# Three Clocks and Two Ages: Days of Cover, Receipt Ageing versus Consumption Ageing, and the Denominator Problem in Inventory Analytics

## Architecture, Computational Methods, and a Modelled Scenario for MileSoft Inventory Control

MileSoft Engineering Research Group

MileSoft Software Technologies

https://milesoft.net/

Corresponding authors: *Deepak Daryani* (deepakdaryani@milesoft.net), *Tanuj Daryani* (tanujdaryani@milesoft.net)

Working paper — Version 1.0 — 23 August 2026

Permanent link: https://milesoft.net/research/products/inventory-control

### Abstract

An enterprise resource planning system answers how much stock is on hand. That is a level, and a level on its own supports no decision: whether 8,256 lakh of stock is too much depends entirely on how fast it is consumed, which parts it sits in, and how long the individual pieces have been standing still.

This paper presents the architecture and computational methods of MileSoft Inventory Control, an analytical layer that synchronises daily receipts, issues and stock from a host ERP into a governed store and reports inventory as days of cover rather than as value. We show that days of cover is Little's Law applied to a stockroom, which makes explicit the assumption every such figure carries and the conditions under which it may be compared between plants.

Three results follow from the system's own reporting structure. Days of cover is not one number but three — a derived figure from a planning norm, a historical average, and today's instantaneous position — and the three diverge precisely when the stock is worst. Ageing must be measured on two clocks, receipt age and consumption age, because a part with recent receipts and no consumption for three years is dead stock that a receipt clock reports as fresh. And a non-moving provision must be capped at the quantity actually on hand, or it books a write-down against stock that is not there.

The classification layer implements four value bands against three variability bands, one band more than the textbook ABC scheme, and we set out what the additional band buys. This product appears in no published case study; Section 5 is therefore a modelled scenario with every assumption printed, and no deployment claim is made anywhere in this paper.

**Keywords:** Inventory control software; ABC-XYZ classification; Days of cover; Inventory ageing analysis; Non-moving stock provision; SAP inventory dashboard; Safety stock and reorder point; Intermittent demand forecasting; Slow-moving inventory; Working capital in manufacturing; Buyer performance analytics; Little's Law inventory

# 1 Introduction

Most enterprise resource planning systems answer one inventory question well: how much is on hand. Buyers, planners and operations leads need five more answers every day — which parts are overstocked weighted by value contribution, which slow movers are tying up capital, what is drifting in demand pattern, which lots are ageing out, and which suppliers are silently inflating safety stock.

None of those five is a level. Each is a rate, a ratio or an age, and each therefore needs a denominator that the transaction record does not supply on its own. That is the gap this class of system fills, and it is why an inventory analytics layer sits above an ERP rather than replacing it.

The product's own material puts the cost of not filling it plainly: a buyer running classification in a monthly spreadsheet is four weeks behind every demand shift. This paper is about what has to be computed, and computed correctly, to close that gap.

## 1.1 The denominator problem

Inventory value is meaningless without a rate to divide it by, and every dispute about whether a plant holds too much stock is in practice a dispute about which rate. Last month's issues? The rolling twelve-month average? The planning norm the buyer was given?

Each is defensible and they give different answers — sometimes by an order of magnitude. The response taken here, and implemented in the product, is not to pick one but to report several and treat their divergence as information. Section 4.2 states the three and Section 6.1 argues that the gap between them is the most useful number on the screen.

> A single days-of-cover figure is a claim with a hidden argument inside it. Three figures side by side put the argument on the screen where a buyer can see it.

## 1.2 Contributions

1. Days of cover derived as Little's Law applied to a stockroom, with the comparison condition that follows from it (Section 4.1).
2. Three days-of-cover figures — derived, average and instantaneous — and a reading of their divergence as diagnostic (Section 4.2).
3. Two ageing clocks, receipt age and consumption age, and the class of dead stock that only the second detects (Section 4.3).

4. A non-moving provision rule capped at on-hand quantity, stated as a minimum rather than left to reconciliation (Section 4.4).
5. An eight-dimension capability reference framework for inventory analytics (Section 6.3).

# 2 Background and Related Work

Inventory theory is old, well settled and largely ignored in practice — not because practitioners disagree with it, but because its standard results assume conditions that a real part master violates part by part.

## 2.1 What the classical results assume

Harris gave the economic order quantity in 1913: the lot size balancing set-up cost against the interest on capital tied up. It assumes constant demand, instant replenishment and a known ordering cost, and it remains the correct answer whenever those hold.

Dickie's 1951 ABC analysis is the first move away from treating parts alike. Applying Pareto and Lorenz to stock value, it observes that a small share of part numbers carries most of the value and argues that control effort should follow value rather than count. Every classification scheme in this field descends from that paper.

The XYZ axis adds demand variability to Dickie's value axis. The combined matrix is what makes differentiated policy possible: a high-value stable part and a high-value erratic part need the same attention and opposite safety stocks.

## 2.2 Why a moving average is the wrong forecast for most part numbers

In an automotive part master, most part numbers have intermittent demand — long runs of zero broken by occasional orders. A simple or exponentially smoothed average over such a series is biased, because it spreads a lumpy quantity evenly across periods that saw nothing.

Croston addressed this in 1972 by forecasting demand size and inter-arrival interval as separate series and combining them. Syntetos and Boylan showed in 2005 that Croston's estimator is itself biased and gave a correction, evaluating both against three thousand real intermittent series drawn from the automotive industry — which is precisely the population a plant part master is drawn from.

> This matters for classification, not only forecasting. Variability computed with an estimator that is wrong for intermittent series will put parts in the wrong XYZ band, and the differentiated policy that follows will be differentiated in the wrong direction.

### 2.3 Where an analytical layer belongs

IEC 62264 separates the enterprise system of record from operations-management activities, of which inventory management is one. The separation is not bureaucratic: it says the transaction history and the analysis over it are different responsibilities with a defined interface between them.

The product follows this. It layers on top of the ERP without replacing or replicating its transaction history, leaving the ERP as the system of record — which also means every figure in this paper is downstream of data the plant already trusts, and inherits its quality.

## 3 System Overview

The architecture is deliberately modest: a background service on a server uploads daily receipts, issues and stock from the host ERP into a centralised database, and authorised users reach the dashboard over the local network. Everything analytical happens on that store.

