

Transport deadlines and common demand batches in regional generator spare stock allocation

Mahmoud Abdelfattah Abbas Mahmoud

mahmoud.hti@gmail.com

High Performance Industry (HPI)

Research Article

Keywords: generator spares, inventory pooling, common demand, compound Poisson, transport deadline, regional stock

Posted Date: October 15th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11269130/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: Competing interest reported. The author is employed at High Performance Industry (HPI). This employment is disclosed as relevant professional context. No company or client operational records, proprietary equipment data or financial data were used. The author has declared no other competing interests.

Abstract

Pooling generator spares can improve stock availability while reducing the fraction of requests delivered within a deadline. This analytical and computational study compares one central pool with three regional stocks under hypothetical demand rates, common three-unit events and deterministic replenishment and transport. Exact compound Poisson probabilities give transport qualified immediate fill, and a first come first served waiting-time extension gives actual delivery by the deadline, excluding installation. At nine spares, seven day replenishment and a common event rate of 0.02 per day, central qualified immediate fill is 0.2131, 0.5119 and 0.9820 at eight, twenty four and forty eight hours. Including backorders replenished before the deadline raises these values to 0.2135, 0.5151 and 0.9917. Regional stock of two, three and four units gives delivery fractions 0.9132, 0.9312 and 0.9542, compared with immediate fill 0.9090. Both endpoints are computed across 108 architecture scenarios. An equivalent mean Poisson pipeline understates the stock needed in a separate high common-demand example. Independent convolution, boundary identities and 1200 additional event histories verify the delivery calculation; individual Monte Carlo interval misses are retained and reported. The contribution is an exact, batch aware comparison of stock and deadline endpoints under declared queue assumptions, not a measured restoration commitment.

1 Introduction

Generator after sales networks must position infrequently demanded critical parts before a failure is discovered. A central stock can consolidate spare units that would otherwise be separated among regions. The longstanding METRIC framework describes recoverable item demand and stock allocation using compound Poisson concepts (Sherbrooke 1968). The present benchmark adopts one for one replenishment and explicitly defines its pipeline count, while omitting the base depot repair structure of METRIC. It is not a replacement implementation of that multi echelon method.

Transportation affects whether available stock can meet an operational deadline. Wong et al. examined pooling with non-zero lateral transshipment times and delayed transshipments (Wong et al. 2005). Their system differs from the mutually exclusive architectures here, but establishes why treating transfer time as zero can change a stocking comparison. This paper compares either a central pool or regional stocks; it does not model lateral transfers among regional warehouses or a combined central and regional inventory.

Demand dependence also matters. Yan et al. investigated inventory performance with pack size constraints and spatially correlated demand (Yan et al. 2009). In a generator service context, concurrent requirements might arise from a common event or a coordinated maintenance campaign. Those examples are motivations, not observed frequencies in this study. A common event here creates exactly three requests for an interchangeable part, one in each modeled region. The resulting central demand is not a Poisson process of independent single requests even though each region's marginal request count is Poisson.

The research question is how a central stock advantage changes when evaluated using a defined transport qualified immediate fill endpoint and batch aware arrival sampling. The contribution is an exact and reproducible architecture comparison, a controlled equivalent mean Poisson approximation audit, and enumeration of regional allocations. Compound Poisson demand, pooling and stock optimization are established methods; the study makes no claim of discovering those methods or representing measured demand in Saudi Arabia.

Reijnen, Tan and van Houtum used a time based fill objective in an inventory network with reachable warehouses and emergency shipments (Reijnen et al. 2009). Their model has an ample central source, individual Poisson customer demands and lateral shipment options. The present benchmark uses common regional batches and FCFS backorders without lateral or emergency supply, and derives the waiting contribution separately. Delivery-time service is an established objective; the comparison here isolates endpoint and batch assumptions.

