



One State Store from Gate to Shelf: Unifying Yard, Warehouse, and Fleet Operations

Architecture, Computational Methods, and Field Evidence from the MileSoft Warehouse Suite

MileSoft Engineering Research Group

MileSoft Software Technologies

<https://milessoft.net/>

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

Working paper — Version 1.0 — 22 August 2026

Permanent link: <https://milessoft.net/research/warehouse>

Abstract

A distribution centre is a queueing system whose servers are dock doors and whose customers are vehicles, wrapped around an inventory system whose accuracy determines whether the queue can be served at all. In most operations these two systems are managed by different tools, or by no tool: yard scheduling lives in messaging apps and paper trip sheets, while the warehouse management system holds a book position that diverges from the shelf. The consequence is not merely inefficiency. It is that neither system can answer the operational question that matters — where is this object right now, and what is it waiting for.

This paper presents the architecture, computational methods, and field evidence of the MileSoft warehouse and logistics suite: six interoperating modules (WMS, Truck Turnaround, Loading Dock Management, Inventory Tracking, Order Fulfillment, and Fleet Optimization) built on a single governed state store in which every physical object — vehicle, dock door, parking slot, unit load, and stock location — holds exactly one current state, and every state transition writes exactly one event.

We describe an identity layer built on GS1 keys and ISO/IEC 15459 transport-unit identifiers; a yard state machine whose system-generated timestamps make dwell time measurable rather than self-reported; slotting and allocation over a reservation model that eliminates the phantom-pick class of error; wave and route construction drawn from the order-picking and vehicle-routing literature; and continuous exception-triggered cycle counting in place of scheduled freeze counts. Performance measures follow the ISO 22400-2 definitions, and the event vocabulary follows the ISO/IEC 19987 (EPCIS) shape.

A production deployment at a third-party logistics operator in Madhya Pradesh, India, reports a 38% reduction in dock idle time, average truck dwell falling from 3 h 10 min to 1 h 20 min, inventory accuracy rising from 91% to 99.4%, and new-client onboarding compressing from six weeks to eight days. We close by proposing a capability reference framework for evaluating warehouse and yard platforms.

Keywords: Warehouse management system; Yard management; Truck turnaround time; Dock scheduling; Order picking; Vehicle routing; Inventory record accuracy; Cycle counting; GS1 SSCC; EPCIS; ISO 22400; Third-party logistics; Cross-docking; Little's Law

1 Introduction

A distribution centre is two coupled systems wearing one roof. The first is a queueing system: vehicles arrive, wait, are served at a dock door, and depart. The second is an inventory system: goods are received, put away, held, picked, and shipped. The coupling is tight — a dock cannot be served if the stock it needs is not where the system says it is, and stock cannot be replenished if the inbound vehicle is still in the gate queue.

Incumbent practice manages the two separately, and often manages the first not at all. Appointments are agreed over messaging apps; gate arrivals are written on a pad; dock assignment is a supervisor's judgement call. Meanwhile the warehouse management system maintains a book position that is reconciled against physical reality once a year, during a shutdown, by counting everything at once.

This paper argues that both problems have the same root and the same fix. The root is that no single system holds the current state of each physical object. The fix is a governed state store in which every vehicle, door, slot, unit load, and stock location has exactly one current state, and every transition writes exactly one event.

1.1 Why dwell time is usually unmeasurable

Operators routinely quote a turnaround time they cannot substantiate. The number is derived from a driver-signed sheet, or from the difference between two entries typed at shift end. Both are reconstructions, and both are systematically biased: the times that get recorded accurately are the ones nobody was too busy to record, which are the good ones.

Little's Law gives the relationship that makes this matter. For a system in steady state the average number of vehicles in the yard equals the arrival rate multiplied by the average time each spends there. If dwell is measured with a bias, yard occupancy predictions inherit that bias, and capacity decisions — how many doors, how large a parking area — are made on a number that was never observed.

The remedy is not a better form. It is to make the timestamp a by-product of a state transition the system performs itself, so that recording is not a task anyone can be too busy to do.

1.2 Why inventory accuracy decays

Inventory accuracy is not lost in a single event; it erodes through unbooked scrap, mis-keyed receipts, put-away to an unrecorded location, and picks taken from the nearest rather than the nominated face. Each is individually small. Their accumulation is what turns a 99% position into a 91% one over a year.

The annual freeze count is a poor instrument against this. It detects the accumulated error long after the process that caused it has been forgotten, it halts operations, and it produces a write-off rather than a diagnosis. Continuous counting triggered by exceptions detects the same error while its cause is still identifiable.