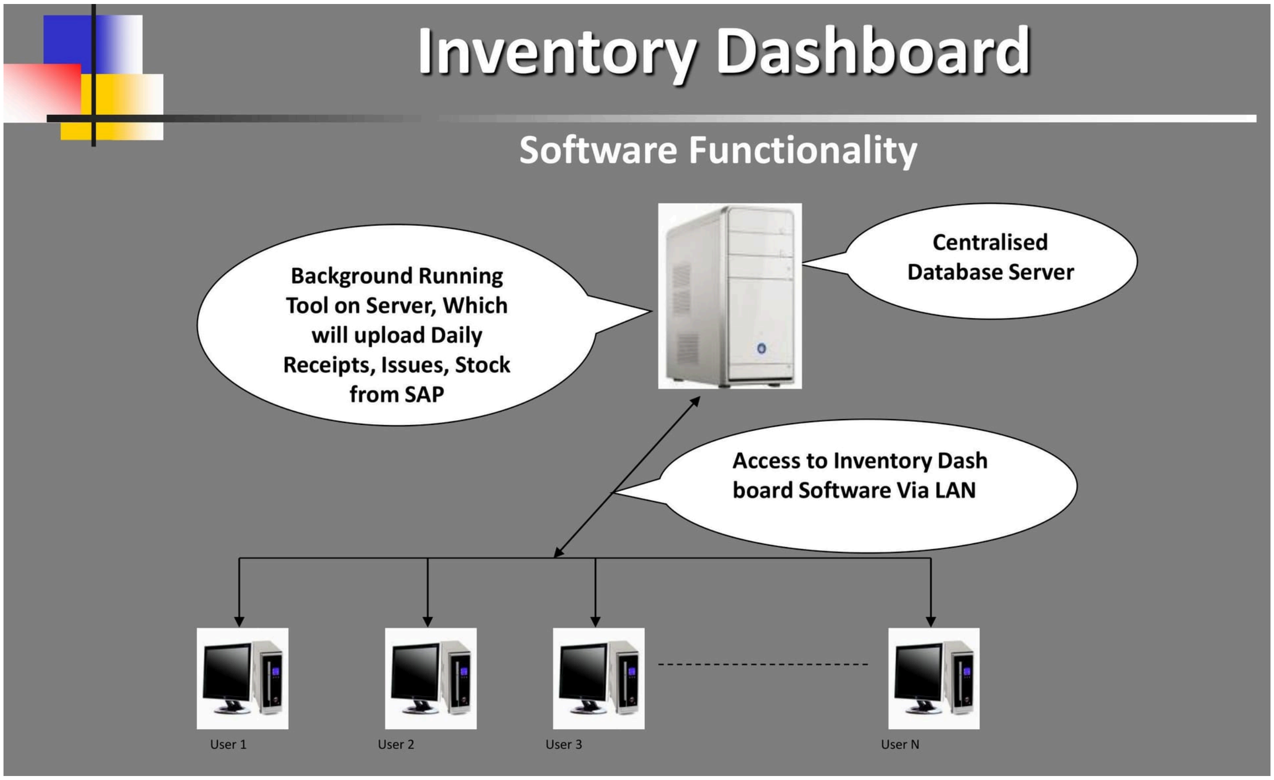


**Figure 1.** Product capture: the deployment architecture. A background service on the server uploads daily receipts, issues and stock from SAP into a centralised database, which multiple users reach over the local network. The cadence is daily batch, not streaming — a property Section 6.2 takes seriously.

### 3.1 Daily batch, and saying so

The refresh cadence is a daily background upload, and the dashboard states the date it last ran. This paper reports that plainly because it bounds what the system can be used for.

A daily cadence is right for every question this system answers. Classification, ageing, trend and provisioning are all decisions on a weekly-to-monthly rhythm, and a figure recomputed hourly would not change any of them. It is the wrong cadence for shop-floor stock confirmation, which is a different product's job.

> The honest description is near-real-time analytics over a daily extract, not real-time inventory. A buyer choosing between suppliers on the strength of the word real-time should ask what is refreshed, and how often.

### 3.2 The dimensions everything is sliced by

Every report in the system is the same underlying position cut along one of a small set of dimensions, which is what allows a question to be pursued from plant level to a single part in a few steps.

**Table 1.** The reporting dimensions, and the question each answers.

| Dimension | The question it answers |
|---|---|
| Plant | Are we above or below target overall, and in which direction are we moving? |
| Procurement class | Which class of part carries the overhang — domestic, group, consumption-driven, export-peculiar or imported? |
| Buyer and manager | Whose portfolio is drifting, and against what target? |
| Project | What is allocated to a given programme, at what value? |
| Model | How is stock distributed across product models? |
| Part | Consumption pattern, stock level and last transaction for one part number |

The procurement class dimension is the one least often found elsewhere and the most useful. Imported and export-peculiar parts have structurally longer lead times and structurally different day targets, and averaging them into a plant figure conceals both.

### 3.3 The part master as the join

A central part master holds the attributes every analysis joins on — group, procurement type, buyer, manager, project and model — searchable on several parameters. Without it the dimensions above would not exist, because the ERP's transaction record carries the movement but not the classification.

This is the least glamorous component and the one that decides whether a deployment succeeds. An analytical layer over a part master with a third of its buyer assignments blank produces reports that are correct and useless.

# 4 Computational Methods

Four computations carry the system. The first two decide whether a plant is holding too much; the third and fourth decide what to do about the part of it that has stopped moving.

## 4.1 Days of cover is Little's Law

Inventory expressed in days is not a convention. It is the flow time of the stockroom, and it obeys the same relation between quantity, rate and time that governs any queueing system in steady state.

$$\text{Little's Law: } L = \text{lambda x } W, \text{ so } W = L \,/\, \text{lambda. Applied to stock: days of cover } d = \text{stock value} \,/\, \text{mean daily consumption value} \quad \text{(doc)}$$

L is the quantity held, lambda the consumption rate and W the average time a unit spends in stock. The identity is exact in steady state, which is the assumption every days-of-cover figure quietly carries.

Two consequences follow immediately. Days of cover is comparable between plants only where the consumption rate used as the denominator was measured over a window in which demand did not shift — so a figure computed across a model changeover means nothing. And reducing stock without reducing days of cover is not an improvement: it is a proportional fall in consumption, which is usually worse news than the stock was.

> The plant target in the product capture is 13.50 days. What that target asserts is that a rupee of material should spend under a fortnight in the building — a statement about flow, not about a balance sheet.

## 4.2 Three figures, and the diagnostic in their gap

The system reports three days-of-cover figures per procurement class rather than one, because the choice of denominator in Equation (doc) is not settled and pretending otherwise loses information.

$$\text{derived days} = \text{stock} \,/\, \text{planning-norm consumption; average days} = \text{mean over the period of stock} \,/\, \text{consumption; today's days} = \text{current stock} \,/\, \text{recent mean consumption} \quad \text{(three)}$$

c(norm) is the consumption rate the planning process assumed, c-bar(t) the rate actually observed in period t, and c-bar(recent) the rate over the trailing window. The three coincide when the plan matched reality and the position is stable.

Their divergence is the diagnostic. Derived far below average means the planning norm assumed a consumption rate the plant never achieved. Average far above today's means the position has recently improved and the historical figure is stale. Today's far above both means something has just stopped moving, and it is the earliest of the three to say so.

The product capture in Figure (dashboard) shows this in a live position: an imported-parts class reading under eight derived days against an average approaching three thousand. No single figure could carry that, and the gap — not either number — is what identifies where to look.