2 Materials and methods

2.1 Demand and replenishment conventions

There are three regions with independent single unit request rates λ_1 , λ_2 and λ_3 per day. A separate Poisson process with rate κ generates common events, each demanding one unit in every region simultaneously. All parts are interchangeable and all requests are backordered when not immediately filled. Every unit demand places a replenishment order regardless of stock availability. Each order arrives after exactly L days. Supplier capacity is unlimited and the lead time is not increased by demand congestion.

For a base stock S , the net inventory just before a demand is S minus the number N of outstanding replenishment units. The immediately available inventory is the positive part of that quantity. Orders generated before the current arrival occupy the pipeline; the current batch is then allocated by an independent uniform random permutation of its three regions. A region with a shorter transport route receives no special stock priority. Deterministic transport for the immediate-fill endpoint occurs only after an immediate fill and does not consume the replenishment lead time specified above. Installation, diagnosis and technician scheduling are excluded.

The central pipeline count has the following decomposition and probability recursion:

$$N = X + 3Y; \quad X \sim \text{Poisson}(\Lambda L); \quad Y \sim \text{Poisson}(\kappa L) \quad (1a)$$

$$p_0 = \exp[-(\Lambda + \kappa)L]; \quad p_n = \frac{\Lambda L p_{n-1} + 3\kappa L p_{n-3}}{n} \quad (1b)$$

Here Λ is the sum of the three independent rates, X and Y are independent Poisson variables, p_n is the probability of N equal to n , and negative subscripts denote zero. A direct convolution of the two

Poisson distributions provides an independent implementation. In a region i , the pipeline count is Poisson with mean λ_i plus κ times L , because each common event requests one regional unit. Dependence between the regional counts does not change their marginal expected fill measures.

2.2 Arrival and deadline qualified fill

The independent Poisson single-request and common-batch streams each sample the stationary prearrival pipeline distribution; the current batch is excluded. This arrival convention is essential to the fill formulas. Let C_j denote the central probability that N is no greater than j , with C_j equal to zero when j is negative. A single unit request is immediately filled with probability C at S minus one. A common batch receives the minimum of three and the available positive inventory. The expected number of immediately filled units in that batch equals the sum of C at S minus one, S minus two and S minus three. Using only the single unit probability for every common request would overstate batch fill.

Define e_i as one when central transport to region i is no longer than the deadline T and zero otherwise. Uniform allocation within a common batch makes each region receive one third of its expected filled units. The central transport qualified immediate fill fraction is:

$$F_{\text{central}} = \frac{\sum_i e_i \lambda_i C_{S-1} + (\kappa/3) (\sum_i e_i) (C_{S-1} + C_{S-2} + C_{S-3})}{\Lambda + 3\kappa} \quad (2)$$

For regional stock s_i and corresponding transport eligibility e_{Li} , the analogous fraction is:

$$F_{\text{regional}} = \frac{\sum_i e_{Li} (\lambda_i + \kappa) \Pr[\text{Poisson}((\lambda_i + \kappa)L) \leq s_i - 1]}{\Lambda + 3\kappa} \quad (3)$$

The denominator is the rate of requested units, Λ plus three κ , rather than the rate of arrival events. The metrics therefore weight regions by request volume. They describe a request that is immediately assigned stock and whose specified transport time meets T . A backordered request might nevertheless reach a site before its deadline if replenishment is imminent. Conversely an immediately delivered spare may not restore the generator in time. Thus neither equation is the probability of full restoration by T ; it is a deliberately restricted logistical endpoint.

2.3 Hypothetical scenarios and allocation search

Table 1 contains all operational inputs. Rates and times are selected scenarios, not estimates obtained from an employer or a supplier. The regions are abstract numbered nodes, not named cities. The budget is an integer count of identical spares; no monetary acquisition, obsolescence or warehouse operating cost is estimated.

