International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 - 5827
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

Event-Driven MRP at Manufacturing Scale: Governed Replanning, Coordinated Replenishment and Dispatch, and an Industrial Implementation

Nisarg Dilipkumar Patel[1]
[1]*Northeastern University, Boston, MA*
Email id: *patelnisarg.21@gmail.com*[1]

Abstract

Material requirements planning (MRP) is still at the core of discrete manufacturing, but many MRP implementations are still done as scheduled batch runs. This paper questions if event-driven architecture (EDA) can be a defensible execution model for specific planning decisions, such as inventory replenishment and outbound shipment consolidation. The search results from a documented search in Scopus, Web of Science, IEEE Xplore and ScienceDirect from January 2020 to December 2025 were processed and integrated in a structured narrative review, consisting of three streams: MRP and production-plan stability; shipment consolidation and coordinated replenishment–dispatch decisions; and event-driven, digital-twin and Industry 4.0 infrastructure. Although demand-driven MRP and rolling-horizon planning and product-driven control can enhance responsiveness, buffer positioning, review frequency and parameter calibration are still important and somewhat subjective. Different cost and service results are obtained with time-, quantity- and hybrid consolidation policies, and the inbound replenishment and outbound dispatch are found to be coupled. Technically possible to have near real-time state synchronisation, dynamic scheduling and replayable An industrial application of this principle, namely an event-triggered, multi-echelon planning platform, is then presented, including event-driven run authorisation, parallel computation under bill-of-materials and network dependency ordering, planning controls, multi-sourcing, shipment consolidation, product-lifecycle handling, and shadow validation against an incumbent planning engine before cutover. A reference architecture and its failure semantics are specified. The remaining research gap is narrowed to coordinated replenishment and dispatch under governed, event-driven planning, and in particular to controlled industrial evaluation.

Keywords: *material requirements planning; event-driven architecture; shipment consolidation; multi-echelon replenishment.*

1. Introduction

MRP considers independent demand and uses the bill of materials, inventory position, scheduled receipts, lot-sizing rules and lead-time assumptions to determine the time-phased requirements for components. The computation is resolved, the context is not. Today, connected equipment, warehouse systems, supplier portals, transport platforms and industrial internet-of-things (IoT) devices are constantly sending signals of operations, but many organisations are still updating material plans on a daily or weekly basis. Receipts, shortages, quality holds, cancellations and accumulating shipments may be known operationally between runs, but not represented in the formal plan [1]. One answer has been to reduce this gap by implementing demand-driven MRP (DDMRP), adaptive buffers, intelligent master production scheduling and digital-twin control. DDMRP reviews show a decrease in nervousness and demand amplification, and a lack of standardisation, subjective buffer positioning, and planning software adaptation costs [1]. Azzamouri et al. demonstrate that the frequency of buffer status checks is a design parameter: continuous and periodic buffer replenishment checks lead to different service and inventory levels for the same buffer parameters [26]. Research on master production scheduling reveals that some of the common causes of schedule instability are demand variability, capacity constraints, supplier performance, and forecast error [2]. Product-driven decision agents decrease the uncertainty in planning, albeit at a little extra expense [3]. Responsive is not just about how fast information

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

comes in, but about how it is modelled in the context of resource-constrained and fuzzy DDMRP models [4,5]. The second development is related to transportation. Manufacturers do not make replenishment decisions, order release and outbound dispatch decisions separately. Orders queue, vehicles have limited capacity and transport costs include a fixed component, all of which incentivize consolidation. Recent models thus integrate inbound replenishment and outbound dispatch or integrate consolidation into inventory-routing decisions [11–17]. One of the findings that is common to this work is that the performance of policy is dependent on the time frame of the state's examination. The load actually dispatched may be quite different from the load assumed in the analytical policy because a quantity threshold may be exceeded long before dispatch is authorised, and only be inspected during a periodic planning run.

The third is architectural development. EDA considers operational changes as events that are consumed by interested services asynchronously. Digital twins and cyber-physical production systems are based on the same concept and synchronise physical resources with computational models for dynamic scheduling [7–10]. Research on the supply chain suggests that digital connectivity, analytics and visibility are related to resilience, especially when disruption renders static plans useless [21-25]. Technical capability does not mean that all events need to be replanned, however. The real cost of schedule changes: resequencing, expediting, supplier set-up losses and supplier confidence erosion. This paper makes four contributions. First, it provides a search protocol that can be reproduced, and precisely specifies the evidence base (Section 2). Second, it defines where assumptions break down at the interfaces among the three streams, and what a consolidation quantity threshold is and when a dispatch timer begins (Section 4). Third, it describes an industrial application of governed event-driven planning, as well as a reference architecture and explicit failure semantics (Sections 6 and 7). Fourth, it focuses the research agenda on coordinated replenishment and dispatch under governed, event-driven planning and to the controlled industrial evaluation still missing (Section 8). The point that is made throughout is that the operational data can change on an ongoing basis, but the re-planning can only be authorised selectively, to ensure that committed orders and schedule stability are maintained.

