Counting Is the Easy Part: Bin Decomposition, Two-Key Putaway, and the Variance a Physical Stock-Take Actually Owes

Architecture, Computational Methods, and a Modelled Scenario for MileSoft Annual Inventory Tagging

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/annual-inventory-tagging

Abstract

The annual physical count is the most disruptive event in a warehouse's year, and most of what it produces is not shrinkage. It is counting error: bins counted twice, bins missed entirely, descriptions transcribed by hand, and a reconciliation performed two weeks later against a stock position that has since moved.

This paper presents the architecture and computational methods of MileSoft Annual Inventory Tagging. Its central argument is that a stock-take is a measurement, and a measurement without a stated error model is not evidence. The system's design decisions are read here as controls on specific error terms rather than as convenience features.

Four results carry the paper. A receipt is decomposed into uniquely tagged bins whose quantities must sum exactly to the invoice quantity, which fixes the granularity at which any later count is checkable. Putaway requires two independent scans — location and container — and stores only when both agree, which makes a misplacement a blocked action rather than an invisible one. Bin tags are locked after scanning and unscanned bins are flagged before the count closes, which converts double-counts and missed bins from procedural risks into structural impossibilities. And the reconciliation must report absolute variance, not net: a count whose positive and negative errors cancel reconciles perfectly in total value while being wrong at every location.

The work is anchored in ISO/IEC 19987 and 19988 for the event and disposition model, GS1 identification keys for the tags themselves, ISO 22400 for the KPI discipline, and ISO 5725 for the accuracy vocabulary a variance report needs. 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: Annual inventory tagging; Physical stock take; Barcode inventory count; RFID stock audit; Inventory record accuracy; Cycle counting; Bin level inventory; FIFO enforcement; Shelf life management; Variance reconciliation; Stock adjustment audit trail; Warehouse count shutdown

1 Introduction

For most warehouses and plants the annual physical count is the most disruptive event of the year. Production halts, shipping halts, temporary counters fan out across the facility with paper sheets, and the variance report that emerges runs to dozens of pages — most of it counting error rather than shrinkage.

That last clause is the whole problem. A variance report conflates two things that need opposite responses: stock that is genuinely missing, and stock that was counted wrongly. A facility that cannot separate them cannot act on either, and will book an adjustment against a discrepancy that never existed.

This paper reads the system's design as a set of controls on specific error terms. A stock-take is a measurement; a measurement without an error model is not evidence; and the design decisions that matter are the ones that remove error terms rather than the ones that speed up counting.

1.1 The three error terms

The product's own material names three failure modes that recur every year, and they map cleanly onto distinct terms in a count's error budget.

Table 1. The three dominant error terms in a manual count, and where each is controlled in Sections 3 and 4.

Error term	How it arises	Control
Double count	The same container is counted twice by different counters	One tag per bin, locked after scan (Section 4.3)
Missed count	A container behind another is never reached	Unscanned tags listed before the count closes (Section 4.3)
Transcription	A counter writes a description that does not match the record	Scan rather than transcribe; identity comes from the tag (Section 3.2)
Post-count drift	Movements between the count and the reconciliation	In-flight transactions rolled back over the count window (Section 3.3)

None of the four is addressed by counting faster. Three are addressed by what is scanned and one by when the reconciliation is drawn — which is why speed is a consequence of this design rather than its objective.

1.2 Contributions

1. Bin decomposition stated as a conservation constraint, fixing the granularity at which a count is checkable (Section 4.1).
2. Two-key putaway as a control function in Shingo's sense, with its error model (Section 4.2).
3. Tag locking and missed-bin detection as structural eliminations of the two dominant count errors (Section 4.3).
4. Absolute against net variance, and why a reconciled total is not evidence of an accurate record (Section 4.4).
5. An eight-dimension capability reference framework for stock-take systems (Section 6.3).

2 Background and Related Work

Physical inventory verification is one of the oldest obligations in commerce and one of the least examined. Its literature sits in three places that rarely cite each other: measurement science, error proofing, and the identification standards that make an object countable at all.

2.1 Trueness and precision are different failures

ISO 5725 separates trueness — the closeness of a mean to a reference value — from precision, the scatter of repeated measurements. The distinction is routinely collapsed in stock-take reporting, and collapsing it is the source of the most common misreading of a variance report.

A count with poor trueness is systematically wrong in one direction: a scanner configured for the wrong pack size, a bin standard that counts crates as pieces. A count with poor precision is scattered: some locations high, some low, no pattern. The first is a configuration defect and is cheap to fix; the second is a process defect and is not.

Section 4.4 shows that net variance measures only the first and absolute variance measures both, which is why reporting one number lets a poor count pass.

2.2 Control against warning at the point of error

Shingo distinguishes the control function of an error-proofing device — the next action becomes impossible — from the warning function, where the operator is told and may proceed. The distinction is decisive in a stock-take because the environment is the worst case for warnings: temporary staff, unfamiliar layout, time pressure, and a lamp nobody has been trained to fear.