Table 1
Declared scenario inputs

Input	Values	Unit
Independent rates by region	0.08, 0.12, 0.20	requests per day
Common event rate kappa	0, 0.02, 0.06	batches per day
Replenishment lead time L	3, 7, 14	days
Total stock budget S	6, 9, 12, 15	identical units
Transport deadline T	8, 24, 48	hours
Regional transport times	2, 4, 6	hours by region
Central transport times	4, 12, 30	hours by region

The full grid has 108 architecture comparisons. Every regional allocation contains at least one unit per region and sums exactly to S. All such allocations are enumerated and the one with greatest regional qualified fill is retained. This local stock requirement is an explicit architecture constraint; removing it can change the low budget optimum. A fixed tie order chooses the first maximizing allocation and is not an economic preference.

A second allocation exercise considers budgets from six to twenty one in increments of three and maximizes the minimum fill over four named scenarios formed by L equal to 3 or 14 days and kappa equal to zero or 0.06. This is finite scenario robustness. No bound over every intermediate demand rate is claimed for the regional allocation. For the central stock at a 48 hour deadline, increasing L or kappa decreases the modeled fill, and the highest values of both determine the lowest central value in the selected range.

2.4 Equivalent mean approximation and verification

The approximation replaces only the central outstanding count by a Poisson variable with mean $\Lambda + 3\kappa L$. Batch aware arrival weights and allocation are retained. This controlled comparison isolates a pipeline distribution error, rather than mixing it with an additional error that treats common arrivals as independent singles. The exact distribution has variance $\Lambda + 9\kappa L$, larger than its mean when kappa is positive.

Calculations use Python, NumPy 2.3.5 and SciPy 1.17.0 (Virtanen et al. 2020). The probability recursion is evaluated through 160 outstanding units, and mass sums are checked against one. A separate verifier recomputes each relevant cumulative probability by convolution and reconstructs all 108 central comparison values. Three Monte Carlo pipeline checks use 600000 independent samples per case. An additional event history calculation independently generates Poisson arrival times, maintains the outstanding replenishment count, fills each single or common batch in order, and discards the first L days. It uses 200 independent histories of 5000 days. Online Resource 1 includes all inputs, scripts,

seeds, tables and outputs. Additional review checks recover ordinary Poisson demand when kappa is zero, zero-stock fill, the eligible-route ceiling at large stock and equal batch allocation by exhaustive priority permutations.

2.5 Delivery by the deadline including backorders

The extension sends every request after stock assignment, whether immediate or after waiting for a replenishment. Unit backorders are served in first come first served order, retaining an independent uniform permutation within each common batch. Each unit request places an order immediately, which arrives after L days. Transport starts when that unit is assigned a spare. Replenishment and transport are successive times, rather than both being called lead time.

For a tagged unit at rank j in its current batch, let W_j be stock waiting time. If $0 \leq w < L$, delivery from stock within w requires no more than S minus j prior demand units in the interval of length L minus w immediately before the batch. Earlier orders have then returned and later demands cannot displace the tagged request under FCFS. Thus the same compound Poisson cumulative probability C_n , now evaluated at a shorter interval, determines the waiting-time distribution:

$$G_j(w) = \Pr(W_j \leq w) = C_{S-j}(L - w), \quad 0 \leq w < L \quad (4a)$$

$$G_j(w) = 0 \text{ for } w < 0; \quad G_j(w) = 1 \text{ for } w \geq L \quad (4b)$$

At w at least L the tagged unit's own replacement has arrived, so the probability is one even at zero base stock or a batch rank larger than S. This boundary must not be computed as a negative-index CDF. At w below zero the probability is zero. Deterministic replenishment permits an atom at W equal to L for low-stock batch ranks. For a region i with transport t_i hours, define w_i equal to (T minus t_i) divided by 24 days and $G_j(w_i)$ as the probability in Eq. 4. Uniform batch priority yields the central delivery fraction:

$$D_{\text{central}} = \frac{\sum_i [\lambda_i G_1(w_i) + \frac{\kappa}{3} (G_1(w_i) + G_2(w_i) + G_3(w_i))] }{\Lambda + 3\kappa} \quad (5)$$

The regional counterpart sums (lambda i plus kappa) times the corresponding single-rank waiting probability and divides by total unit request rate. For $0 \leq w_i < L$, its prior pipeline is Poisson with mean (lambda i plus kappa)(L minus w_i). All regional allocations with at least one unit per region are enumerated again for this new endpoint. The same 108 scenario grid is used, and independent convolution checks every central value. Boundary checks recover zero waiting, impossible transport and the own-order arrival at L.