2. Review Scope and Method

2.1.Search protocol

The review adopts a structured narrative design, which was selected due to the fact that the relevant work is spread across production planning, operations research, transportation, information systems and manufacturing digitalisation, and that the outcome measures are incommensurable and thus do not allow for meta-analysis. Four bibliographic databases were searched: Scopus, Web of Science, IEEE Xplore and ScienceDirect; and backward and forward citation chasing was performed on the included records using Google Scholar. The search was carried out on [search date to be confirmed] and includes material published from January 2020 to December 2025. Three search strings were executed, one for each stream, in the title, abstract and keywords fields: S1 (planning): (“material requirements planning” OR “MRP” OR “demand-driven MRP” OR “DDMRP” OR “master production schedul*”) AND (nervousness OR stability OR buffer* OR “rolling horizon” OR replanning). S2 (logistics): (“shipment consolidation” OR “order consolidation” OR “dispatch polic*” OR “inventory routing”) AND (“time-based” OR “quantity-based” OR hybrid OR replenishment OR dispatch). S3 (architecture): (“event-driven architecture” OR “event stream*” OR “digital twin” OR “Industry 4.0”) AND (manufactur* OR “production control” OR “supply chain” OR scheduling). The titles and abstracts were scanned for inclusion, and full texts were read for those that met inclusion criteria. The following had to be included: (i) published between 2020 and 2025; (ii) a peer-reviewed journal article, review or analytical model, or a conference paper that explicitly addresses system architecture or integration with production control; (iii) relevance to discrete manufacturing or to a decision transferable to it; and (iv) availability in English or Spanish. Records were removed when the only application was

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

process-industry campaign scheduling, when the material was pedagogical or vendor-authored, or when no model, data or architecture was reported. If multiple papers were published by the same group the most detailed model was kept. Twenty-five publications met these criteria. During the revision process one additional paper [26] was included, as it directly answers the review-frequency question and there was no equivalent source available, leaving the synthesis based on 26 publications. A previous version of this paper stated the coverage as 2020-2026. There is no publication included dated 2026 and the window is stated here as 2020-2025. Screening was done by the author only and no inter-rater agreement was calculated. The number of records at each screening stage was not kept and should be reconstructed and reported prior to submission (Appendix A, item F2).

2.2.Scope limits

Three limits apply. First, the review addresses discrete manufacturing, where materials, bills of materials and shipment units are enumerable; process-industry campaign scheduling is out of scope. Second, "event-driven MRP" denotes an execution architecture around MRP planning logic, not a replacement for the MRP calculation itself. Third, the review separates demonstrated technical feasibility from demonstrated operational benefit: a platform may process events in milliseconds while the appropriate planning cadence remains hourly or exception-based.

Table 1 Contemporary evidence streams relevant to event-driven MRP

Stream	Focus	References	Implication for event-driven MRP
A	DDMRP buffers	[1,4–6,26]	Reliable buffer settings and regular adjustment
A	Planning uncertainty	[2,3]	Balance responsiveness and schedule stability
B	Replenishment and dispatch	[11–13]	Coordinate inventory and dispatch timing
B	Intraday demand	[14,15]	Capture demand changes within the day
B	Consolidation and routing	[16,17]	Combine live data with transport constraints
C	Scheduling and digital twins	[7–10]	Enable rescheduling with reliable, integrated models
C	Event processing	[18–20]	Ensure timely, ordered and duplicate-safe processing; architectural evidence
C	Digital resilience	[21–25]	Support recovery through visibility and decision capability

3. Stream A: The Plan Stability

3.1.Periodic MRP with demand-driven control

Recent MRP research does not refute requirements planning, it simply moves the point of absorbing variability and the point at which replenishment is triggered. DDMRP places buffers at strategic locations and replenishes them based on buffer status. Butturi et al. attribute the following advantages to the method: better responsiveness, less bullwhip effect, and a lack of standardisation, subjective buffer choice and adaptation costs of the planning software [1]. Xu et al. add resource constraints to demand-driven production planning and demonstrate that demand signals can't be translated to feasible plans without representing capacity [4]. Araz et al. use fuzzy logic for model inputs, and clearly state that buffer calculations are