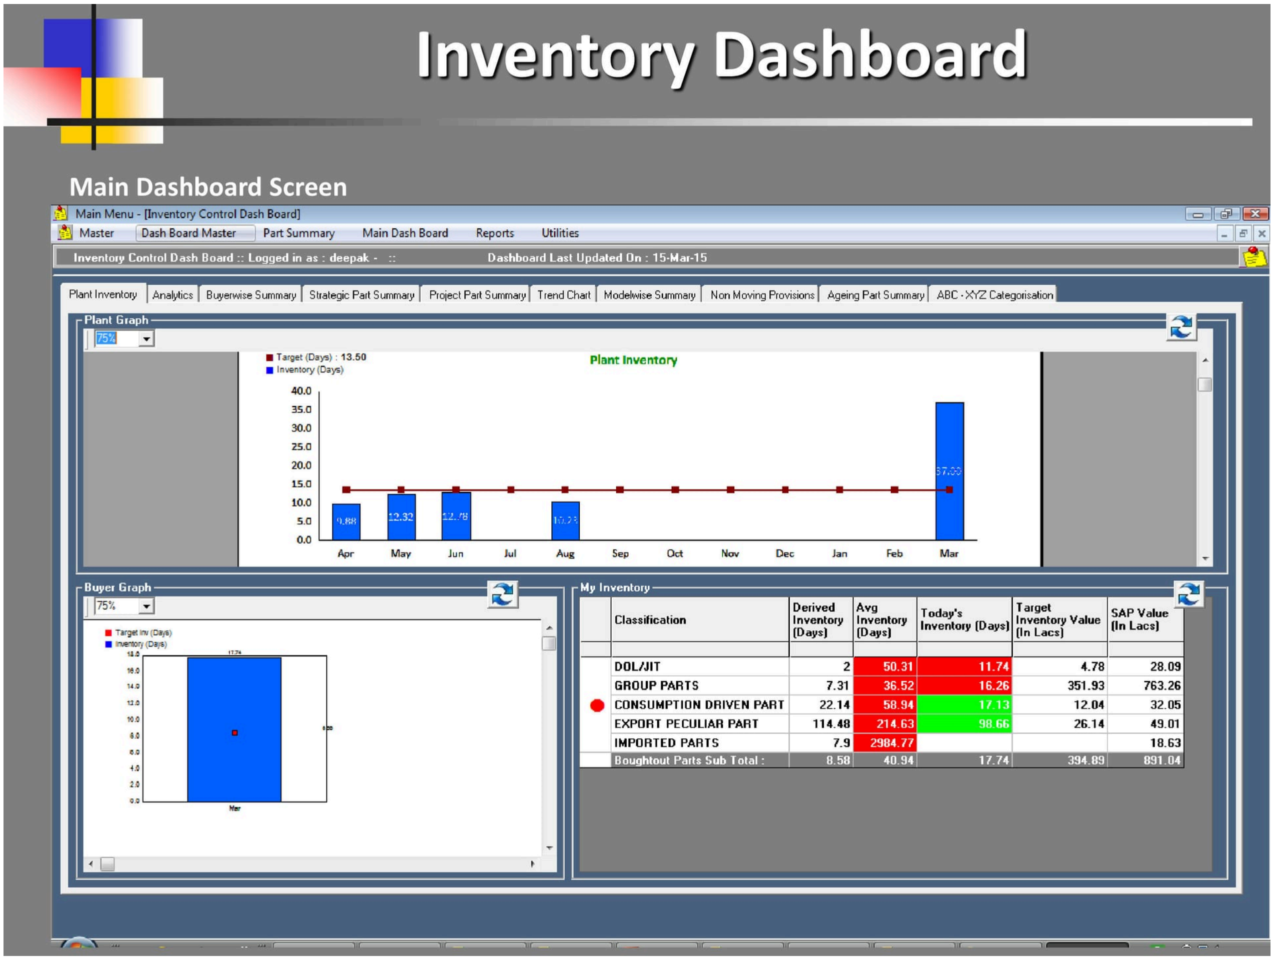


**Figure 2.** Product capture: the plant dashboard. The upper chart plots monthly inventory days against a target line of 13.50 days. The lower table reports derived, average and today's inventory days per procurement class — domestic and just-in-time, group, consumption-driven, export-peculiar and imported — beside target and actual value, with conditional shading where a class is outside tolerance.

## 4.3 Two ageing clocks

Ageing is usually reported as one number, and that number is almost always receipt age. It is the wrong one on its own, and the reason is a class of stock it cannot see.

$$\textit{receipt age of part } p = \textit{now} - \textit{date of its last receipt; consumption age of part } p = \textit{now} - \textit{date of its last issue} \qquad \text{(ages)}$$

both are reported per part, and bracketed into 0-30, 31-60, 61-90 and above-90-day bands. A part is dead stock when consumption age is large, regardless of receipt age.

The failure mode is a part still on an automatic replenishment schedule that nothing consumes. Receipts keep arriving, so receipt age stays small and the part never appears on an ageing report. Consumption age grows without bound. The product capture in Figure (ageing) shows a live row of exactly this shape: last receipt seven days ago, last consumption 1,328 days ago.

Three and a half years without an issue, and a receipt-age report would file that part as under thirty days old. The replenishment schedule that keeps it looking fresh is the same schedule that is quietly buying more of it.