The putaway design of Section 4.2 is a control function. It does not tell the operator that the container and location disagree; it refuses to store the record until they agree. The distinction is the difference between an error rate that falls and one that is documented.

2.3 What makes an object countable

A count presupposes a set of distinguishable objects. GS1 provides the Global Trade Item Number for a trade item class and the Serial Shipping Container Code for an individual logistic unit; ISO/IEC 15459-1 covers transport units outside a GS1 scheme; the General Specifications govern symbology selection and print quality, which decide whether a tag scans at all on the first attempt.

The class-against-instance distinction is what makes bin-level counting possible. A tag identifying only the part number tells a counter what is in front of them; a tag identifying this container tells the system which of forty identical containers has been counted, which is the fact tag locking depends on.

The event model is EPCIS, standardised as ISO/IEC 19987, with its vocabulary in ISO/IEC 19988. A count observation is an event with a business step and a disposition like any other, which means a stock-take integrates into a traceability record rather than sitting beside it as a spreadsheet.

3 System Overview

The system is a warehouse spine with a stock-take mode. Everyday inward, storage and outward flows create and maintain the tagged containers; the annual count is a verification pass over the structure those flows built. This ordering matters: a count is only as good as the granularity established at receipt, months earlier.

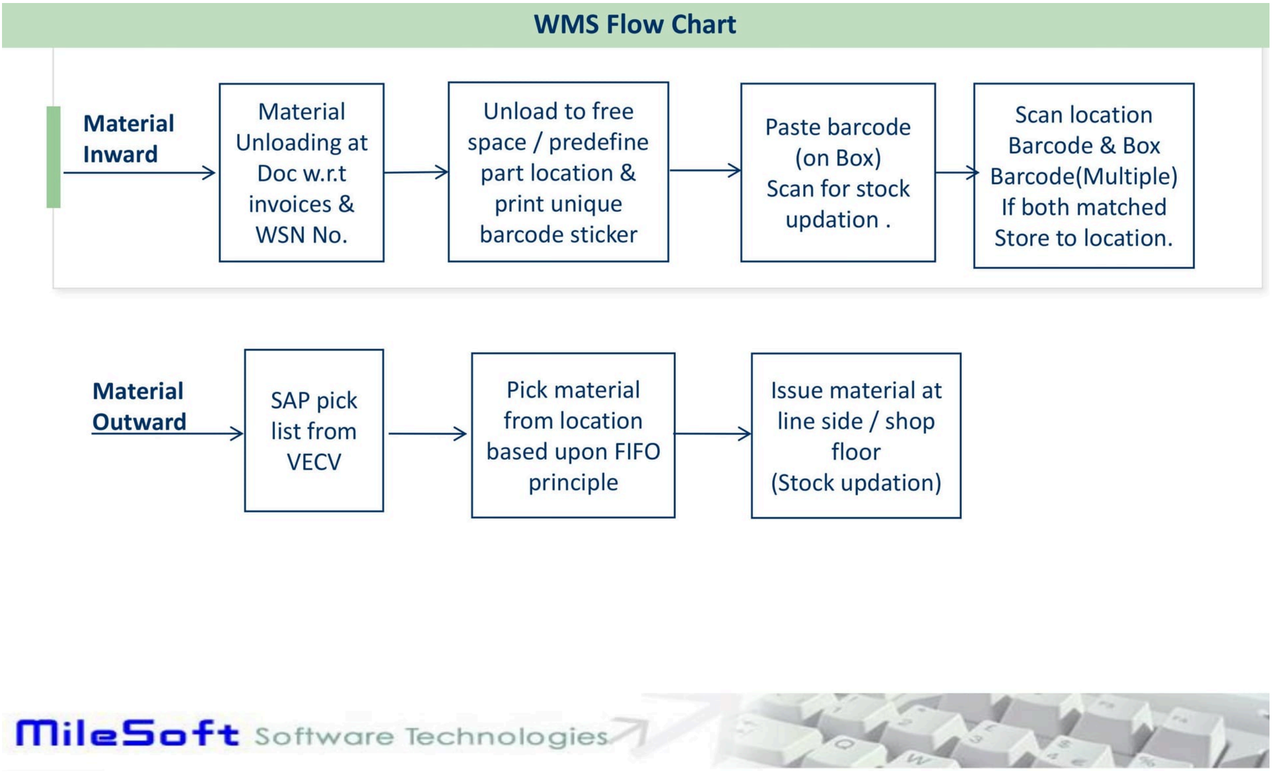

Figure 1. Product capture: the material flow. Inward runs from unloading against the document and invoice, through allocation to a free or predefined location with a unique barcode printed, to barcode application and stock update, and finally a location-and-container scan that stores only when both match. Outward runs from the host system's pick list, through location picking under first-in-first-out, to issue at the line with stock update.