1.3 Contributions

1. A single-state-store architecture spanning yard and warehouse, in which identity, state, and event history are properties of one shared layer rather than of six applications.
2. Module-level computational methods stated against published standards and the operations-research literature: identity resolution, yard state transitions and dwell decomposition, reservation-based allocation, wave and route construction, and exception-triggered cycle counting.
3. Field evidence from a multi-client third-party logistics deployment, with every reported figure traced to its source.
4. An eight-dimension capability reference framework for evaluating warehouse and yard platforms by demonstration.

2 Background and Related Work

Four bodies of work inform the design: identification standards, the warehouse operations literature, vehicle routing, and standardised performance measurement.

2.1 Identification standards

The GS1 system supplies the identification keys that let independent organisations refer to the same physical object without prior agreement. The Serial Shipping Container Code is an eighteen-digit key identifying a logistic unit — any combination of goods assembled into a case, a pallet, or a truck load that must be managed as a unit through the supply chain. It is carried in a GS1-128 barcode, or an RFID tag, under application identifier 00, and it is the only mandatory field on a GS1 logistic label. The Global Trade Item Number identifies what is inside; the Global Location Number identifies where.

ISO/IEC 15459-1 provides the complementary rule set for a unique identifier on an individual transport unit, independent of which agency issued it. Taken together these define the boundary condition for the architecture: identity must be resolvable from the label alone, without reference to a shipment document that may not have arrived.

ISO/IEC 19987 standardises EPC Information Services, the event vocabulary for visibility data — what happened, when, where, and why, with the business step and disposition. Its contribution here is structural rather than protocol-level: it establishes that a movement record is a typed event about an identified object, which is exactly the shape the state store writes.

2.2 Warehouse operations

Order picking dominates warehouse operating cost, and the design space — layout, storage assignment, zoning, batching, and routing — has been surveyed comprehensively by de Koster, Le-Duc and Roodbergen, and the broader operational problem by Gu, Goetschalckx and McGinnis. Two findings from that literature shape the implementation described here.

First, routing and batching interact: the best picker route depends on which orders were batched together, so optimising them independently leaves value on the table. Second, the marginal return from sophisticated routing heuristics collapses when storage assignment is poor, because travel is then dominated by distance between unrelated items rather than by route construction. The practical consequence is that slotting quality is prerequisite to routing quality, which is why the allocation model is described before the wave engine.

2.3 Vehicle routing

Outbound distribution is an instance of the vehicle routing problem, introduced by Dantzig and Ramser as the truck dispatching problem, and made tractable at operational scale by the savings heuristic of Clarke and Wright. The suite uses a savings-based construction with capacity and time-window feasibility checks, followed by local improvement, rather than an exact method.

This is a deliberate choice rather than a limitation. Route plans in a live distribution centre are invalidated by events — a late inbound, a cancelled order, a vehicle breakdown — on a timescale shorter than an exact solver's runtime. A fast constructive heuristic that can be re-run on every material change produces better realised routes than an optimal plan computed against stale data.

2.4 Standardised performance measurement

ISO 22400-2 specifies key performance indicators for operations management with their formulae, constituent elements, and time behaviour. Although written for manufacturing, its treatment of utilisation and availability transfers directly to dock doors, which are servers with planned and unplanned unavailability exactly as a machine is.

The relevant discipline it imposes is that the time state must be captured, not inferred. A door that is unavailable because it is under maintenance is not idle, and counting it as idle understates utilisation and overstates the case for building more doors.

2.5 Incumbent tooling and its failure modes

- Split state. The yard knows where the truck is; the WMS knows where the stock is; neither knows whether the truck can be served. Every dock assignment is therefore a human cross-check.
- Reconstructed timestamps. Arrival and departure times entered by people after the fact are biased toward the uneventful cases, making measured dwell systematically optimistic.

- Book-only inventory. A position reconciled annually cannot support exception detection, so a discrepancy surfaces when a picker reaches an empty bin.
- Per-client integration. Where a WMS exposes no API, every new client is a bespoke data pipeline, and onboarding cost scales with client count rather than amortising across it.

Each failure follows from state living in more than one place. Consolidating state is what removes them; the modules above it are then thin.

3 System Overview and Architecture




Figure 1. Schematic. Platform architecture. One governed state store holds the current state of every physical object and the ledger of every transition. Six modules read and write through it; none holds private state. The physical estate below writes exactly one event per transition, and the integration layer above exchanges plans and confirmations with carriers, shippers, and client systems.