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

only estimates and not facts [5]. Fernandes et al. compare buffer-status signals in shop-floor scheduling and demonstrate that the selection of the signal has a significant impact on performance [6]. Azzamouri et al. introduce the element most pertinent here: under the assumption of constant buffer parameters, the difference between continuous and periodic buffer status evaluation has an impact on the service and inventory performance [26]. Frequency of review is thus a decision variable in its own right and one that an event-driven architecture can manipulate at will — and that's why it must be governed. The results of these tests are relevant to EDA. The architecture can be notified of an inventory change instantaneously, but the decision will still be based on buffer definitions, demand interpretation, capacity and order status. The idea that each raw transaction is a planning event is mixing relevance of data to a decision with the availability of data. A more defensible design would be to have business events (such as "net-flow position entered the red zone", "confirmed receipt missed tolerance", "available-to-promise fell below committed demand") be derived from lower level transactions and plan on those.

3.2.Plan nervousness and schedule stability

Forecasts, orders, capacity and supply conditions change and master production schedules are revised accordingly. Luna-Victoria-León et al. demonstrate that uncertainty is inherent in manufacturing planning and that a number of interacting factors influence the quality of the schedule [2]. According to Sáez et al. [3] intelligent, product-driven interventions reduce that uncertainty at higher costs. The trade-off is important: a plan might look more responsive but have hidden costs in extra set-ups, resequencing, procurement changes, overtime and expediting.

With batch execution, changes are stored and processed at predefined time intervals. Event-driven execution, on the other hand, provides lots of little signals. When all signals lead to a complete re-explosion and order revision, the number of proposed changes can increase quickly even if there has been little net change in demand. Stability therefore has to be designed into the execution architecture and cannot be taken for granted. Some examples of suitable controls are: time fences, minimum change thresholds, hysteresis bands, debounce windows, event aggregation, explicit cost-of-change penalties, and approval rules for orders already committed externally.

The idea is that information state can be changed continuously, but decision state only when an event crosses an economically meaningful boundary – this is called selective responsiveness. If a one-unit change in inventory is made, then the projected availability should be updated immediately, but not necessarily the supplier schedule; if a machine breaks down, and that capacity is no longer available, then replanning may be necessary. Events need to be categorized according to severity, horizon, reversibility and planning level impacted.

3.3.Continuous review in a planning hierarchy

Digital-twin research shows that dynamic production scheduling with real-time synchronisation is achievable [7]. Ullah and Younas present digital-twin control for flexible manufacturing systems [8], Ojstersek et al. optimise smart manufacturing systems through digital twins [9] and Lin combines simulation and scheduling in a cyber-physical digital-twin environment [10]. These studies outline architectures in which the actual conditions are used to update a model and invoke consideration of scheduling alternatives. But, the decision making process has varying natural rhythms in the manufacturing hierarchy. In seconds inventory forecasts are changed, in minutes shop-floor sequences are changed, purchasing commitments are frozen for days, and aggregate capacity plans are changed every month. One layer accelerated does not provide end-to-end responsiveness, it typically moves the constraint. Event-driven MRP should thus be based on a hierarchy of triggers. Local events are first fed to state and then rules are applied to decide whether a policy should be recalculated, work centre rescheduled, material family replanned, or exception escalated to a higher planning level. This helps to avoid enterprise-wide nervousness to each operational variation.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

4. Stream B: The shipment Consolidation & Dispatch Policy

4.1.Policies and recent evidence

Consolidation of shipments is a compromise between transport economies and delivery delay. A quantity-based policy sends dispatches when a certain number of customers have ordered; a time-based policy sends dispatches after a specified time period; a hybrid policy sends dispatches when either of these two events occurs. These policies have come back into the spotlight in data-rich, capacity-constrained environments. Wei, Çetinkaya and Cline develop an analytical model to coordinate inbound replenishment and outbound dispatch decisions under hybrid consolidation and compare the policies [11]. Memon et al. suggest hybrid time and quantity algorithms which minimize delivery cost [12]. Malmberg and Marklund propose evaluation and control approaches for quantity-based consolidation in inventory-distribution systems [13]. As a whole, these studies suggest that a combination of service and economic triggers should be used, as opposed to a single rule. The issue is continued in inventory routing. Alarcon Ortega et al. model the stochastic dynamic demand with intraday depletion and demonstrate that the period-level aggregation can hide operationally important timing [14]. Alvarez et al. consider inventory routing in the presence of uncertain supply and demand [15]. The works of Yang et al. [16] and Yosephine et al. [17] consider order consolidation and vehicle-routing optimisation and dynamic e-commerce demand, respectively. While these settings are not factory MRP, they do define a transferable outcome: dispatch thresholds cannot be optimised without considering routes, capacities, inventory implications, and changing demand.