3.1 Receipt, decomposition and tagging

A receipt is captured against its document and invoice, with the part, its line-side location, its standard bin quantity and its shelf life in days. Two quantities are recorded: the quantity the invoice claims and the quantity actually received.

The two-quantity pattern recurs throughout this library and does the same work here as elsewhere. A single received quantity records what was accepted and destroys the evidence of a supplier discrepancy; two record both, and the difference is a supplier performance fact rather than a rounding error absorbed at the dock.

The received quantity is then decomposed into containers at the standard bin quantity, with a residual container for the remainder, and each container is given a unique barcode. Section 4.1 states the conservation constraint this decomposition must satisfy.

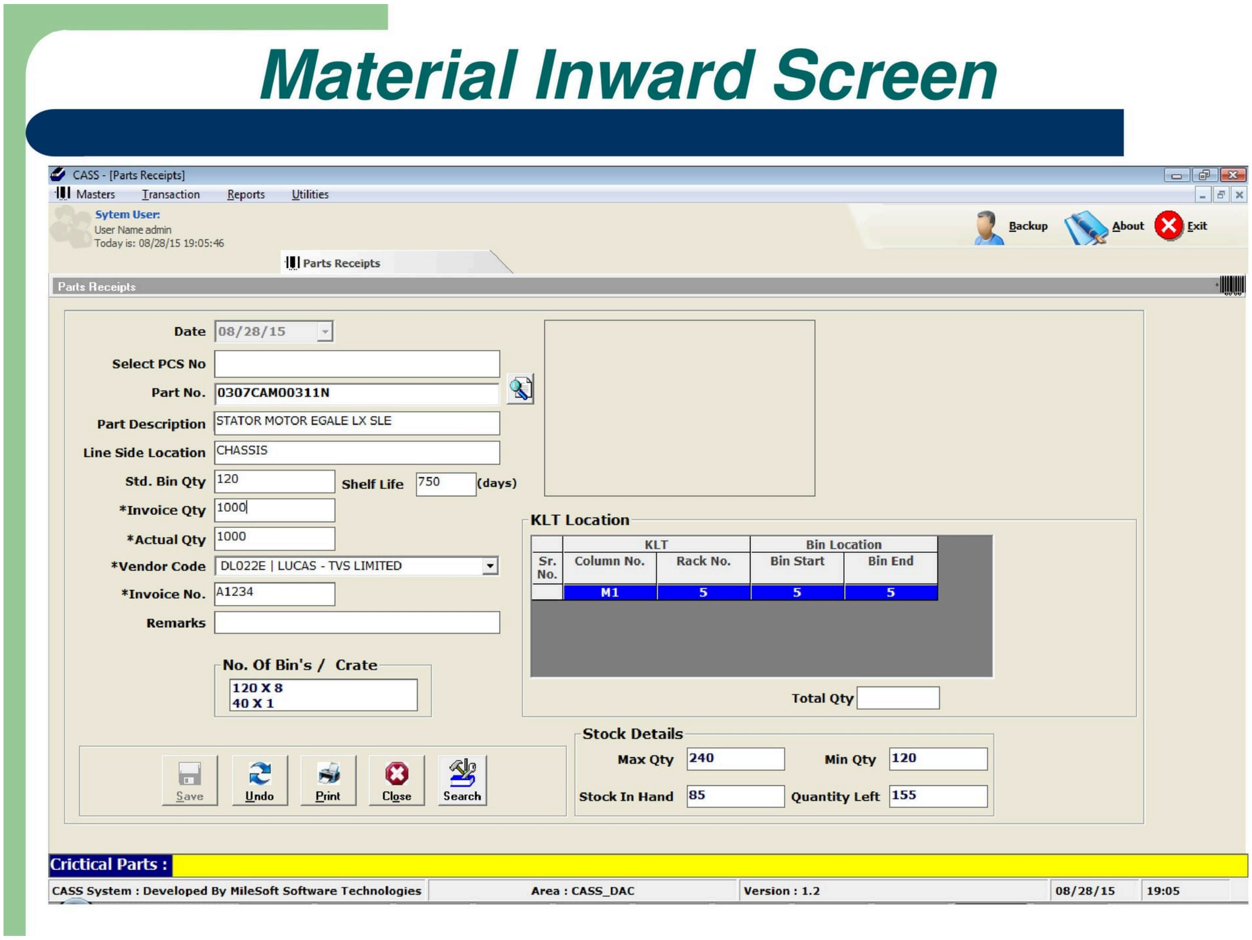

Figure 2. Product capture: the material inward screen. The part carries a line-side location, a standard bin quantity and a shelf life in days; invoice quantity and actual quantity are separate fields; the container decomposition is shown as counts by quantity and must reconcile to the received total; and the allocated location is expressed as a column, rack and bin range with the part's minimum, maximum, on-hand and remaining capacity beside it.