3.1 Design goals

- G1 — One state per object. Every vehicle, door, slot, unit load, and stock location has exactly one current state, in one store. No module keeps a private copy.
- G2 — System-generated transitions. Timestamps are produced by the system at the moment of transition, never typed in afterwards.
- G3 — Identity from the label. A unit load is resolvable from its SSCC alone, without a matching shipment document.

- G4 — Reservation before release. Stock is reserved at allocation, so no two tasks can be released against the same physical unit.
- G5 — Standards-defined measures. Utilisation and availability follow ISO 22400-2 definitions computed from captured time states.
- G6 — Integration as configuration. Client onboarding is API mapping, not a bespoke pipeline.

3.2 The governed state store

The store has seven components. The identity layer resolves GS1 keys and ISO/IEC 15459 transport-unit identifiers from scans and RFID reads (G3). The yard state machine owns vehicle state and is the only writer of yard transitions (G1, G2). Slotting and allocation own location and reservation state (G4). The wave and route engine constructs picking waves and outbound routes. The event ledger records every transition in the EPCIS shape. The KPI engine computes measures from captured time states (G5). The integration layer exposes REST endpoints, EDI gateways, and carrier APIs (G6).

The consequential property is negative rather than positive: no module may write a state another module owns. A dock-assignment request from Loading Dock Management is a request to the yard state machine, which validates it against door availability, staging occupancy, and the EN 1398 interlock state before granting it. The validation cannot be bypassed because there is no other write path.

3.3 The module portfolio

The WMS owns master data, inbound, put-away, picking, dispatch, and traceability. Truck Turnaround owns the gate and the yard queue. Loading Dock Management owns door scheduling and assignment. Inventory Tracking owns location confirmation and cycle counting. Order Fulfillment owns wave release, pack, and carrier handoff. Fleet Optimization owns outbound route construction.

DashboardMaster DataBarcode/QRInboundBOM/PPCPickingDispatchReportsTraceability

Master Data Management

UsersLocationsItemsVendors

Item Master

Item Type

Select Type

Item Code

Auto-generated

Item Name

Enter item name

Description

Enter description

Unit of Measure

PCS

Minimum Stock Level

Enter minimum stock

Save Item

Item Code	Name	Type	UOM	Current Stock	Min Level	Actions
RM001	Steel Rod 8mm	Raw Material	PCS	500	100	Edit View Stock

Export PDF

Date	Item Code	Item Name	Quantity	UOM	Vendor	Reference
2025-09-20	RM001	Steel Rod 8mm	500	PCS	ABC Steel Suppliers	GRN001

Figure 2. WMS master data: the Item Master. Item, location, user, and vendor masters are the governed reference data behind every transaction; the module tab strip (Dashboard, Master Data, Barcode/QR, Inbound, BOM/PPC, Picking, Dispatch, Reports, Traceability) shows the operational span served from the same store.




Figure 3. Vendor and supplier master, with an inbound goods-receipt line resolved against it. Because vendor identity is master data rather than free text on a receipt, inbound discrepancies can be attributed to a supplier rather than merely counted.

3.4 Deployment and multi-client operation

A third-party logistics operator runs several clients in one building, and each client requires its own stock segregation, its own service levels, and its own reporting, while sharing doors, staging, and labour. The store models client as a partition of stock and orders, not of the physical estate: a door is a shared resource whose assignment is arbitrated centrally.

This is what makes onboarding an integration exercise rather than a build. A new client contributes master data and an API mapping; it does not contribute a new copy of the warehouse model (G6). Section 5 reports the effect on onboarding duration.

4 Computational Methods by Module

Notation is collected in Appendix A; worked numerical examples in Appendix B.

4.1 Yard state and dwell decomposition

A vehicle occupies exactly one of six states. Transitions are performed by the system in response to an observed event — an RFID read at the gate, a slot sensor, a door interlock release, a load confirmation — and each writes its own timestamp (G2).




Figure 4. Schematic. Yard state machine. Dwell is the interval between the Arrived and Departed transitions; queuing time is the interval between Arrived and At dock. Because both endpoints are system-generated, dwell is observed rather than reported.

Figure 5. Gate entry as deployed, with live yard statistics beside it. The RFID tag identifier is captured at the barrier, so the Arrived transition and its timestamp are produced by the read rather than by the clerk.