Two separately seeded chronology runs retain 200 and 1000 histories of 5000 days each after a seven day warmup. Supply assignment uses the timestamp of the request S unit positions earlier plus L, independently of the distribution formulas. An additional finite-path check compares every central waiting classification with the corresponding prior demand count and current batch rank. All Monte Carlo

outcomes, including individual interval misses, are retained. The endpoint excludes diagnosis, installation, technician delay and successful generator restart.

2.6 AI assistance in study preparation

OpenAI ChatGPT, accessed on 4 and 5 October 2026, assisted substantially with topic development, inventory formulation, software, numerical and simulation checks, literature discovery, drafting, figure preparation from computed outputs and Arabic translation. Simulations and numerical findings were executed by the supplied scripts and are not represented as observed transactions. AI is not an author and does not replace human accountability. Further assistance on 5 October 2026 covered journal-specific formatting, reference presentation and submission preparation.

3 Results

3.1 Central availability and transport qualification

At S equal to nine, L equal to seven days and κ equal to 0.02 per day, the optimized regional allocation is two, three and four units. With the 48 hour deadline, all central and regional transport routes qualify. The central fill is 0.981996 and the regional fill is 0.908955. At shorter deadlines, the regional routes remain eligible while some central routes do not, producing the reversal in Table 2 and Fig. 1.

Table 2
Deadline sensitivity for the nine unit example

Deadline hours	Regional qualified fill	Central qualified fill	Eligible central regions
8	0.908955	0.213126	First only
24	0.908955	0.511906	First and second
48	0.908955	0.981996	All three

As central stock grows, its qualified fill at an eight hour deadline remains limited by transport eligibility. Its upper bound in this example is the first region's share of total requests, 0.10 divided by 0.46, or 0.217391. Additional stock cannot overcome that deterministic route constraint. This statement concerns the defined immediate fill metric; it does not assess alternative expedited transport, substitution or forward stocking.

3.2 Common batches and stock thresholds

At fourteen day replenishment and a 48 hour deadline, the exact central stock needed for fill of at least 0.98 rises from twelve units with no common demand to fourteen units at κ equal to 0.02 and eighteen at κ equal to 0.06. Table 3 compares these thresholds with the controlled equivalent mean Poisson approximation. Figure 2 displays the resulting fill curves.

Table 3
Central stock for a 0.98 fill target at a 48 hour deadline

Common batch rate per day	Exact stock	Approximation stock	Actual exact fill at approximation stock
0.00	12	12	0.987513
0.02	14	13	0.967896
0.06	18	16	0.959196

For kappa equal to 0.06, using the approximation stock is short by two units and misses the target by 0.020804 in fill fraction. The comparison does not estimate a monetary stocking saving or an observed rate of service failure. It shows that matching the mean outstanding count is insufficient for this high service target under the declared common batches.

3.3 Finite scenario allocation and independent checks

At a nine unit budget, the allocation maximizing nominal regional fill is two, three and four. Its lowest value over the four named supply scenarios is 0.494984. Maximizing the minimum instead selects one, three and five units and raises that lowest value to 0.514424. The difference is 0.019440 in the modeled fraction. At budgets of twelve, fifteen and eighteen units, the nominal and finite scenario allocations coincide. Thus robustness does not always require a different allocation in this example.

The probability recursion and independent mass convolution differ by at most 1.94 times ten to the minus sixteen. Reconstructing the 108 central fill comparisons directly from cumulative convolutions gives a maximum discrepancy of 5.55 times ten to the minus sixteen. The truncated probability mass error is below ten to the minus twelve in all evaluated demand and lead time cases.