3.2 Putaway, and the second key

The tagged container is carried to its allocated location. The operator scans the location barcode and the container barcode, and the system stores the container only when the pair agrees with the allocation.

The design decision here is that both keys are read from the physical world. A workflow that pre-selects the location on screen and asks the operator to confirm by scanning the container is checking one fact against a screen, not two facts against each other, and it fails silently whenever the screen and the aisle disagree.

> This is the moment a warehouse error becomes invisible. A container in the wrong bin is not lost — the stock exists and the system says it exists — so nothing surfaces until a picker cannot find it or a count reports two compensating variances.

3.3 Issue under first-in-first-out, with shelf life alongside

Outward begins with a pick list from the host system. Material is picked from location under a first-in-first-out rule and issued to the line, with stock updated on issue.

Shelf life is carried per part — 750 days on the captured part record — and the interaction with first-in-first-out is worth being explicit about, because the two rules are not the same rule. First-in-first-out selects by receipt date; expiry-based selection selects by remaining life. They coincide only where every receipt of a part carries the same shelf life, which is the common case and not the general one.

The system reports first-in-first-out violations as a standing report rather than assuming its own enforcement worked. That is the correct posture for any enforcement mechanism operating through human hands: enforce, then verify that the enforcement held.

3.4 The count window and in-flight movement

The count works best with a brief operational pause — four to eight hours for the initial physical pass, on the product's published figure. In-flight transactions across the count window are rolled back so that goods received or shipped during the count are excluded from the variance calculation.

This addresses the fourth error term of Section 1.1, and it is the one most often mishandled. A variance drawn against a book position that has moved since the count attributes post-count drift to shrinkage, and the attribution is invisible because both look like a quantity difference.

4 Computational Methods

Four computations carry the paper. The first three are controls applied while stock moves; the fourth is what the count owes its reader once it closes.

4.1 Bin decomposition and its conservation constraint

A receipt of quantity Q is split into containers at the part's standard bin quantity, with one residual container holding the remainder.

$$n = \mathit{floor}(\text{received quantity} / \text{standard bin quantity}),\ r = \text{the remainder, and } n \text{ full bins plus the residual must equal the received quantity exactly} \qquad \text{(decomp)}$$

q(std) is the standard bin quantity carried on the part record. The identity is a constraint the system must enforce at receipt, not an arithmetic observation — a decomposition that does not close indicates a mis-keyed quantity before any stock is put away.

For the captured receipt, 1,000 pieces at a standard bin quantity of 120 give n = 8 and r = 40: eight containers of 120 plus one of 40, which is exactly what the screen records.

The decomposition's real function is granularity. A count can verify only what is independently identified, so a receipt held as one undifferentiated quantity can be verified only in total, while nine tagged containers can be verified nine times. This is the same argument as the traceability tier structure elsewhere in this library: resolution is decided at capture and cannot be recovered afterwards.

> The unit of a stock-take is not the part. It is the smallest independently tagged container, and every accuracy figure a count produces is an accuracy figure about those.

4.2 Two-key putaway and its error model

A putaway is accepted only when both scanned keys agree with the allocation. Stating this as a condition makes the residual error visible.

store only if the scanned location equals the allocated location and the scanned container equals the allocated container; the undetected-error probability is the product of the two individual error rates and their correlation factor (putaway)

p(l) and p(c) are the probabilities of a wrong location and wrong container scan respectively, and rho accounts for the correlation between them. Independent keys give rho near 1; keys that can both be satisfied by one wrong action give rho far larger.

The correlation term is where designs of this class succeed or fail. Two keys read from two physical labels are close to independent, so a putaway error requires two simultaneous mistakes. A workflow where the screen pre-selects the location and only the container is scanned collapses rho, because a single wrong screen selection satisfies both conditions at once.

The control-against-warning distinction of Section 2.2 then determines what happens on a mismatch. Refusing to store means the container cannot enter the record in the wrong place; warning and permitting means it can, with a note nobody reads.

4.3 Coverage: making double and missed counts impossible

The two dominant errors in a manual count are both coverage errors: a container counted more than once, and a container never reached. Both are eliminated by treating the count as a set operation over tags rather than a walk over aisles.

S is the set of tags scanned, M the missed set - every tag not in S - and tag locking guarantees the number of scans equals the number of distinct tags scanned (coverage)

T is the set of tags expected in the count scope. Locking makes a second scan of the same tag a no-op rather than a second count, so the scan count cannot exceed the population; the missed set is computed by difference before the count is allowed to close.

Two consequences follow that procedure alone cannot deliver. Double-counting is impossible rather than discouraged, because the second scan changes nothing. And the missed set is known while the counters are still on the floor, so flagged bins are re-counted before the team leaves rather than the following week from memory.

This is the mechanism behind the published reduction in count duration. The saving is not from scanning being faster than writing; it is from never running the count twice.