Dwell decomposes additively, which is what makes it actionable: a reduction claim must say which component moved.

$$W = W(\text{queue}) + W(\text{dock}) + W(\text{release}) \quad (\text{dwell})$$

where $W(\text{queue})$ is Arrived to At dock, $W(\text{dock})$ is At dock to Loaded, and $W(\text{release})$ is Loaded to Departed. Re-sequencing acts on $W(\text{queue})$; dock productivity acts on $W(\text{dock})$; paperwork and seal checks act on $W(\text{release})$.

Yard occupancy follows from Little's Law, which holds for any stable system regardless of arrival or service distribution. It is used here as a capacity check rather than a forecast: if measured occupancy diverges from the product of arrival rate and dwell, the yard is not in steady state and a capacity conclusion drawn from it would be unsound.

$$L = \text{lambda} \times W \quad (\text{little})$$

where L is the mean number of vehicles in the yard, λ the mean arrival rate, and W the mean dwell time. Parking capacity must exceed L with margin for the arrival process's variability, not merely equal it.

4.2 Dock assignment as a re-sequenced queue

A fixed appointment grid holds a door for a vehicle that has not arrived while a vehicle that has arrived waits. The assignment logic here treats appointments as priority inputs to a re-sequencing decision rather than as reservations of a time slot.

On each arrival or door release, the eligible set is computed: vehicles whose cargo type matches the door's capability, whose client stock is available, and whose staging lane is free. Within that set, vehicles are ordered by a score combining appointment adherence, service-level exposure, and expected service time, and the highest-scoring vehicle is assigned.

$$s(v) = w1 \times \text{delta}(v) + w2 \times \text{rho}(v) - w3 \times \text{t-hat}(v) \quad (\text{dockscore})$$

where $\text{delta}(v)$ is lateness against the vehicle's appointment, $\text{rho}(v)$ its service-level exposure, and $\text{t-hat}(v)$ the estimated service time for its load. Weights are configured per site; the third term is negative so that, all else equal, a short job fills a short gap.

Door utilisation follows the ISO 22400-2 treatment of availability, with maintenance and blocked states excluded from available time rather than counted as idle.

$$U(\text{door}) = T(\text{busy}) / T(\text{available}), \text{ where } T(\text{available}) = T(\text{open}) - T(\text{maintenance}) - T(\text{blocked}) \quad (\text{dockutil})$$

$T(\text{blocked})$ covers intervals in which the door was serviceable but its staging lane was full, which is a warehouse constraint rather than a dock one and must not be charged to the door.




Figure 6. Queue dashboard. The queue is a live ordering recomputed on each transition, not a printed list; a vehicle that arrives early is re-sequenced into an available gap rather than held to its nominal slot.

4.3 Slotting, allocation, and the reservation model

Allocation binds a demand line to specific physical stock. The binding is a reservation held in the state store, not a note on a pick list, which is what makes a phantom pick — two tasks released against the same unit — impossible rather than merely unlikely (G4).

Slotting assigns stock to locations. The objective is to minimise expected travel, which for a class-based assignment reduces to placing high-turn items in the locations with the lowest travel cost from the depot or dispatch point.

$$\text{minimise SUM over items } i \text{ of } [f(i) \times c(\text{sigma}(i))] \quad (\text{slot})$$

where $f(i)$ is the pick frequency of item i , $\text{sigma}(i)$ the location assigned to it, and $c(l)$ the travel cost of location l . Solved as a class-based assignment — items ranked by frequency, locations ranked by cost, then matched — rather than exactly, because the ranking is re-derived weekly as demand shifts.

Slotting quality bounds routing quality. A picker route through a badly slotted warehouse is a well-optimised tour of unrelated locations, which is why allocation is solved before wave construction rather than after it.

4.4 Wave construction and picker routing

A wave is a set of orders released together so that their picks can be batched and their travel shared. Wave construction balances two costs that move in opposite directions: batching more orders reduces travel per line but increases sortation effort downstream and delays the earliest order in the batch.

Orders are grouped by shared aisle footprint and by carrier cut-off, then picker routes are constructed within each batch. Following the order-picking literature, routing uses an S-shape traversal with return-aisle exceptions rather than an exact shortest-tour solve, because the difference between the two is small relative to the variance introduced by congestion, and the heuristic can be recomputed when a wave changes mid-pick.