The three independent pipeline sampling comparisons have standardized differences of minus 1.93, 2.24 and minus 1.12 standard errors. One exact value lies outside its nominal 95 percent sampling interval; all three discrepancies are within three standard errors, and no sample was replaced or rerun to select a favorable result. The independent event history calculation retains 460914 requested units after warmup and gives fill 0.982012, with standard error 0.000396, compared with exact fill 0.981996. Its discrepancy is 0.041 standard errors. These checks validate the implementation of the assumed pipeline and batch rules, not an actual supply network.

3.4 Stock waiting changes the deadline endpoint

Table 4 and Fig. 3 compare qualified immediate fill with delivery including waiting at S equal to nine, L equal to seven and kappa equal to 0.02. The optimized regional stock remains two, three and four in these three examples. At forty eight hours the regional delivery fraction exceeds its immediate fill by 0.045264, or 4.53 percentage points; the central increase is 0.009719. The eight hour central

disadvantage and forty eight hour central advantage both persist, so including wait does not reverse the architecture ordering in this example.

Table 4
Immediate fill and delivery including stock waiting

Deadline hours	Central immediate fill	Central delivery	Regional immediate fill	Regional delivery
8	0.213126	0.213525	0.908955	0.913232
24	0.511906	0.515113	0.908955	0.931202
48	0.981996	0.991715	0.908955	0.954219

[Insert Fig. 3 here]

All 108 central delivery values agree with independent convolution within 3.33 times ten to the minus sixteen. Forty boundary checks pass. The new simulations retain 459740 and 2297879 requested units. Finite-path waiting identities agree for 8272857 central request/deadline classifications. In the first run, the central twenty four hour estimate is 0.517100 against 0.515113 analytically, a discrepancy of 3.09 standard errors. Its estimated 95 percent interval misses the exact value. An independently seeded larger replication gives 0.515774, or 2.20 standard errors above the exact value, and also misses its individual 95 percent interval. The other five endpoints in each run contain the exact values in their estimated intervals. The original and replication are both retained; they are not selected to force interval coverage. These sampling discrepancies limit claims of Monte Carlo precision. The exact convolution and finite-path identities provide separate implementation checks and do not validate the field assumptions.

4 Discussion

The central versus regional comparison depends on the endpoint. A high unconstrained immediate fill may coexist with a low transport qualified fill if a large demand region is too far from the central warehouse under the assumed service deadline. This is consistent with the importance of nonzero transshipment time in pooling research (Wong et al. 2005), although the present architectures omit lateral transfers. It also distinguishes stock availability from the complete generator restoration process, which includes identification, compatibility, mobilization, installation and testing.

The compound Poisson construction separates common events from independent failures. Such dependence can increase tail exposure even with the same mean pipeline count. Spatial correlation has also been examined in inventory studies (Yan et al. 2009), but this paper uses a particular three unit batch mechanism and does not claim that all correlated demand has the same effect. Batch sizes, unequal regional exposure and demand concentration after environmental events require different input distributions.

The one for one fixed lead time convention is deliberately restrictive. Supplier shortages, repair capacity, order cancellation, lost sales, variable lead time, transportation reliability, multiple part numbers and compatibility constraints can materially change the result. Equal unit cost is imposed by the identical spare budget and cannot be generalized to a heterogeneous catalog. An optimization that minimizes expected downtime or total cost would require additional service and economic information. The finite scenario allocation search should not be described as a guarantee over every possible rate or lead time.

To calibrate the model, a service network would need dated part requests and their relationships, replenishment completion times, immediate allocations, transport times and whether each arrival actually enabled restoration. No such records were used. The results identify the kinds of data needed to decide whether a central warehouse, regional stock or a mixed architecture merits more detailed analysis.

The waiting extension resolves the distinction between qualified immediate fill and delivery by T under the declared supply model. The historical delivery-time objective in (Reijnen et al. 2009) prevents treating that objective itself as novel. FCFS priority, deterministic supply and unlimited capacity are strong assumptions: expediting, random lead times, deadline priority or congestion can change the waiting distribution. Delivery of a part remains different from successful repair.

