next. Stock levels that come from those economics, rather than from a uniform percentage, buy more readiness for the same money. This note describes a Bit Bros engagement with a U.S. Department of Defense organization, the reasoning behind it, and how we coach clients to make these calls themselves.

$0.10
part that can ground a fleet asset

$1.70
extra stock that removes most of the risk

$360K
freed on one expensive, non-critical part

1 rule
applied to all three, wrong for all three

A Bit Bros engagement. The client's identity, data, models and results are not disclosed. Every number in this note comes from an illustrative worked example built to show the mechanism, not from client data.

The planning problem

The organization supported maintenance, repair and overhaul for a fleet through a logistics network in which the right part at the right place decides whether an aircraft is available to fly. Demand for repair parts is intermittent and driven by failures, which are uncertain by nature. Lead times are long and variable. And the cost of running out is measured not in a lost sale but in an aircraft that cannot perform its mission.

Settings like this are often governed by a single target: hold enough of every part to meet, say, 95 percent of demand from stock. The rule is simple to state, easy to audit and identical for every part. That is exactly the problem.

THE DECISION

How many of each part should be held at each location so that the fleet is available when it is needed, at the least total cost, including the cost of the aircraft that cannot fly?

1. One rule, three parts, three wrong answers

Consider three parts stocked at one operating base. A $0.10 airframe bolt is mission critical: without it the aircraft does not fly. An $85 hydraulic filter is also mission critical. A $180,000 cabin display unit is not: the aircraft can fly its missions without it. Apply the same 95 percent fill-rate target to all three, then compare the stock that minimizes total cost.

Part	At 95%	Econ.	What changes
Bolt, $0.10	16	33	−3.9 grounded aircraft-days a year for $1.70 more stock
Filter, $85	11	20	−2.0 grounded aircraft-days a year for $765 more stock
Display, $180K	3	1	$360K of capital freed; +2.6 degraded days a year

The uniform rule is wrong for every part, in both directions at once: it starves the cheap critical parts and overstocks the expensive non-critical one.

Why the bolt is the whole argument

Holding one more bolt for a year costs about two and a half cents; running out grounds an aircraft until a replacement is expedited. The two costs differ by a factor of millions, and at that ratio the right service level is not 95 or 99 percent. It is as close to certain as the arithmetic allows.

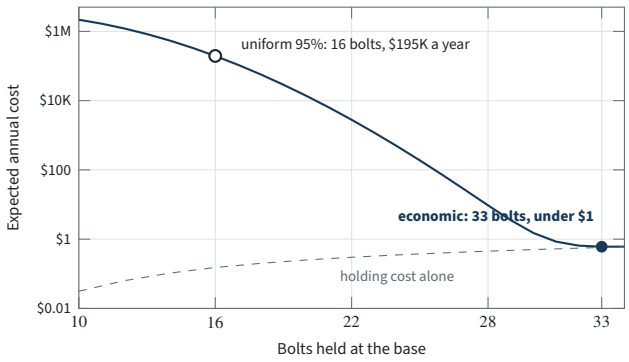

Exhibit 1. Total expected annual cost of bolt stock at one base (solid), mostly the cost of grounded aircraft, against holding cost alone (dashed). Log scale. The 95% target sits where the curve is still steep.

The model is a continuous-review base-stock policy, one unit reordered for each unit used, with Poisson demand over the resupply time. The illustration assumes 40 bolts used a year, a three-month resupply cycle (so lead-time demand averages 10), a two-day expedite when a bolt is missing, an aircraft-day valued at $50,000 and a 25 percent annual holding rate. Fill rate is the share of demands met from stock, which for unit Poisson demand equals $\Pr\{X \le s-1\}$ for lead-time demand X. Scaling any one assumption by a factor of ten moves the economic stock between 11 and 166 bolts, but the extra protection never costs more than $5.

Discreteness matters too. For the display unit, two spares give a 93.8 percent fill rate, so a 95 percent target forces a third $180,000 unit and a realized 99.2 percent. Uniform targets overshoot most on exactly the slow-moving, expensive items where each unit is dearest.

In a surge, the price moves and the rule does not

In peacetime, a grounded aircraft costs a training sortie. In a contingency or a war, it can cost a mission, and the value of an aircraft-day rises by orders of magnitude. An economic stock level rises with it, automatically, because it is computed from that value. A fixed 95 percent says the same thing in a crisis as on a quiet Tuesday, which is precisely when a planning rule most needs to change its mind.