Fulfilment performance is reported as the fraction of order lines shipped complete within the promised window, which is the measure a client's service-level agreement is actually written against.

$$FR = (\text{lines shipped complete within window}) / (\text{lines due}) \quad (\text{fillrate})$$

computed on a line basis rather than an order basis, because a single short line fails an entire order and an order-basis measure therefore conceals how much of the shipment was correct.

4.5 Outbound route construction

Outbound routing is a capacitated vehicle routing problem with time windows. Routes are constructed with the savings heuristic: the saving from serving customers i and j on one route rather than two separate out-and-back trips is the distance avoided.

$$s(i,j) = d(0,i) + d(0,j) - d(i,j) \quad (\text{savings})$$

where 0 denotes the depot and $d(a,b)$ the travel cost between a and b. Pairs are merged in decreasing order of saving, subject to vehicle capacity and time-window feasibility, then improved by local exchange.

Route plans are recomputed on every material change rather than fixed at a morning cut-off. The justification is given in Section 2.3: an optimal plan against stale demand is worse in realisation than a good plan against current demand.




Figure 7. Yard analytics. Turnaround is reported per carrier and per traffic pattern, which lets a persistent dwell excess be attributed to a carrier's behaviour rather than absorbed into the site average.

4.6 Inventory tracking and exception-triggered counting

Rather than a scheduled freeze count, counting is triggered by exceptions: a pick that found less than expected, a location whose last confirmed scan is older than its policy age, a negative theoretical balance, or a put-away to an unexpected location. Each trigger enqueues a count task for the affected location, executed on a handheld alongside normal operations.

Accuracy is reported on a location basis, which is the basis that predicts pick failure. A value-weighted measure can look healthy while a large number of low-value locations are wrong, and it is those locations that stop a picker.

$$IRA = (\text{number of locations } L \text{ where } | \text{counted}(L) - \text{book}(L) | \leq \text{tolerance}(L)) / N \quad (\text{ira-wh})$$

where N is the number of locations counted in the period. Tolerance is per location and is zero for serialised or lot-controlled stock.

Every discrepancy is attributed — to receipt error, put-away error, pick error, unbooked scrap, or count error — because an unattributed variance produces a write-off, while an attributed one identifies the process to correct.

5 Field Evidence: A Production Deployment

The deployment described here is at a third-party logistics operator in Madhya Pradesh, India, running a multi-client distribution centre. All figures are operator-reported from the production system and are itemised with their provenance below.

5.1 Context and prior workflow

Before deployment, dock and yard scheduling ran on messaging apps and paper trip sheets. Carriers arrived early or late; doors sat idle while drivers waited in the gate queue. The incumbent warehouse management system was an original-equipment application with no API, so every client onboarding required roughly six weeks of custom data plumbing.

Inventory accuracy ran at approximately 91%, meaning about one pick in eleven involved an exception, and service-level penalties from one anchor client were consuming margin. Last-mile routes were planned once daily at 06:00 and never re-optimised; fuel cost per kilogram had risen for two consecutive quarters.

5.2 What was deployed

The WMS was deployed as the distribution-centre backbone with REST APIs for client integration. Loading Dock Management and Truck Turnaround were added for end-to-end yard visibility — appointment, gate check-in, dock assignment, departure against a service-level clock. Inventory Tracking moved counting onto handhelds and RFID so cycle counts run without halting operations. Order Fulfillment and Fleet Optimization closed the loop, so wave plans feed route plans that re-optimize as orders mature.

5.3 Reported outcomes

Table 1. Operator-reported outcomes following deployment, measured against the pre-deployment baseline at the same site.

Measure	Before	After
Average truck dwell, gate to gate	3 h 10 min	1 h 20 min
Dock idle time	Baseline	38% reduction
Inventory accuracy	91%	99.4%
New client onboarding	6 weeks	8 days
Route re-optimisation	Once daily, 06:00	On every material change

The dwell reduction and the dock-idle reduction are two views of the same mechanism. Sharing one yard state lets an early arrival be re-sequenced into a gap rather than held to its nominal appointment: the vehicle waits less, and the door that would have sat empty is filled. Neither result requires additional doors or additional labour.

The accuracy improvement is attributable to continuous exception-triggered counting replacing the annual freeze count, with handheld and RFID confirmation at put-away and pick. Errors are detected while their cause is still identifiable rather than accumulating until a shutdown.

The onboarding figure is the clearest architectural signal. It did not fall because integrations became simpler in substance; it fell because they stopped being bespoke. Where every workflow is exposed as an API with standard mappings, a new client is a configuration exercise (G6).