4.4 The variance the reconciliation owes

Once the count closes, the system compares counted quantity to book quantity per container and reports the difference. How that difference is aggregated decides whether the report is evidence or reassurance.

$$\textit{per-container variance v(i); net value variance sums the signed differences; absolute value variance sums their magnitudes; record accuracy alpha is the share of containers with zero variance} \qquad \text{(variance)}$$

u(i) is the unit value of the part in container i. Net variance is what a financial reconciliation needs; absolute variance and record accuracy are what an operational judgement needs, and the three answer different questions.

The failure this prevents is specific and common. A facility with one location over by 400 pieces and another under by 400 has a net variance of zero and an absolute variance of 800. Reported by net alone, the count reconciles perfectly and the facility concludes its records are sound; in fact it has two errors, both of which will cause a pick failure, and neither of which it now knows about.

Record accuracy — the fraction of containers whose count matched exactly — is the figure a warehouse should manage against, because it is the probability that any given pick will find what the system promised. It is also the only one of the three that cannot be improved by compensating errors.

> Ask a stock-take report for three numbers: net variance, absolute variance, and the proportion of locations that matched exactly. A system that offers only the first is reporting to the ledger and not to the warehouse.

The variance report carries counter name, location, timestamp and before-and-after quantities per part, with discrepancies stratified by value impact — which is what makes an adjustment attributable rather than anonymous, and what a statutory audit asks for.

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 either an observed property of the software or a modelled scenario with its assumptions printed.

5.1 Observed properties of the software

Table 2. Properties read directly from the delivered software, with where each was observed.

Property	Value	Where observed
Two receipt quantities	invoice quantity and actual quantity as separate fields	Material inward screen
Container decomposition	recorded as counts by quantity, reconciling to the receipt	Material inward screen
Shelf life on the part	carried in days at receipt	Material inward screen
Location granularity	column, rack and bin range	Material inward screen
Min / max / on-hand / remaining	reported together on the receipt screen	Material inward screen
Two-key putaway	location and container scanned; stored only if both match	Material flow chart
First-in-first-out verification	a standing report of violations, distinct from enforcement	Reports menu
Reports available	19 in the reports menu	Reports menu

The last row is worth a note on honesty in both directions. The product page claims over fifteen reports; the captured menu holds nineteen. Where a count can be made it has been made here rather than repeating the marketing range.