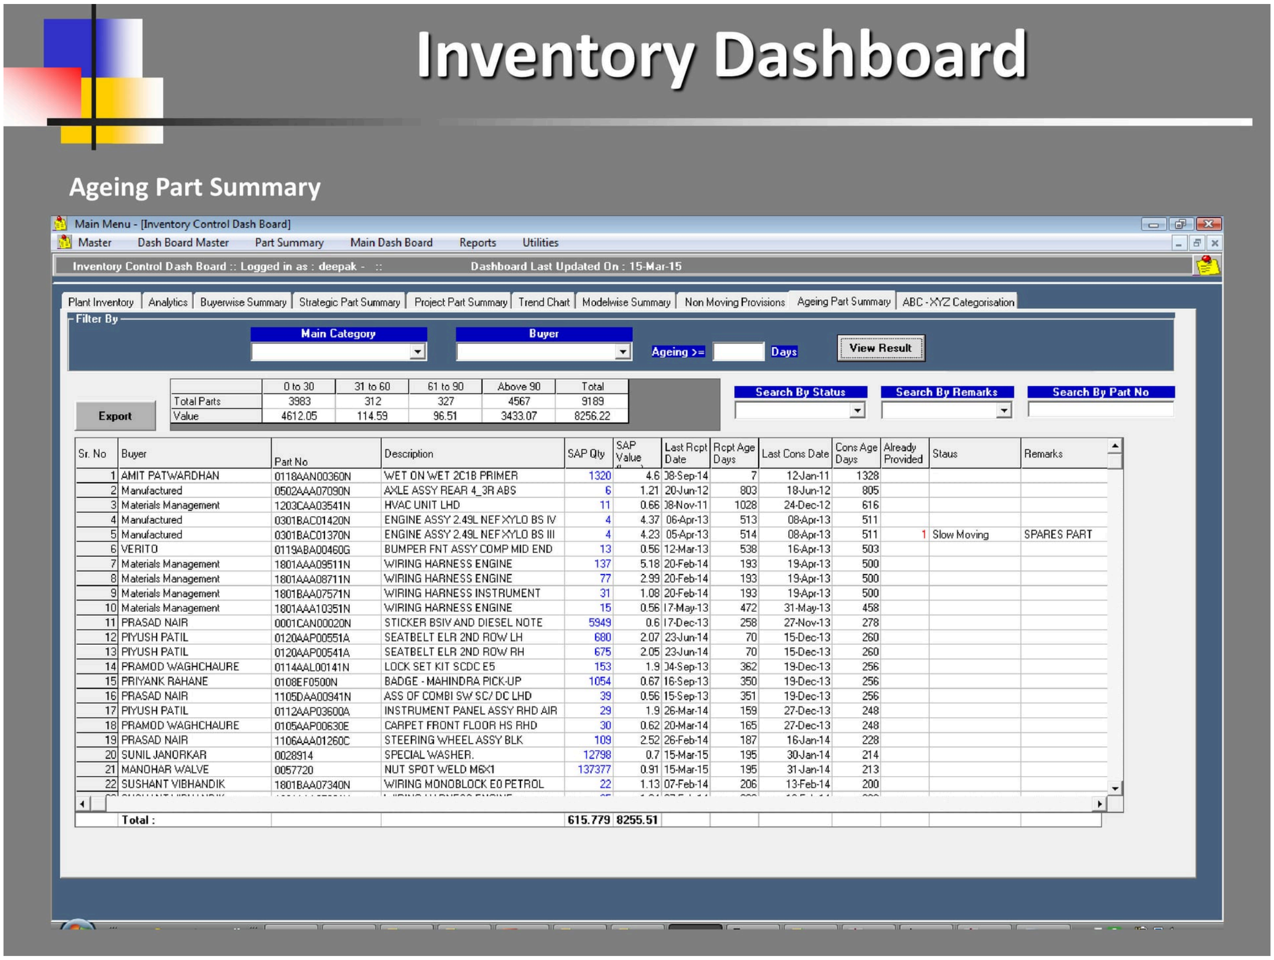


**Figure 3.** Product capture: the ageing part summary. Band totals across 0-30, 31-60, 61-90 and above-90 days appear above a part-level grid carrying both clocks — last receipt date with receipt age in days, and last consumption date with consumption age in days — alongside buyer, quantity, value, status and remarks.

Reporting both clocks also fixes the reverse error. A part received in bulk two years ago and consumed steadily since has a large receipt age and a small consumption age; it is working stock with an unlucky purchase history, not dead stock, and a receipt-age report would flag it for disposition wrongly.

## 4.4 The non-moving provision, and the cap that keeps it honest

Once a part is identified as non-moving, an accounting provision follows. The quantity a policy calls for and the quantity the system of record says exists are different numbers, and the difference matters.

$$\textit{non-moving quantity} = \min(\textit{provision quantity called for, quantity the ERP holds}), \qquad \text{(provision)}$$
$$\textit{provision value} = \textit{sum over parts of that quantity times unit value}$$

q(prov) comes from the ageing policy applied to the part, q(erp) from the daily extract, and v(p) is the unit value carried in the host system.

The minimum is the whole point. Where the policy calls for a provision against 219 pieces and the ERP holds none, the provision is zero — the stock was already scrapped or consumed, and booking against it would double-count a write-down that has happened. Where the ERP holds more than the policy covers, only the policy quantity is provided for.

Stating this as a computation rather than a spreadsheet convention is a small thing that survives audit. The provision is reproducible from two stored quantities, and the reason any given part contributes nothing is visible in the same row.