6 Discussion

6.1 Shared state is the load-bearing decision

Of the six design goals, G1 produces most of the observed benefit. Once the yard and the warehouse read the same state, dock assignment can be validated rather than judged, and the supervisor cross-check that previously stood between the two systems disappears — along with the delay it imposed.

The corollary is that adding modules to a split-state estate does not help. A yard management product bolted onto a warehouse management product that does not share its state reproduces the cross-check at the integration boundary; the human is removed from the decision but the round trip is not.

6.2 Measurement discipline as a precondition, not a report

The dwell figures reported in Section 5 are meaningful only because the timestamps are system-generated. An operator migrating from clerk-entered times to system-generated ones will often see measured dwell rise on day one, because the reconstructed baseline was optimistic. This is worth stating plainly to anyone evaluating such a change: the first honest number can look like a regression.

The same applies to door utilisation. Excluding maintenance and blocked intervals from available time, as ISO 22400-2 requires, typically raises reported utilisation and weakens the case for capital expenditure on more doors — which is the correct conclusion when the binding constraint was staging rather than doors.

6.3 A capability reference framework for warehouse and yard platforms

Table 2. Capability reference framework. Each dimension is a question answerable by demonstration against a running system within an hour, without access to source code.

Dimension	Question the platform must answer by demonstration
D1 Single state	Do the yard and the warehouse read one state store, or two synchronised copies?
D2 Generated timestamps	Is every state-transition time produced by the system, with no user-entered path?
D3 Label-resolvable identity	Can a unit load be resolved from its SSCC alone, with no shipment document?
D4 Reservation integrity	Can two tasks ever be released against the same physical unit? Demonstrate the attempt.
D5 Standards-defined measures	Are utilisation and availability computed from captured time states per ISO 22400-2?
D6 Dwell decomposition	Can dwell be split into queue, dock, and release components for any past vehicle?
D7 Counting model	Is counting exception-triggered and concurrent with operations, or scheduled and blocking?
D8 Onboarding cost	Is a new client an API mapping or a new data pipeline? Demonstrate with a fresh client.

6.4 Generalisability

The evidence base is one multi-client third-party logistics site. Third-party logistics is the setting where shared-state benefits are largest, because clients contend for the same doors and labour while requiring separate stock and reporting. A single-client in-plant warehouse with abundant dock capacity should expect the same architectural properties with materially smaller measured gains, particularly on dock idle time.

7 Threats to Validity and Limitations

1. Single-site evidence. All field figures come from one distribution centre run by one operator. There is no control site and no matched comparison against an alternative platform.
2. Baseline measurement asymmetry. The pre-deployment dwell baseline was reconstructed from clerk-entered records, while the post-deployment figure is system-generated. As Section 6.2 notes, this biases the comparison in the deployment's favour by an unquantified amount.
3. Operator-reported metrics. Figures are reported from the production system by the operator rather than independently audited.
4. Confounded rollout. Six modules were deployed as one programme alongside process and training change; no single module's contribution to any single figure can be isolated.

5. Seasonality not controlled. The comparison periods were not matched for seasonal volume, and third-party logistics volumes are seasonal.
6. No cost data. Operational outcomes are reported; total cost of ownership is not, so no return-on-investment claim is made or implied.

The second limitation is the most consequential and the least commonly acknowledged in vendor literature. A dwell improvement measured against a self-reported baseline is partly an artefact of the measurement change itself.

8 Future Work

- Quantifying the baseline bias. Running clerk-entered and system-generated capture in parallel for one period would bound the measurement artefact described in Section 7 and let the dwell improvement be stated net of it.
- Appointment policy under re-sequencing. If arrivals are re-sequenced anyway, the value of a tight appointment window becomes an open question; measuring dwell against appointment-adherence weighting would indicate whether tight windows still earn their coordination cost.
- Joint batching and routing. Wave construction and picker routing are currently solved in sequence; the order-picking literature indicates a joint treatment recovers further travel, at a computational cost that may now be affordable.
- EPCIS interchange. The event ledger follows the ISO/IEC 19987 event shape but does not yet emit conformant EPCIS documents for exchange with clients and carriers.
- Counting-policy learning. Trigger thresholds are configured per site; learning them from observed discrepancy rates per location class would target counting effort where accuracy actually decays.

9 Conclusion

A distribution centre fails at the seam between its yard and its warehouse, and the seam exists because the two keep separate state. This paper has described a platform that removes it: one governed state store in which every vehicle, door, slot, unit load, and stock location holds exactly one current state, every transition writes exactly one system-generated event, and six modules operate through that store rather than beside it.