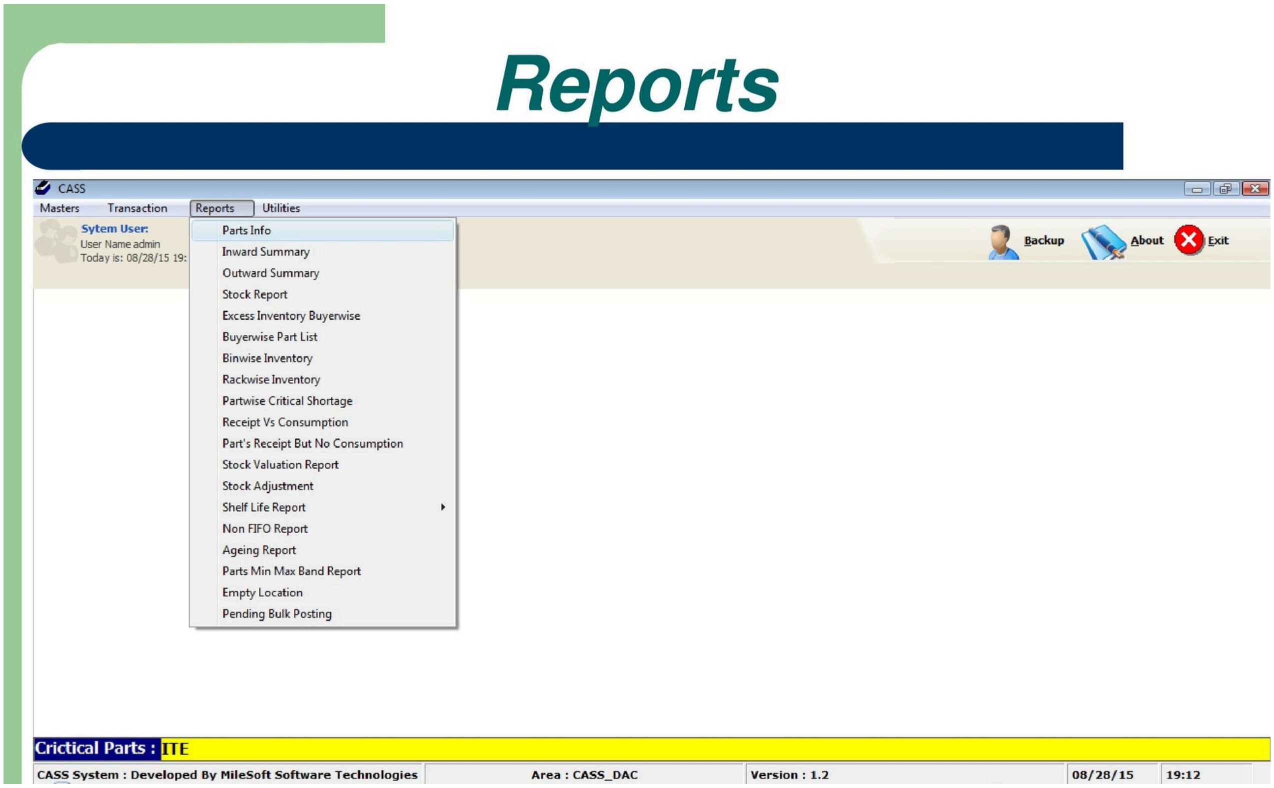

Figure 3. Product capture: the reports menu. Nineteen reports covering stock, ageing, valuation, bin and rack level positions, empty locations, minimum-maximum bands, shelf life, stock adjustment, receipt against consumption, parts received but not consumed, and first-in-first-out violations.

5.2 Figures published for this module

Table 3. Published figures for this module, with their provenance.

Figure	Value	Provenance
Physical count phase	8-16 hours, against 48-72 on paper	Product FAQs — vendor-reported, no sample or facility size
Operational pause required	4-8 hours for the initial pass	Product FAQs — vendor-reported
Reports for auditors	15 or more	Product page — a floor, and the captured menu shows 19

The count-duration figure is the one a buyer will weigh most heavily and the one carrying least evidence. It is a range with no facility size, no container population and no definition of what the phase includes. Section 4.3 gives the mechanism that would produce such a reduction — never running the count twice — which makes the claim structurally plausible without making it measured.

A four-to-one reduction in count duration is a large claim. Ask a vendor making it for the container population counted, the number of counters, and whether the figure includes the re-count of flagged bins.

5.3 Modelled effort avoided

The published return model for this product is stated as stock-take and reconciliation time saved. Its assumptions are printed here so a reader can substitute their own.

Table 4. Modelled scenario — not a deployment result. Assumptions are the published defaults of the return-on-investment model for this product.

Assumption	Value
People engaged in the audit	8
Audit hours per person per day	4
Audit window	60 days
Baseline effort	1,920 hours per audit year
Share of that effort automated	70%
Modelled hours avoided	1,344 per audit year

Two observations. The model's window is sixty days, not the count weekend — so what it prices is the reconciliation tail rather than the count itself, which is consistent with the product's own framing that the expensive part is the fortnight of spreadsheet archaeology afterwards. And it prices only labour: the shutdown hours of production and shipping are excluded, and in most facilities those dominate.

A buyer computing this properly should add the shutdown. If a facility loses forty-eight hours of production to a paper count and eight to a scanned one, the recovered output is likely to exceed the entire labour saving the published model reports.

6 Discussion

6.1 The annual count is a symptom

A facility that must halt for two days each year to discover what it holds is telling you that its records were not trusted for the preceding three hundred and sixty-three. The count is a consequence of that distrust, not a cause of accuracy.

This reframes what a tagging system is for. Its most valuable output is not the annual variance report but the record accuracy figure of Section 4.4, because that figure is what determines whether counts can eventually become sampled rather than exhaustive. A facility at 99.5% record accuracy can verify a stratified sample and defend it; one at 92% cannot, and will be counting everything for as long as that holds.

The controls of Sections 4.1 to 4.3 are what move that figure, and all three operate during the year rather than during the count. A system evaluated only on how fast it counts is being evaluated on the least important thing it does.

> The goal is not a faster annual count. It is a facility whose records are good enough that the annual count stops being the only thing anyone believes.

6.2 Enforcement and verification are different jobs

The system enforces first-in-first-out at picking and separately reports violations of it. That is worth drawing out, because carrying both is unusual and correct.

Any enforcement that passes through human hands has a bypass. A picker facing an empty designated location and a full one behind it will take the accessible stock, and a well-designed system lets them — blocking the pick would stop the line. The violation report is what converts that necessary flexibility from an untracked leak into a measured exception with a location attached.

The same logic recommends caution with the shelf-life interaction of Section 3.3. Where receipts of a part carry different remaining lives, first-in-first-out and expiry-first selection diverge, and a system enforcing the first while an operator believes it enforces the second will issue expiring stock in strict receipt order.

6.3 A capability reference framework for stock-take systems

Table 5. Capability reference framework. Each dimension separates a stock-take system from a digital counting sheet, and each is answerable by demonstration during a live count rather than by reading a specification.