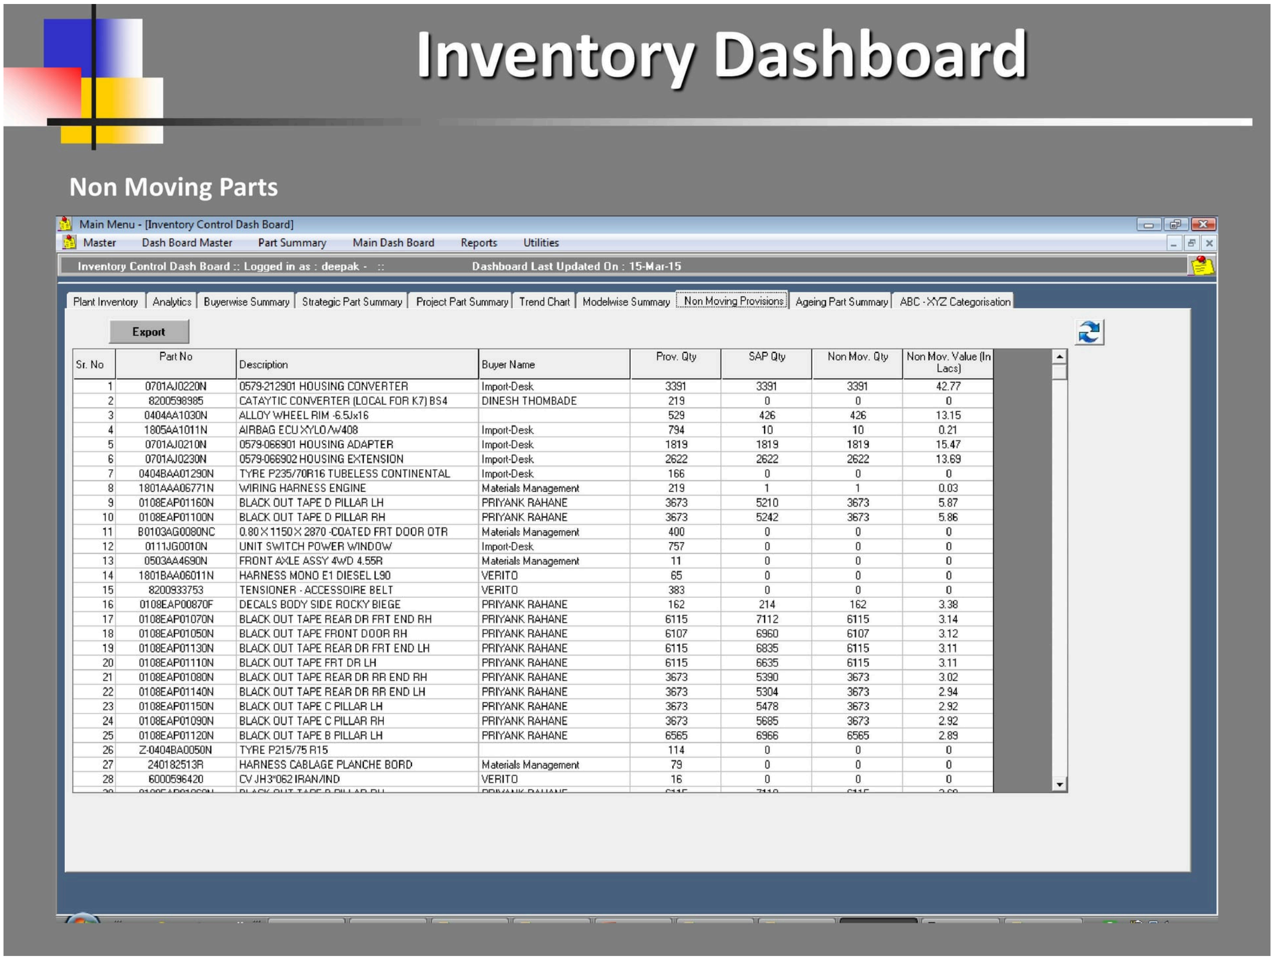


| Sr. No | Part No | Description | Buyer Name | Prov. Qty | SAP Qty | Non Mov. Qty | Non Mov. Value (In Lacs) |
|---|---|---|---|---|---|---|---|
| 1 | 0701AJ0220N | 0579-212901 HOUSING CONVERTER | Import-Desk | 3391 | 3391 | 3391 | 42.77 |
| 2 | 8200598985 | CATAYTIC CONVERTER (LOCAL FOR K7) BS4 | DINESH THOMBADE | 219 | 0 | 0 | 0 |
| 3 | 0404AA1030N | ALLOY WHEEL RIM -6.5Jx16 | | 529 | 426 | 426 | 13.15 |
| 4 | 1805AA1011N | AIRBAG ECU XYLO/W408 | Import-Desk | 794 | 10 | 10 | 0.21 |
| 5 | 0701AJ0210N | 0579-066901 HOUSING ADAPTER | Import-Desk | 1819 | 1819 | 1819 | 15.47 |
| 6 | 0701AJ0230N | 0579-066902 HOUSING EXTENSION | Import-Desk | 2622 | 2622 | 2622 | 13.69 |
| 7 | 0404BAA01290N | TYRE P235/70R16 TUBELESS CONTINENTAL | Import-Desk | 166 | 0 | 0 | 0 |
| 8 | 1801AAA06771N | WIRING HARNESS ENGINE | Materials Management | 219 | 1 | 1 | 0.03 |
| 9 | 0108EAP01160N | BLACK OUT TAPE D PILLAR LH | PRIYANK RAHANE | 3673 | 5210 | 3673 | 5.87 |
| 10 | 0108EAP01100N | BLACK OUT TAPE D PILLAR RH | PRIYANK RAHANE | 3673 | 5242 | 3673 | 5.86 |
| 11 | B0103AG0080NC | 0.80 X 1150 X 2870 -COATED FRT DOOR OTR | Materials Management | 400 | 0 | 0 | 0 |
| 12 | 0111JG0010N | UNIT SWITCH POWER WINDOW | Import-Desk | 757 | 0 | 0 | 0 |
| 13 | 0503AA4690N | FRONT AXLE ASSY 4WD 4.55R | Materials Management | 11 | 0 | 0 | 0 |
| 14 | 1801BAA06011N | HARNESS MONO E1 DIESEL L90 | VERITO | 65 | 0 | 0 | 0 |
| 15 | 8200933753 | TENSIONER - ACCESSOIRE BELT | VERITO | 383 | 0 | 0 | 0 |
| 16 | 0108EAP00870F | DECALS BODY SIDE ROCKY BIEGE | PRIYANK RAHANE | 162 | 214 | 162 | 3.38 |
| 17 | 0108EAP01070N | BLACK OUT TAPE REAR DR FRT END RH | PRIYANK RAHANE | 6115 | 7112 | 6115 | 3.14 |
| 18 | 0108EAP01050N | BLACK OUT TAPE FRONT DOOR RH | PRIYANK RAHANE | 6107 | 6960 | 6107 | 3.12 |
| 19 | 0108EAP01130N | BLACK OUT TAPE REAR DR FRT END LH | PRIYANK RAHANE | 6115 | 6835 | 6115 | 3.11 |
| 20 | 0108EAP01110N | BLACK OUT TAPE FRT DR LH | PRIYANK RAHANE | 6115 | 6635 | 6115 | 3.11 |
| 21 | 0108EAP01080N | BLACK OUT TAPE REAR DR RR END RH | PRIYANK RAHANE | 3673 | 5390 | 3673 | 3.02 |
| 22 | 0108EAP01140N | BLACK OUT TAPE REAR DR RR END LH | PRIYANK RAHANE | 3673 | 5304 | 3673 | 2.94 |
| 23 | 0108EAP01150N | BLACK OUT TAPE C PILLAR LH | PRIYANK RAHANE | 3673 | 5478 | 3673 | 2.92 |
| 24 | 0108EAP01090N | BLACK OUT TAPE C PILLAR RH | PRIYANK RAHANE | 3673 | 5685 | 3673 | 2.92 |
| 25 | 0108EAP01120N | BLACK OUT TAPE B PILLAR LH | PRIYANK RAHANE | 6565 | 6966 | 6565 | 2.89 |
| 26 | Z-0404BA0050N | TYRE P215/75 R15 | | 114 | 0 | 0 | 0 |
| 27 | 240182513R | HARNESS CABLAGE PLANCHE BORD | Materials Management | 79 | 0 | 0 | 0 |
| 28 | 6000596420 | CV JH3*062 IRAN/IND | VERITO | 16 | 0 | 0 | 0 |

**Figure 4.** Product capture: the non-moving provisions report. Provision quantity, the quantity the host system holds, and the resulting non-moving quantity appear as separate columns, with the third the smaller of the first two — so rows where the stock no longer exists contribute nothing to the provision value.

# 5 Modelled Scenario

This product appears in no published MileSoft case study. There is therefore no field evidence to report, and this section makes no deployment claim. Everything below is a modelled scenario with its assumptions printed.

## 5.1 What can be stated from the product itself

Three things in this paper are observed properties of the delivered software rather than claims about outcomes, and they are separated here so a reader can tell which is which.

**Table 2.** Observed properties of the software, with where each was read.

| Property | Value | Where observed |
|---|---|---|
| Plant days-of-cover target | 13.50 days | Plant inventory screen |
| Classification matrix | four value bands against three variability bands | ABC-XYZ categorisation screen |
| Ageing bands and clocks | 0-30 / 31-60 / 61-90 / above 90, on receipt age and consumption age | Ageing part summary screen |
| Source refresh cadence | daily background upload of receipts, issues and stock | Architecture screen |
| Provision rule | non-moving quantity is the lesser of policy and on-hand | Non-moving provisions screen |

None of these is a performance claim. They are what the software does, and they are stated because a buyer comparing systems needs the structure before the benefits.

## 5.2 The fourth value band

The classification matrix carries four value bands — A+, B, C and D — against the three variability bands X, Y and Z. Dickie's scheme and every textbook since use three value bands, so the additional one is a deliberate choice worth examining.