2. The model behind the argument

Symbol	Meaning
λ_i	Demand (failure) rate of item i, per year
T_i	Mean resupply or repair turnaround, years
X_i	Units of i in the pipeline; $m_i = \mathbb{E}X_i = \lambda_i T_i$
s_i	Spares (base-stock level) of item i
$\mathrm{EBO}_i(s_i)$	Expected backorders of i
h, c_i	Holding cost per unit-year; unit price
b, p	Cost per backorder-year; cost per unit short
N, Z_i	Aircraft in the fleet; installs of i per aircraft
A, B	Fleet availability; spares budget

The pipeline is Poisson, and backorders are what ground aircraft

Under a one-for-one (base-stock) policy every failure triggers a resupply or repair order, so the stock position s_i is constant and on-hand stock and backorders are determined entirely by the pipeline X_i:

$$\text{on hand} = (s_i - X_i)^+, \qquad \text{backorders} = (X_i - s_i)^+.$$

Palm's theorem. If demands arrive as a Poisson process with rate λ_i and turnaround times are independent and identically distributed with mean T_i (any distribution), then in steady state $X_i \sim \mathrm{Poisson}(\lambda_i T_i)$. The pipeline behaves like an $M/G/\infty$ queue, so only the mean turnaround matters, not its shape.

Expected backorders are $\mathrm{EBO}_i(s) = \mathbb{E}(X_i - s)^+ = \sum_{x>s}(x-s)\Pr\{X_i = x\}$, with $\mathrm{EBO}_i(0) = m_i$. Taking first differences gives the recursion used in every sparing tool,

$$\mathrm{EBO}_i(s+1) = \mathrm{EBO}_i(s) - \Pr\{X_i > s\}.$$

The decrement $\Pr\{X_i > s\}$ falls as s rises, so EBO_i is decreasing and convex: each extra spare removes less backorder than the one before. Each backorder of a mission-critical item is a hole in an aircraft, and this single convexity property drives everything that follows.

The service level is computed from the costs

Backorder-time cost. If each backorder costs b per year (an aircraft-day's value times 365 when one backorder grounds one aircraft), expected annual cost is

$$C(s) = h\big(s - m + \mathrm{EBO}(s)\big) + b\,\mathrm{EBO}(s),$$

because on-hand stock equals $s - m + \mathrm{EBO}(s)$ in expectation. Using the recursion,

$$\Delta C(s) = C(s{+}1) - C(s) = h - (h+b)\Pr\{X > s\} = (h+b)\Pr\{X \le s\} - b.$$

ΔC is nondecreasing, so C is convex and the optimum is the smallest s with

$$\Pr\{X \le s\} \;\ge\; \frac{b}{b+h},$$

the newsvendor-type condition. For the bolt that ratio differs from one in the ninth decimal place.

Per-shortage cost. When a shortage instead triggers an expedite costing p per unit, as in Exhibit 1, the shortage term is $p\,\lambda\Pr\{X \ge s\}$, since a demand finds no stock when the pipeline already holds at least s units. Its first difference is $-p\lambda\Pr\{X = s\}$, so

$$\Delta C(s) = h\Pr\{X \le s\} - p\lambda\Pr\{X = s\} \;\ge\; 0$$

defines the optimum. The Poisson distribution is log-concave, so the reversed hazard $\Pr\{X = s\}/\Pr\{X \le s\}$ decreases in s; the condition, once met, stays met, and the first s that satisfies it is optimal. This is no longer a fixed percentile, but it is still an output of the costs, computed part by part.

From parts to aircraft: availability

What leadership manages is aircraft availability, not part fill rates. For a fleet of N aircraft in which item i appears Z_i times per aircraft, and assuming backorders are spread across aircraft and not cannibalized, Sherbrooke's approximation is

$$A(\mathbf{s}) = \prod_i \Big(1 - \frac{\mathrm{EBO}_i(s_i)}{NZ_i}\Big)^{Z_i}.$$

Taking logarithms makes the objective separable:

$$\log A(\mathbf{s}) = \sum_i g_i(s_i), \qquad g_i(s) = Z_i \log\Big(1 - \frac{\mathrm{EBO}_i(s)}{NZ_i}\Big).$$

Each g_i is concave and increasing in s: it composes the concave, decreasing function $x \mapsto Z_i \log(1 - x/NZ_i)$ with the convex, decreasing EBO_i.