Dimension	Question the system must answer by demonstration
D1 Container identity	Is the tag on the box unique to that box, or does it identify only the part?
D2 Decomposition closes	Key a receipt whose container quantities do not sum to the total. Is it accepted?
D3 Two independent keys	At putaway, are both keys read from physical labels, or is one pre-selected on screen?
D4 Control, not warning	Scan a mismatched location and container. Is the putaway refused or merely noted?
D5 Tag locking	Scan the same bin twice. Does the count change?
D6 Missed set before close	Can the system list unscanned tags while counters are still on the floor?
D7 Three variance figures	Does the report give net, absolute and the proportion of locations matching exactly?
D8 Window integrity	Receive goods during the count. Do they appear in the variance?

> D5 and D7 take a minute each and settle most of the question. A system where a second scan increases the count is a counting sheet; one that reports only net variance is reporting to the ledger rather than to the warehouse.

6.4 Generalisability

The coverage argument of Section 4.3 generalises to any exhaustive verification over an identified population: asset registers, tool cribs, library stock, calibration inventories. The absolute-against-net variance argument generalises further still, to every reconciliation where compensating errors are possible — which is most of them.

What does not generalise is the count-duration figure. It depends on facility size, container population, counter numbers and layout, none of which the published range states, and a facility adopting it as a plan would be planning against another site's geometry.

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 absence is itself informative.
2. The count-duration figures are vendor-reported. Eight to sixteen hours against forty-eight to seventy-two carries no facility size, container population, counter count or definition of what the phase includes.
3. The modelled 70% is an assumption. It prices the reconciliation tail over a sixty-day window and excludes the production and shipping shutdown, which in most facilities is the larger cost.
4. Tag quality is an unmeasured dependency. Every control in Section 4 assumes the tag scans first time; a label degraded by oil, abrasion or cold storage silently converts a scan workflow back into a transcription workflow.
5. Two-key putaway does not prevent all misplacement. It prevents storing a container against a location that disagrees with the allocation; it does not prevent a container being physically left somewhere other than where it was scanned.
6. Record accuracy is a property of tagged containers only. Bulk stock, open stores and line-side material that never entered the container structure are outside every figure in Section 4.4.
7. The roll-back of in-flight transactions is described but not observed. Its correctness under partial receipts or split shipments during the count window is not something this paper can attest to.
8. The screens are one deployment at one date. Standard bin quantities, shelf lives and location schemes are configuration, and no value read from a capture is offered here as a benchmark.

> The fourth limitation is the one facilities underestimate. Print quality and label survivability are treated as a procurement detail and are in fact the foundation every control in this paper rests on.

8 Future Work

- A measured deployment. A published count with facility size, container population, counter numbers and the three variance figures would replace the strongest claim in the product's material with evidence.

- Record accuracy as a continuous metric. Publishing the proportion of containers matching exactly, tracked between counts rather than only at them, is the figure that decides whether exhaustive counting can eventually become sampled.
- Sampling plans with stated confidence. Once record accuracy is high and measured, a stratified sample defensible to an auditor is the natural next step, and it is the only route to ending the annual shutdown rather than shortening it.
- Label survivability testing per environment. Publishing scan-first-time rates by storage condition would put a number on the dependency Section 7 identifies as the least examined.
- Expiry-first selection alongside first-in-first-out, so that parts whose receipts carry different remaining lives are issued by the rule the operator believes is in force.

9 Conclusion

A physical count is a measurement, and most variance reports are not evidence because they carry no error model. This paper has read one system's design as a set of controls on the specific terms that make up that error.

Four results are worth carrying away. Resolution is set at receipt: a count can verify only what was independently tagged, so the decomposition of a receipt into containers — and the constraint that those containers sum exactly to what arrived — fixes what any later count is capable of proving. Putaway with two independently read keys makes a misplacement a blocked action rather than an invisible one, and the correlation between those keys is what a buyer should examine. Tag locking and a missed set computed before the count closes eliminate the two dominant coverage errors structurally, which is also the mechanism behind the count-duration claim. And a reconciliation owes three numbers, not one: a net variance of zero is compatible with every location being wrong.

The framework of Section 6.3 is offered as the durable contribution. Its fifth dimension takes a minute: scan the same bin twice, and watch whether the count moves.

Appendix A Nomenclature

Table 6. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
Q	Quantity actually received against a document
q(std)	Standard bin quantity carried on the part record
n, r	Number of full containers and the residual quantity in the last one
l(scan), l(alloc)	Location scanned at putaway and location allocated by the system
c(scan), c(alloc)	Container scanned at putaway and container expected there
p(l), p(c)	Probability of a wrong location scan and of a wrong container scan
rho	Correlation factor between the two scan errors; near 1 when the keys are independent
T, S, M	Tags expected in scope, tags scanned, and the missed set T minus S
v(i)	Per-container variance: counted quantity less book quantity
V(net)	Net value variance — the signed sum, what a ledger reconciliation needs
V(abs)	Absolute value variance — the sum of magnitudes, what a warehouse needs
alpha	Record accuracy: the share of containers whose count matched exactly
u(i)	Unit value of the part in container i
FIFO / FEFO	First-in-first-out by receipt date; first-expired-first-out by remaining life
SSCC / GTIN	Serial Shipping Container Code and Global Trade Item Number

Appendix B Worked Numerical Examples

Appendix B.1 Decomposition, and what it buys

A receipt of 1,000 pieces arrives for a part whose standard bin quantity is 120.