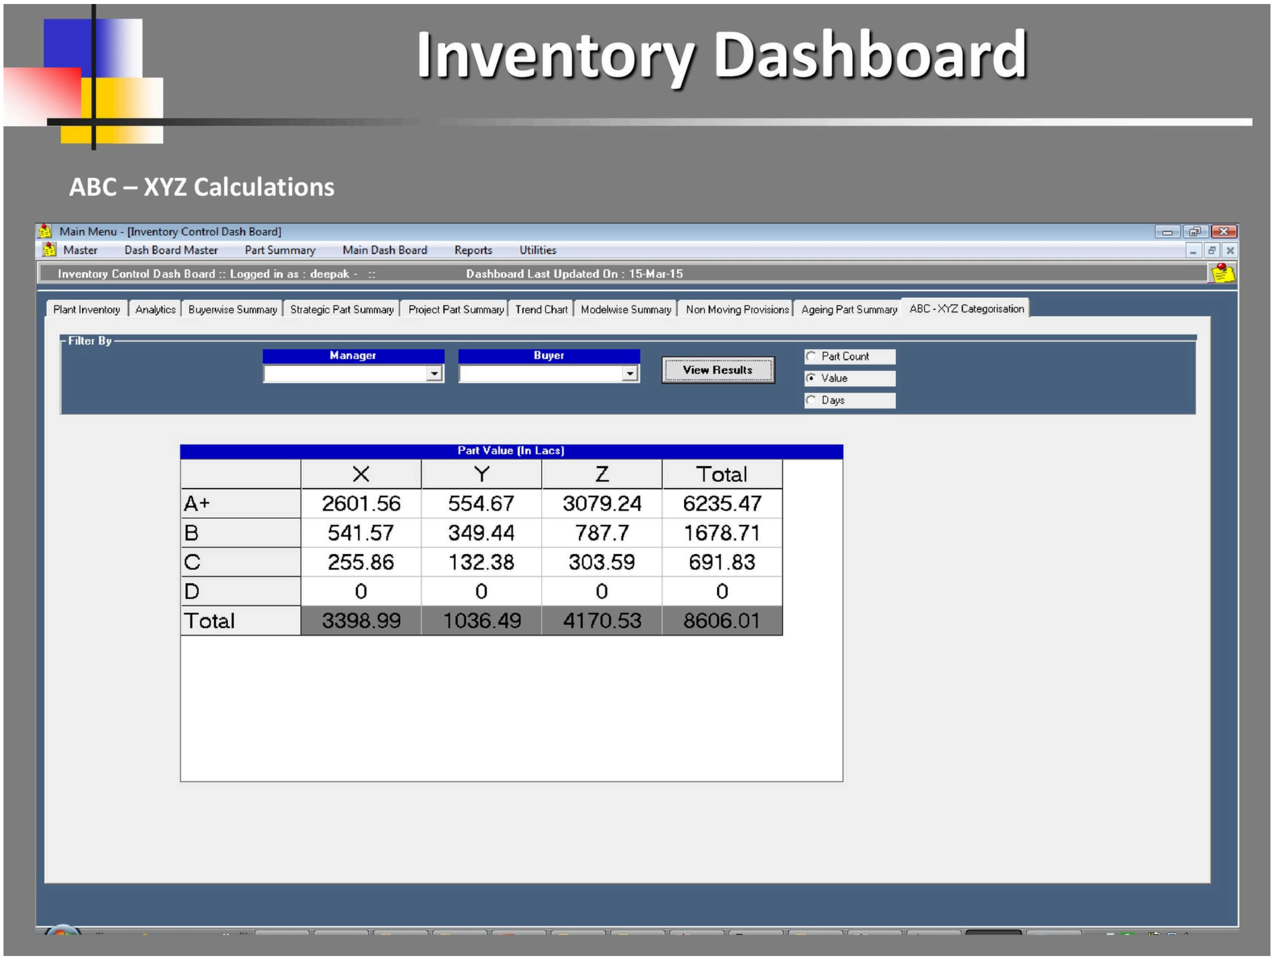


| | X | Y | Z | Total |
|---|---|---|---|---|
| A+ | 2601.56 | 554.67 | 3079.24 | 6235.47 |
| B | 541.57 | 349.44 | 787.7 | 1678.71 |
| C | 255.86 | 132.38 | 303.59 | 691.83 |
| D | 0 | 0 | 0 | 0 |
| Total | 3398.99 | 1036.49 | 4170.53 | 8606.01 |

**Figure 5.** Product capture: the ABC-XYZ categorisation screen. Four value bands are crossed with three variability bands, with the cell contents switchable between part count, value and days. The D band exists to hold parts that carry value but no consumption signal at all.

The role the D band plays follows from Section 2.2. A part with no consumption in the classification window has no computable variability, so placing it in the Z band asserts high variability where the truth is no signal at all. Separating it keeps the Z band meaning erratic rather than meaning unknown — and the parts that land in D are the same population Section 4.3 identifies by consumption age.

The cell contents switch between part count, value and days, which is the same denominator discipline as Section 4.2. A cell holding 3% of parts and 38% of value is a different management problem depending on which is being read, and both readings are one click apart.

## 5.3 Modelled effort avoided

The published return model for this product is stated as manual effort automated. Its assumptions are printed here in full so a reader can substitute their own.

**Table 3.** Modelled scenario — not a deployment result. Assumptions are the published defaults of the return-on-investment model for this product.

| Assumption | Value |
|---|---|
| People handling inventory data | 5 |
| Hours per person per day on that handling | 3 |
| Working days per year | 300 |
| Baseline effort | 4,500 hours per year |
| Share of that effort automated | 60% |
| Modelled hours avoided | 2,700 per year |

The 60% figure is an assumption, not a measurement, and a reader should treat it as the parameter to interrogate. What the model captures well is that the effort being replaced is extraction and reconciliation rather than judgement: the buyer still decides what to do about a slow mover, and the system removes the two hours spent finding it.

The product's own material frames the same point as a lag rather than a headcount: classification run monthly in a spreadsheet leaves a buyer four weeks behind every demand shift. That framing is the better one, because the cost of the lag falls on stock decisions rather than on labour, and stock is where the capital is.

# 6 Discussion

## 6.1 The gap between the three figures is the product

A conventional reading of a dashboard showing three days-of-cover figures per class is that two of them are redundant and the reporting should be tidied. That reading is wrong, and inverting it is the most useful thing in this paper.

The three agree when the plan matched reality and the position is stable, which is exactly when nobody needs to look. They separate when the planning norm was wrong, when the position has just changed, or when consumption has stopped — the three conditions a buyer is being paid to notice. The divergence is therefore an alarm that requires no threshold to be configured.

The same logic runs through the ageing report: two clocks that agree describe ordinary stock, and two clocks that disagree describe the two failure modes of Section 4.3. In both cases the information is in the difference, and a system that reports one number has thrown it away before the user sees the screen.

## 6.2 What a daily cadence can and cannot support

The daily extract of Section 3.1 sets a hard boundary on the product's claims, and it is worth drawing precisely rather than defensively.

**Table 4.** What a daily analytical extract supports, and what it does not.

| Supported | Not supported |
|---|---|
| Classification and re-banding | Shop-floor stock confirmation before a pick |
| Ageing and non-moving identification | Allocation of a specific piece to a specific order |
| Days-of-cover reporting against target | Preventing a stockout that develops within the day |
| Buyer and supplier performance review | Real-time available-to-promise |
| Trend and forecast input | Any decision whose window is shorter than the refresh |

Nothing in the right column is a deficiency; each belongs to a different layer, and the product's own material places physical movement with a warehouse management system and the transaction record with the ERP. The mistake to avoid is buying an analytical layer expecting a movement layer, which is a mistake the word real-time invites.