Readiness-based sparing by marginal analysis

The sparing problem is a separable concave knapsack,

$$\max_{\mathbf{s}\in\mathbb{Z}_+^n} \sum_i g_i(s_i) \quad \text{s.t.} \quad \sum_i c_i s_i \le B.$$

Marginal analysis buys, one unit at a time, the spare with the largest gain per dollar, $\delta_i(s_i) = \big(g_i(s_i{+}1) - g_i(s_i)\big)/c_i$. Because each g_i is concave, every δ_i is nonincreasing, and the greedy sequence passes through exactly the undominated points: each is optimal for the budget it spends (Fox, 1966), and together they trace the convex hull of the availability-investment frontier. Between those budgets the greedy plan is near-optimal, with a gap bounded by the value of a single unit.

ALGORITHM: MARGINAL ANALYSIS FOR SPARING

$s_i \leftarrow 0$ for all i; spend $\leftarrow 0$
compute $\delta_i(0)$ for all i; place items in a max-heap keyed by δ_i
while heap not empty **do**
  $j \leftarrow \arg\max_i \delta_i(s_i)$ ▷ pop the heap
  if spend $+ c_j > B$ **then remove** j; **continue**
  end if
  $s_j \leftarrow s_j + 1$; spend $\leftarrow$ spend $+ c_j$
  record point (spend, $\exp\sum_i g_i(s_i)$)
  update EBO_j by the recursion; push j with $\delta_j(s_j)$
end while

Each step costs $O(\log n)$ for the heap and $O(1)$ for the backorder update, so tracing the whole frontier for thousands of items takes seconds. Exhibit 2 is that frontier on a synthetic catalog.

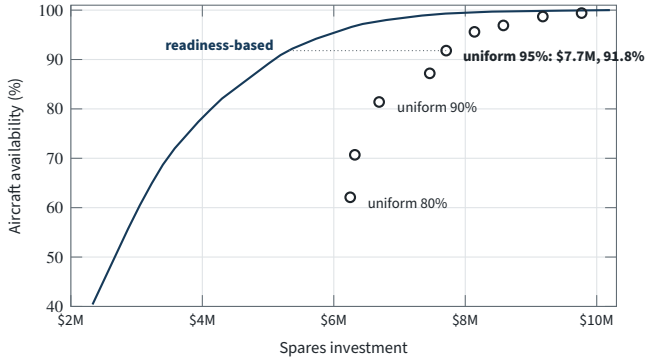

Exhibit 2. Aircraft availability against spares investment for a synthetic catalog of 60 mission-critical items and 24 aircraft: marginal analysis (line) against uniform fill-rate targets from 80% to 99.5% (circles). The dotted line marks equal availability.

On this synthetic catalog, a uniform 95 percent target spends $7.7 million for 91.8 percent availability. Readiness-based sparing reaches the same availability for $5.3 million, about 31 percent less, or 99.3 percent availability for the same money. The uniform rule buys expensive spares to hit percentages that do not move availability, and under-buys the cheap parts that do.

What a real system adds

The single-site model is the core; a fielded system handles the structure around it.

Feature	How it is modeled
Repairable items	Pipeline is repair turnaround plus a condemnation fraction resupplied from new buys
Multiple echelons	METRIC: base pipelines include delay from depot backorders; VARI-METRIC corrects the variance
Indentured parts	Module and sub-module (LRU and SRU) stock both feed the same aircraft availability
Cannibalization	Raises availability above the formula, at a maintenance cost; simulate it rather than assume it away
Lateral supply	Transshipment between bases shortens effective pipelines for slow movers
Demand drivers	Rates per flying hour, so a surge in flying scales demand directly

Two echelons: the METRIC equations

With a depot supporting bases j, let r_j be the fraction of failures repaired at the base (turnaround T_j) and $1 - r_j$ the fraction sent to the depot, which ships a replacement after order-and-ship time O_j if it has stock. The depot sees demand $\lambda_0 = \sum_j \lambda_j(1 - r_j)$ and holds s_0 spares against a pipeline of mean $m_0 = \lambda_0 T_0$. By Little's law the average extra delay a base order waits at the depot is

$$\delta = \frac{\mathrm{EBO}_0(s_0)}{\lambda_0},$$

so base j's pipeline has mean

$$m_j = \lambda_j\big[r_j T_j + (1 - r_j)(O_j + \delta)\big].$$

METRIC treats each base pipeline as Poisson with this mean. VARI-METRIC corrects the variance, because depot backorders are shared out among bases: with $f_j = \lambda_j(1 - r_j)/\lambda_0$ and B_0 the depot's backorders,

$$\mathrm{Var}X_j = \lambda_j r_j T_j + \lambda_j(1 - r_j)O_j + f_j(1 - f_j)\,\mathbb{E}B_0 + f_j^2\,\mathrm{Var}B_0,$$

and a negative binomial with that mean and variance replaces the Poisson. Optimizing s_0 and the s_j jointly by marginal analysis over this model is the classical multi-echelon sparing method.