We stated the computational methods of each module against published standards and the operations-research literature, and reported a production deployment in which dock idle time fell 38%, average gate-to-gate dwell fell from 3 h 10 min to 1 h 20 min, inventory accuracy rose from 91% to 99.4%, and client onboarding compressed from six weeks to eight days — while being explicit that the dwell comparison is partly confounded by the measurement change itself.

The capability reference framework of Section 6.3 is offered as the durable contribution. Platforms will change; the eight questions — about shared state, generated timestamps, label-resolvable identity, reservation integrity, standards-defined measures, dwell decomposition, counting model, and onboarding cost — will remain the right ones to put to any warehouse or yard platform.

Appendix A Nomenclature

Table 3. Symbols and abbreviations used in this paper.

Symbol / term	Meaning
W	Mean dwell time, gate to gate
W(queue), W(dock), W(release)	Dwell components: arrival to dock, dock to loaded, loaded to departure
L	Mean number of vehicles present in the yard
lambda	Mean vehicle arrival rate
s(v)	Dock-assignment score for vehicle v
delta(v)	Lateness of vehicle v against its appointment
rho(v)	Service-level exposure carried by vehicle v
t-hat(v)	Estimated service time for the load on vehicle v
U(door)	Door utilisation, busy time over available time
T(open), T(maint), T(blocked)	Opening, maintenance, and staging-blocked intervals for a door
f(i)	Pick frequency of item i
sigma(i)	Location assigned to item i by slotting
c(l)	Travel cost of location l
s(i,j)	Clarke-Wright saving from serving i and j on one route
d(a,b)	Travel cost between points a and b; 0 denotes the depot
FR	Fill rate, order lines shipped complete within window
IRA	Inventory record accuracy, on a location basis
SSCC	Serial Shipping Container Code, the GS1 logistic-unit key
GTIN / GLN	Global Trade Item Number / Global Location Number
3PL	Third-party logistics operator

Appendix B Worked Numerical Examples

Appendix B.1 Yard capacity from Little's Law

A site receives 96 vehicles across a 16-hour operating day, so $\lambda = 6$ vehicles per hour. Before deployment, mean dwell was 3 h 10 min, or $W = 3.167$ h.

Applying Equation (little): $L = 6 \times 3.167 = 19.0$ vehicles resident on average. After deployment, $W = 1.333$ h, so $L = 6 \times 1.333 = 8.0$ vehicles.

The operational consequence is larger than the dwell figure suggests. Parking must be sized for the peak, not the mean; at roughly twice the mean for a Poisson-like arrival process, the requirement falls from about 38 slots to about 16. The dwell reduction released parking capacity that would otherwise have been a construction project.

Appendix B.2 Door utilisation with correct time states

A door is open for 16 h in a day. It was under planned maintenance for 1.5 h and blocked by a full staging lane for 2.0 h. It was actively loading or unloading for 9.6 h.

Applying Equation (dockutil): $T(\text{available}) = 16 - 1.5 - 2.0 = 12.5$ h, so $U(\text{door}) = 9.6 / 12.5 = 0.768$, or 76.8%.

Had maintenance and blocking been charged to the door as idle time, utilisation would read $9.6 / 16 = 0.60$, or 60% — a figure that invites a proposal to build more doors. The correct decomposition shows the door was busy 77% of the time it could have been used, and that 2.0 h were lost to staging. The investment case is a staging-lane change, not a dock.

Appendix B.3 Route construction by savings

A depot at 0 serves customers A and B. Distances are $d(0,A) = 18$ km, $d(0,B) = 22$ km, and $d(A,B) = 9$ km.

Serving them separately costs $2 \times 18 + 2 \times 22 = 80$ km. Applying Equation (savings): $s(A,B) = 18 + 22 - 9 = 31$ km. The combined route 0 to A to B to 0 costs $18 + 9 + 22 = 49$ km, and $80 - 49 = 31$ km, matching the saving.

Merges are taken in decreasing order of saving subject to capacity and time-window feasibility. A pair with a large saving that would breach a customer's delivery window is skipped rather than accepted and repaired later, because a plan that is infeasible on paper will be repaired by the driver in a way the system cannot observe.

Appendix B.4 Why accuracy is reported per location

A count covers 2,000 locations. 12 locations are outside tolerance. Applying Equation (ira-wh): $IRA = (2000 - 12) / 2000 = 0.994$, or 99.4%.