### 6.3 A capability reference framework for inventory analytics

**Table 5.** Capability reference framework. Each dimension separates inventory analytics from an inventory ledger, and each is answerable by demonstration against live plant data rather than by reading a specification.

| Dimension | Question the system must answer by demonstration |
|---|---|
| D1 Stated denominator | For any days-of-cover figure on screen, which consumption rate is underneath it? |
| D2 Multiple denominators | Are derived, average and instantaneous figures shown together, or reconciled into one? |
| D3 Two ageing clocks | Can the system list parts received this month that have not been consumed in two years? |
| D4 No-signal band | Where does a part with zero consumption in the window land — in the erratic band, or its own? |
| D5 Provision capped at on hand | Show a part whose policy provision exceeds host quantity. Which number is booked? |
| D6 Procurement class | Can imported and domestic parts carry different day targets, or is there one plant number? |
| D7 Intermittent-aware variability | Which estimator computes the variability band for a part with eleven months of zeros? |
| D8 Stated cadence | How old is the data on screen right now, and does the screen say so? |

D8 takes ten seconds and settles more than the rest. A system that displays its own extract timestamp has decided to be honest about a limitation; one that does not has decided the opposite.

### 6.4 Generalisability

The three-denominator argument generalises to any ratio KPI whose denominator is contested — which, following ISO 22400's insistence that a KPI carry its definition, is most of them. The two-clock argument generalises to any asset with independent arrival and departure processes: spare parts, tooling, returnable packaging.

What does not generalise is the 13.50-day target. It is a property of one plant's mix, lead times and service commitments, and a plant adopting it because it appeared in a screenshot would be adopting an answer to somebody else's question.

## 7 Threats to Validity and Limitations

1. No field evidence exists for this product. It appears in no published case study, so this paper reports no deployment outcome and the reader should treat the absence as informative.

2. The modelled 60% effort reduction is an assumption. It is the published default of the return model, carries no measurement, and is the parameter most worth replacing with a local estimate.
3. Every figure inherits the host system's data quality. The analytical layer neither validates nor corrects the ERP's transaction record; a wrong issue posting produces a wrong consumption age with no indication that anything is amiss.
4. Days of cover assumes steady state. Little's Law is exact only where the consumption rate is stationary over the averaging window, so figures spanning a model changeover or a demand step are not comparable with figures either side of it.
5. Classification quality depends on the variability estimator. Where variability is computed with an estimator unsuited to intermittent demand, parts will be banded wrongly and the differentiated policy will be applied in the wrong direction.
6. The part master is an unmeasured dependency. Buyer, project and procurement-class assignments are maintained outside the system, and every dimensional report is only as good as they are.
7. Daily cadence bounds the claims. Nothing whose decision window is shorter than the refresh interval can be supported, whatever the interface suggests.
8. The screenshots are one deployment at one date. They establish that the software computes these things; they establish nothing about typical values, and no figure read from them is presented here as a benchmark.

> The third limitation is the one that bites in practice. An analytical layer makes an ERP's data quality visible for the first time, and a plant should expect its first months of output to be an audit of its master data rather than of its stock.

## 8 Future Work

- A measured deployment. The most valuable next step for this product is a published case study, so that Section 5 can report an outcome instead of a model.
- Publishing the variability estimator. Naming which estimator computes the XYZ band — and applying a Croston or Syntetos-Boylan variant where demand is intermittent — would let a buyer judge whether the classification can be trusted for the parts it matters most on.
- Divergence as an alert rather than a reading. The gap between the three days-of-cover figures is already computed; surfacing it as a ranked exception list would turn Section 6.1 from an interpretation into a workflow.
- Consumption-age thresholds per procurement class. A ninety-day silence means something different for an imported part and a domestic one, and a single band boundary treats them alike.
- Closing the loop to procurement. Supplier reliability is described in the product's material as feeding the safety-stock model; measuring that feedback in a live deployment would make the claim checkable.

## 9 Conclusion

How much is on hand is a level, and a level supports no decision on its own. Everything a buyer needs is a ratio or an age, and every ratio needs a denominator that the transaction record does not supply.

Three results carry this paper. Days of cover is Little's Law applied to a stockroom, which makes explicit the stationarity assumption behind every such figure and the condition under which two plants' numbers may be compared at all. The three days-of-cover figures the system reports are not redundant — they agree when nothing is wrong and separate when something is, which makes their gap an alarm with no threshold to configure. And ageing needs two clocks, because a part replenished last week and unconsumed for three and a half years is dead stock that a receipt clock files as fresh.

The framework of Section 6.3 is offered as the durable contribution, and its last dimension takes ten seconds to test: ask the screen how old its data is, and see whether it answers.

# Appendix A Nomenclature

**Table 6.** Symbols and abbreviations used in this paper.

| Symbol / term | Meaning |
|---|---|
| L, lambda, W | Little's Law terms: quantity held, throughput rate, and average time in system |
| d | Days of cover — stock value divided by mean daily consumption value |
| d(der) | Derived days of cover, computed against the planning norm |
| d(avg) | Average days of cover over the reporting period |
| d(today) | Today's days of cover, against the trailing consumption rate |
| c(norm) | The consumption rate the planning process assumed |
| c-bar | Observed mean consumption rate over the window |
| a(rcpt) | Receipt age — days since the part's last receipt |
| a(cons) | Consumption age — days since the part's last issue |
| q(prov) | Provision quantity called for by the ageing policy |
| q(erp) | Quantity the host system of record holds |
| q(nm) | Non-moving quantity — the lesser of the two above |
| ABC / A+BCD | Value bands; this product implements four rather than the textbook three |
| XYZ | Demand-variability bands: X stable, Y moderate, Z erratic |
| Lakh | Indian numbering unit of 100,000, the unit values are reported in on screen |

# Appendix B Worked Numerical Examples

## Appendix B.1 Reading three days-of-cover figures

A procurement class holds stock valued at 763 lakh. The planning norm assumed consumption of 104 lakh per day; the trailing twelve-month observed mean was 20.9 lakh per day; the last thirty days averaged 46.9 lakh per day.

Applying Equation (three): derived days = 763 / 104 = 7.3. Average days = 763 / 20.9 = 36.5. Today's days = 763 / 46.9 = 16.3.

The reading follows directly. Derived far below average says the planning norm assumed a consumption rate five times what the plant achieved — a planning error, not a stock error. Today's below average says consumption has recently recovered and the position is improving. A plant looking only at the derived figure would conclude it was comfortably inside a 13.50-day target while actually holding more than twice it.

> The corrective action differs by which figure is wrong. A bad derived figure is fixed in the planning norm; a bad today's figure is fixed on the shop floor. One number cannot tell a buyer which conversation to have.

### Appendix B.2 Two clocks on four parts

Four parts, with days since last receipt and days since last issue: P1 (7, 1328), P2 (803, 805), P3 (513, 511), P4 (195, 30).

By receipt age alone, P1 falls in the 0-30 band and looks like fresh working stock; P2 and P3 look equally aged; P4 falls in the above-90 band and looks like a disposition candidate.

By both clocks the picture inverts. P1 is dead stock still being replenished — the worst case on the list, and the only one a receipt-age report would miss entirely. P2 and P3 are genuinely dormant, both clocks agreeing. P4 was bought in bulk six months ago and consumed last month: ordinary working stock with an unlucky purchase history, and flagging it for disposition would be an error.

One clock ranks these P4, P2, P3, P1 — exactly backwards at both ends. Two clocks rank them P1, P2, P3, and leave P4 alone.