Intermittent demand needs its own estimators

Many parts fail a few times a year, with long gaps between demands, and simple averages are biased and noisy for such series. Croston's method smooths demand size z and the interval q between demands separately, updating only in periods with a demand:

$$\hat{z} \leftarrow \alpha z + (1 - \alpha)\hat{z}, \qquad \hat{q} \leftarrow \alpha q + (1 - \alpha)\hat{q},$$

and forecasts the demand rate as $\hat{z}/\hat{q}$. That ratio is biased upward; the Syntetos–Boylan approximation corrects it to $(1 - \alpha/2)\,\hat{z}/\hat{q}$. For sparing, rates are better expressed per flying hour and fitted as Poisson or negative binomial models, which pool information across fleets that share a part and scale directly with a surge in flying. Removals later found to be "no fault found" should be separated from true failures before any rate is estimated.

3. How the engagement approached it

We treated parts demand as intermittent and uncertain, driven by failures rather than by a steady usage rate. We framed the objective around availability and readiness of the assets served, not around inventory turns or cost alone. And we used simulation to test candidate policies against realistic variability before recommending any change, so that each recommendation arrived with evidence of how it behaved in bad conditions as well as average ones.

Simulation earns its place technically too. Analytic availability models rest on approximations (Poisson pipelines, independence across items, no cannibalization); a simulation of failures, repairs and resupply checks them against the actual network and replays surge profiles in which flying hours rise while pipelines lengthen.

REFERENCE DECISION FLOW

Failure-driven demand → price of a shortage, by criticality → readiness optimization → simulation under stress → recommended stock levels

Specific outcomes are not published, consistent with the nature of the engagement.

4. We coach the economics, not just deliver the model

A model handed over without its reasoning is a black box the organization will override at the first surprise, and it should. Bit Bros works the other way: we sit with maintainers, logisticians and finance and price the shortage together, part class by part class, and work the optimality condition by hand on a dozen real parts before any software is involved. The model is the second deliverable. The first is an organization that can set, defend and revise its own stock levels by asking what a shortage actually costs.

5. Limitations to state plainly

The value of an aircraft-day is a judgment, not a measurement: elicit it with the people who own readiness and test how much the answer depends on it. Demand is rarely stationary across fleet ageing, modifications and operating tempo. A sound system makes these assumptions visible rather than burying them in parameters.

6. Lessons that transfer beyond defense

Criticality is an economic quantity

Labels such as “critical” and “non-critical” are a start, but they hide a range of several orders of magnitude. Put a value on the consequence of a shortage, even a rough one, and the stock level follows from it.

Let the price move with the situation

When the value of availability changes, in a surge or a seasonal peak, the stock policy should change with it. A rule computed from economics does that automatically. A fixed percentage cannot.

A USEFUL EXECUTIVE DIAGNOSTIC

Which of your cheapest parts can stop your most expensive asset, and what service level are they set to today?

Further reading

Sherbrooke, C. C. *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*, 2nd ed. Springer, 2004.

Muckstadt, J. A. *Analysis and Algorithms for Service Parts Supply Chains*. Springer, 2005.

Graves, S. C. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science* 31(10), 1985.

Fox, B. Discrete optimization via marginal analysis. *Management Science* 13(3), 1966.

Sherbrooke, C. C. VARI-METRIC: improved approximations for multi-indenture, multi-echelon availability models. *Operations Research* 34(2), 1986.

Croston, J. D. Forecasting and stock control for intermittent demands. *Operational Research Quarterly* 23(3), 1972.

Syntetos, A. A. and Boylan, J. E. The accuracy of intermittent demand estimates. *International Journal of Forecasting* 21(2), 2005.

About the case and Bit Bros

This paper describes a Bit Bros engagement with a U.S. Department of Defense organization. The client’s identity, data, models, operational details and results are not disclosed, and the worked example uses illustrative numbers chosen to show the mechanism. Bit Bros is a boutique applied economics and optimization firm. We help operating teams price their decisions, build the systems that make them repeatable, and coach the people who own them to reason economically.

Adam DeJans Jr.
Technical Managing Director, Bit Bros LLC
Decision scientist specializing in optimization, simulation and decisions under uncertainty.

John Brandon Elam
Managing Director, Bit Bros LLC
Leads client engagements, working with executives to find the decisions that carry the most value.

Price the shortage. The service level follows.