5 Conclusions

An exact batch aware stock calculation separates the effect of pooling from transport eligibility. In the nine unit example, central fill exceeds regional fill when all routes qualify but falls sharply below it at an eight hour deadline. Under common demand and fourteen day replenishment, an equivalent mean Poisson pipeline understates the stock needed for a 0.98 target. Complete allocation enumeration, probability convolution and independent simulations support the results. Practical use requires calibrated demand, replenishment and transport data and a metric that includes the full restoration process when that is the decision objective.

Including stock waiting raises the selected forty eight hour delivery fractions to 0.9917 centrally and 0.9542 regionally. The exact extension therefore supports evaluating deadline service directly when backorders are possible; the immediate-fill metric alone is insufficient.

6 Nomenclature

Symbol	Meaning	Unit
λ i Lambda	Regional independent rate and its total	requests per day
κ	Common three unit event rate	batches per day
L	Deterministic replenishment lead time	days
S and s_i	Central budget and regional stocks	identical units
$N \times Y$	Outstanding units independent requests and common events	counts
p_n and C_j	Pipeline probability and cumulative probability	dimensionless
T	Deadline for transport qualification or delivery as specified	hours
e_i and e_{Li}	Central and regional route eligibility	zero or one
F_c and F_l	Central and regional qualified immediate fill	dimensionless
W_j w_i G_j	Tagged stock waiting available waiting time and its cumulative probability	days days and dimensionless
$D_{central}$	Central delivery fraction including stock waiting and transport	dimensionless

Declarations

Competing Interests

The author is employed at High Performance Industry (HPI). This employment is disclosed as relevant professional context. No company or client operational records, proprietary equipment data or financial data were used. The author has declared no other competing interests.

Author Contribution

M.A.A.M. supplied the industrial context and research direction. Substantial OpenAI ChatGPT assistance covered model formulation, software, formal calculations, verification scripts, visualization, original drafting and translation, as disclosed in the manuscript. M.A.A.M. is responsible for scientific review, interpretation and the final manuscript.

Data Availability

The declared scenarios, allocation tables, probability calculations, simulation seeds, code and checks are supplied in Online Resource 1. There are no company demand measurements.

References

1. Reijnen IC, Tan T, van Houtum GJ (2009) Inventory planning for spare parts networks with delivery time requirements. BETA Working Paper 280, Eindhoven University of Technology. <https://pure.tue.nl/ws/files/2963049/652106.pdf>
2. Sherbrooke CC (1968) METRIC: A Multi-Echelon Technique for Recoverable Item Control. Operations Research 16:122–141. <https://doi.org/10.1287/opre.16.1.122>
3. Virtanen P, Gommers R, Oliphant TE et al (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods 17:261–272. <https://doi.org/10.1038/s41592-019-0686-2>
4. Wong H, Cattrysse D, Van Oudheusden D (2005) Inventory pooling of repairable spare parts with non-zero lateral transshipment time and delayed lateral transshipments. European Journal of Operational Research 165:207–218. <https://doi.org/10.1016/j.ejor.2003.09.031>
5. Yan XS, Robb DJ, Silver EA (2009) Inventory performance under pack size constraints and spatially-correlated demand. International Journal of Production Economics 117:330–337. <https://doi.org/10.1016/j.ijpe.2008.11.006>

Figures




Figure 1

Central and optimized regional transport qualified immediate fill at two deadlines, L equal to seven days and κ equal to 0.02 per day. All stocks use the same integer unit budget. The central eight hour ceiling reflects route eligibility rather than a shortage of stock




Figure 2

Central fill at fourteen day replenishment for three common batch rates. Solid lines use the exact compound Poisson pipeline and dashed lines the equivalent mean Poisson pipeline, retaining batch aware arrivals. The dotted line marks the assumed 0.98 target




Figure 3

Deadline service with and without waiting for imminent replenishment. The regional comparison keeps stock at two, three and four. Delivery includes stock waiting and transport; it excludes installation and restoration

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ESM1.zip](#)