### Appendix B.3 Where the provision cap bites

Four parts carry a policy provision quantity and a host-system quantity: P1 (3391, 3391), P2 (219, 0), P3 (794, 10), P4 (162, 214). Unit values are 1.26, 0.42, 0.021 and 2.09 in hundredths of a lakh respectively.

Applying Equation (provision), the non-moving quantities are 3391, 0, 10 and 162. Provision values follow: 42.77, 0, 0.21 and 3.38, for a total of 46.36.

Taking the policy quantity uncapped would have given 42.77, 0.92, 16.67 and 3.38 — a total of 63.74, or 37% too high. The overstatement comes entirely from P2 and P3, where the stock the policy would provide against is not in the building. P4 shows the other direction of the cap: the host holds 214 pieces but only 162 meet the non-moving criterion, so the remainder is untouched.

## Data provenance

Every quantitative claim in this paper resolves to one of the entries below. Figures marked as modelled estimates are not deployment measurements; their assumptions are stated.

| Metric | Reported value | Provenance |
| --- | --- | --- |
| Plant inventory target published on the dashboard, expressed in days of cover | 13.50 days | **Product specification** /products/inventory-control — Read from the Plant |

| Metric | Reported value | Provenance |
|---|---|---|
| | | Inventory screen captured on the product page; the target line is per plant and configurable. |
| Classification matrix implemented for value against demand variability | 4 value bands (A+, B, C, D) x 3 variability bands (X, Y, Z) | **Product specification** /products/inventory-control — Read from the ABC-XYZ categorisation screen captured on the product page; four value bands rather than the textbook three. |
| Ageing brackets and the clocks reported against them | 0-30, 31-60, 61-90, above 90 days, on both receipt age and consumption age | **Product specification** /products/inventory-control — Read from the Ageing Part Summary screen captured on the product page, which reports Rcpt Age Days and Cons Age Days as separate columns. |
| Source-data refresh cadence from the host ERP | daily background upload of receipts, issues and stock | **Product specification** /products/inventory-control — Stated on the architecture screen captured on the product page; a batch cadence, not a streaming feed. |
| Planning lag when ABC analysis is run monthly in a spreadsheet | 4 weeks behind every demand shift | **Product specification** /products/inventory-control |
| Hours of manual inventory data handling automated per year | 2,700 h (modelled) | **Modelled estimate** 5 people spending 3 hours per day on inventory data, across 300 working days; 60% of that effort automated; Model and defaults published in src/data/roiModels.ts |

## References


**[1]** F. W. Harris (1913). How Many Parts to Make at Once. Factory, The Magazine of Management, vol. 10, no. 2, pp. 135-136, 152; the first statement of the economic order quantity, reprinted in Operations Research 38(6), 1990, pp. 947-950. https://pubsonline.informs.org/doi/10.1287/opre.38.6.947

**[2]** H. F. Dickie (1951). ABC Inventory Analysis Shoots for Dollars, Not Pennies. Factory Management and Maintenance, vol. 109, no. 7, July 1951, pp. 92-94; the origin of ABC classification, applying Pareto and Lorenz to stock value.

**[3]** J. D. Croston (1972). Forecasting and Stock Control for Intermittent Demands. Operational Research Quarterly, vol. 23, no. 3, pp. 289-303. https://link.springer.com/article/10.1057/jors.1972.50

**[4]** A. A. Syntetos and J. E. Boylan (2005). The Accuracy of Intermittent Demand Estimates. International Journal of Forecasting, vol. 21, no. 2, pp. 303-314; evaluates moving average, single exponential smoothing and Croston's method against 3,000 real intermittent automotive series. https://doi.org/10.1016/j.ijforecast.2004.10.001

**[5]** Little, J. D. C. (1961). A proof for the queuing formula: L = lambda W. Operations Research, 9(3), 383-387.

**[6]** International Organization for Standardization (2014). ISO 22400-1: Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 1: Overview, concepts and terminology. ISO, Geneva. https://www.iso.org/standard/56847.html

**[7]** International Organization for Standardization (2014). ISO 22400-2: Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 2: Definitions and descriptions. ISO, Geneva; defines OEE as availability x effectiveness x quality ratio. https://www.iso.org/standard/54497.html

**[8]** International Electrotechnical Commission / International Society of Automation (2025). Enterprise-control system integration - Part 1: Models and terminology. IEC 62264-1; ANSI/ISA-95.00.01-2025 (IEC 62264-1 Mod).

**[9]** International Society of Automation (2013). Enterprise-control system integration - Part 3: Activity models of manufacturing operations management. ANSI/ISA-95.00.03-2013 (IEC 62264-3 Modified).

**[10]** International Electrotechnical Commission (2025). IEC 62541-1: OPC unified architecture - Part 1: Overview and concepts. IEC, Geneva. https://webstore.iec.ch/en/publication/81513

**[11]** International Organization for Standardization (2022). ISO 28000: Security and resilience - Security management systems - Requirements. ISO, Geneva; published 15 March 2022. https://www.iso.org/standard/79612.html

**[12]** GS1 (2024). GS1 General Specifications. GS1 AISBL; defines the GTIN, SSCC and GLN identification keys. https://ref.gs1.org/standards/genspecs/

**[13]** GS1 (2024). Serial Shipping Container Code (SSCC). GS1 identification key series; an 18-digit key for a logistic unit, carried under application identifier 00. https://www.gs1.org/standards/id-keys/sscc

**[14]** International Organization for Standardization / International Electrotechnical Commission (2014). ISO/IEC 15459-1: Information technology - Automatic identification and data capture techniques - Unique identification - Part 1: Individual transport units. ISO/IEC, Geneva. https://www.iso.org/standard/54779.html

**[15]** International Organization for Standardization / International Electrotechnical Commission (2024). ISO/IEC 19987: Information technology - EPC Information Services (EPCIS) Standard, version 2.0. ISO/IEC, Geneva. https://www.iso.org/standard/85557.html

**[16]** International Organization for Standardization / International Electrotechnical Commission (2024). ISO/IEC 19988:2024 - Information technology - GS1 Core Business Vocabulary (CBV). ISO/IEC, Geneva, edition 3, March 2024; publishes GS1 CBV 2.0, ratified June 2022, which fixes the vocabularies for the EPCIS business step, disposition and business transaction type fields. https://www.iso.org/standard/85558.html

**[17]** Nakajima, S. (1988). Introduction to TPM: Total Productive Maintenance. Productivity Press, Cambridge, MA.

**[18]** Kagermann, H., Wahlster, W., and Helbig, J. (2013). Securing the future of German manufacturing industry: Recommendations for implementing the strategic initiative INDUSTRIE 4.0. Final report of the Industrie 4.0 Working Group, acatech - National Academy of Science and Engineering. https://www.acatech.de/wp-content/uploads/2018/03/Final_report__Industrie_4.0_accessible.pdf

**[19]** Lee, J., Bagheri, B., and Kao, H.-A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3, 18-23.

**[20]** Shingo, S. (translated by A. P. Dillon) (1986). Zero Quality Control: Source Inspection and the Poka-Yoke System. Productivity Press, Cambridge, MA. Originally published in Japanese, 1985.