Suppose those 12 locations hold low-value fasteners worth 0.3% of stock value. A value-weighted accuracy measure would report 99.7% and imply the position is healthier than it is. But a picker sent to any of those 12 locations still stops, and the resulting exception costs the same as one on a high-value line. Location-basis accuracy predicts pick failure; value-basis accuracy predicts write-off magnitude. Both are reported, and they are not interchangeable.

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
Dock idle time	38% reduction	Operator-reported — Tier 1 Manufacturers (3PL, Madhya Pradesh, India) /case-studies/warehouse
Average truck dwell time	1 h 20 min (from 3 h 10 min)	Operator-reported — Tier 1 Manufacturers (3PL, Madhya Pradesh, India) /case-studies/warehouse — Gate-to-gate, measured across all carriers at the distribution centre.
Inventory accuracy	99.4% (from 91%)	Operator-reported — Tier 1 Manufacturers (3PL, Madhya Pradesh, India) /case-studies/warehouse — Continuous exception-triggered cycle counting replaced scheduled freeze counts.
New client onboarding time	8 days (from 6 weeks)	Operator-reported — Tier 1 Manufacturers (3PL, Madhya Pradesh, India) /case-studies/warehouse — REST APIs replaced per-client custom data pipelines.
Truck turnaround time reduction	40-60%	Product specification /products/truck-turnaround

References

- [1] de Koster, R., Le-Duc, T., and Roodbergen, K. J. (2007). Design and control of warehouse order picking: A literature review. *European Journal of Operational Research*, 182(2), 481-501.
- [2] Gu, J., Goetschalckx, M., and McGinnis, L. F. (2007). Research on warehouse operation: A comprehensive review. *European Journal of Operational Research*, 177(1), 1-21.
- [3] Dantzig, G. B., and Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1), 80-91.
- [4] Clarke, G., and Wright, J. W. (1964). Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4), 568-581.
- [5] Little, J. D. C. (1961). A proof for the queuing formula: $L = \lambda W$. *Operations Research*, 9(3), 383-387.
- [6] GS1 (2024). GS1 General Specifications. GS1 AISBL. <https://ref.gs1.org/standards/genspecs/>

- [7] GS1 (2024). Serial Shipping Container Code (SSCC). GS1 identification key series. <https://www.gs1.org/standards/id-keys/sscc>
- [8] GS1 (2023). GS1 Logistic Label Guideline. GS1 AISBL. <https://www.gs1.org/standards/gs1-logistic-label-guideline/1-3>
- [9] International Organization for Standardization / International Electrotechnical Commission (2014). Information technology - Automatic identification and data capture techniques - Unique identification - Part 1: Individual transport units. ISO/IEC 15459-1:2014. <https://www.iso.org/standard/54779.html>
- [10] International Organization for Standardization / International Electrotechnical Commission (1999). Information technology - Unique identification of transport units - Part 2: Registration procedures. ISO/IEC 15459-2:1999. <https://www.iso.org/standard/31374.html>
- [11] International Organization for Standardization / International Electrotechnical Commission (2024). Information technology - EPC Information Services (EPCIS) Standard, version 2.0. ISO/IEC 19987:2024. <https://www.iso.org/standard/85557.html>
- [12] International Organization for Standardization / International Electrotechnical Commission (2015). Information technology - EPC Information Services (EPCIS) Standard, version 1.1. ISO/IEC 19987:2015. <https://www.iso.org/standard/66796.html>
- [13] Roodbergen, K. J., and de Koster, R. (2001). Routing methods for warehouses with multiple cross aisles. International Journal of Production Research, 39(9), 1865-1883.
- [14] Petersen, C. G. (1997). An evaluation of order picking routeing policies. International Journal of Operations & Production Management, 17(11), 1098-1111.
- [15] International Organization for Standardization (2014). Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 1: Overview, concepts and terminology. ISO 22400-1:2014. <https://www.iso.org/standard/56847.html>
- [16] International Organization for Standardization (2014). Automation systems and integration - Key performance indicators (KPIs) for manufacturing operations management - Part 2: Definitions and descriptions. ISO 22400-2:2014. <https://www.iso.org/standard/54497.html>
- [17] 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).
- [18] 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).
- [19] International Organization for Standardization (2022). Security and resilience - Security management systems - Requirements. ISO 28000:2022, published 15 March 2022. <https://www.iso.org/standard/79612.html>
- [20] European Committee for Standardization (2009). Dock levellers - Safety requirements. EN 1398:2009.