Applying Equation (decomp): n = floor(1000 / 120) = 8, r = 1000 — 960 = 40. Nine tagged containers — eight of 120 and one of 40 — and the identity closes: 8 x 120 + 40 = 1,000.

Now consider what each arrangement permits at count time. Held as one quantity, the receipt admits exactly one verification, which passes if the total is right however the stock is distributed. Held as nine containers, it admits nine, and a container in the wrong aisle produces a variance in two places rather than none.

> Nine tags cost nine labels and roughly a minute. They are the difference between a count that can locate an error and one that can only detect its absence.

Appendix B.2 What the second key is worth

Suppose an operator scans the wrong location with probability 0.004 and picks up the wrong container with probability 0.006, over 20,000 putaways in a year.

With one key — container only, location pre-selected on screen — the undetected misplacement rate is 0.004, giving roughly 80 misplaced containers a year, each invisible until a pick fails.

With two independently read keys and rho near 1, Equation (putaway) gives 0.004 x 0.006 = 0.000024, or about one misplacement every two years at the same volume.

Now collapse the independence. If the screen pre-selects the location and the operator confirms by scanning whatever container is in hand, a single wrong screen selection satisfies both conditions, rho rises toward 1/p(c), and the undetected rate returns to roughly the single-key figure. The second scan happened; the second key did not.

Appendix B.3 Three numbers from one count

A count covers 2,400 containers. 2,352 match exactly. Of the remaining 48: 22 are over by an average of 31 pieces at 18 currency units each, and 26 are under by an average of 27 pieces at 19 units each.

Applying Equation (variance): the positive value variance is 22 x 31 x 18 = 12,276 and the negative is 26 x 27 x 19 = 13,338. Net variance is -1,062. Absolute variance is 12,276 + 13,338 = 25,614. Record accuracy is 2,352 / 2,400 = 98.0%.

Reported by net alone the count looks excellent: a shortfall of 1,062 on a population of 2,400 containers is well inside any materiality threshold, and the ledger reconciles. Reported honestly, the facility has 48 wrong locations, an absolute error twenty-four times the net, and a two-percent probability that any given pick will not find what the system promised.

> The two reports describe the same count. One of them tells a warehouse manager that forty-eight picks are going to fail before the next count.

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
Physical count phase duration with barcode-driven capture	8-16 h, against 48-72 h on paper	**Product specification** /products/annual-inventory-tagging — Vendor-reported in the product FAQs; no sample size, facility size or count population is stated.
Operational pause the count works best with	4-8 h for the initial physical pass	**Product specification** /products/annual-inventory-tagging
Double-count and missed-bin control	one tag per bin, locked from re-scan, unscanned bins flagged before close	**Product specification** /products/annual-inventory-tagging
Operational shutdown imposed by a paper-based annual count	48-72 hours	**Product specification** /products/annual-inventory-tagging
Stock, ageing and valuation reports produced for auditors	15+	**Product specification** /products/annual-inventory-tagging
Shelf life carried on a part at receipt	750 days on the captured part record	**Product specification** /products/annual-inventory-tagging — Read from the material inward screen captured on the product page; one part's value, not a product default.
Stock-take and reconciliation hours automated per audit year	1,344 h (modelled)	**Modelled estimate** 8 people spending 4 audit hours per day across a 60-day audit window; 70% of stock-take and reconciliation effort automated; Model and defaults published in src/data/roiModels.ts

References

[1] 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

[2] 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

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

[4] 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

[5] GS1 (2023). GS1 Logistic Label Guideline. GS1 AISBL; the SSCC is the only mandatory field on a GS1 logistic label. https://www.gs1.org/standards/gs1-logistic-label-guideline/1-3

[6] 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

[7] International Organization for Standardization (2023). ISO 5725-1: Accuracy (trueness and precision) of measurement methods and results - Part 1: General principles and definitions. ISO, Geneva (superseding ISO 5725-1:1994).

[8] International Organization for Standardization (2019). ISO 5725-2: Accuracy (trueness and precision) of measurement methods and results - Part 2: Basic method for the determination of repeatability and reproducibility. ISO, Geneva.

[9] 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

[10] 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

[11] 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).

[12] 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).

[13] International Organization for Standardization (2007). ISO 22005:2007 - Traceability in the feed and food chain - General principles and basic requirements for system design and implementation. ISO, Geneva; defines traceability as the ability to follow the movement of a feed or food through specified stages of production, processing and distribution, and states the design requirements a traceability system must meet. https://www.iso.org/standard/36297.html

[14] 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

[15] 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.

[16] 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

[17] 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.

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

[19] 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

[20] 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.