4.2.The quantity measure

The analytical treatments typically define a time limit T and a threshold Q, without specifying the units of measurement for each and the time at which the clock begins. Both are implementation decisions and both affect the policy that is actually implemented. The total number of dispatchable demand for a consolidation group from the last dispatch from that group until the present time is quantity threshold Q. There are three options which need to be made clear. Firstly, the grouping key, usually origin, destination, service level and, if applicable, carrier or mode: order lines with different grouping keys must not be consolidated and must not be accumulated together. Second, the unit of measure, whether order lines, pieces, weight, cubic volume, or handling units like pallets or containers: Only a measure commensurate with the binding transport constraint is meaningful; otherwise, the number of order lines loaded on a weight-limited vehicle is not meaningful, because the load is expressed in a different unit of measure. Third, the reset rule: the value of the accumulator resets at dispatch and a partially dispatched group keeps its remaining value. Timer T is an age, NOT a calendar schedule. It begins at the time of the earliest eligibility of the oldest line in the group that has not yet shipped (when the line first became dispatchable) and resets when that line ships. A different policy is to start the timer at a fixed clock epoch (or at the time of the last dispatch), in which case a line created just before the window closes gets almost the entire waiting allowance, and a line created just after gets almost nothing. If the service commitment is stated against a customer promise date, and not an order date, the trigger should be the time remaining until the customer promise date, thus turning the trigger from elapsed time into a deadline. These formulations are not the same and the formulation in force should be reported with any result.

4.3.Eligible demand - load selected for dispatch

A trigger is for a candidate not a shipment. There are three sets to be distinguished. Eligible demand set includes all order lines with the same consolidation key, with inventory allocated or allocatable and with no hold, quality or credit screen. The set that is triggered is the set that exists when Q or T fires. It is usually not the first two, but the dispatched load, that is tendered. There are four mechanisms which separate them. The minimum and maximum quantity windows ensure that a group is not tendered below the minimum quantity (which is uneconomic) and that a group is not tendered above the maximum quantity (which is

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

the capacity of the equipment or tender agreement) and where the quantity is split, a residual is left. Capacity (vehicle weight, cubic volume, dock slots and carrier cut-off times) limits the selectable load, regardless of the trigger. Allocation can remove inventory promised to a higher-priority line between evaluation and tender, while pull-ahead can add a line that is not yet due, but is within a specified pull-ahead horizon, to fill any remaining capacity, thereby lowering the demand in later periods. It makes sense that a time trigger isn't a dispatch trigger. The system is obligated to evaluate, NOT ship, at the time of expiration of T. The right answer is to record and explain a deferral with an exception raised against the service commitment, not to force a tender, if the qualified load is less than the minimum window, if there is no available capacity or dock slot prior to carrier cut-off, or if inventory has been reallocated. It is crucial to distinguish between trigger events/evaluation results and actual tenders, however. Combining them is hyperbole on the provision of services and a cover for the true cost of the policy. As shown in Table 2. Shipment-consolidation policies under event-driven execution.

Table 2 Shipment-consolidation policies under event-driven execution

Policy	Trigger	Cited evidence	Proposed extension	Implementation risk
Immediate dispatch	Each released line; Q = 1	No-consolidation baseline [16,17]	Control policy; check minimum window	High costs, low utilisation, congestion
Quantity-based	Eligible quantity reaches Q	Quantity control and grouping [13,16]	Use projected quantity; limit overshoot	Premature triggers; capacity constraints
Time-based	Oldest line reaches age T	Intraday timing affects service [14,15]	Apply deadlines; record deferrals	Low loads; missed arrivals
Hybrid time-and-quantity	Q or T reached first	Joint replenishment and dispatch [11,12]	Coordinate events; stabilise replenishment	Clock mismatch; duplicate processing
Predictive hybrid	Q or T plus forecast conditions	Demand-informed routing [14,16,17]	Pull ahead within firm commitments	Forecast error; delays; untracked overrides

Columns 3 and 4 are separated deliberately: column 3 reports what the cited sources establish, while .4 Execution semantics Order-created, order-released, cancellation, inventory-available, vehicle-assigned, dock-capacity and carrier-cut-off events should only affect a single consistent consolidation state. Idempotence is a must: no order should be allocated twice in the event of a duplication. Ordering per shipment or per destination is the same, as a cancelled order that is cancelled before the order it cancels will create an invalid load. The separation mentioned in §4.3 is evident when evaluating and tendering, so dispatch needs to be atomic, meaning that lines selected are

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

reserved in one transactional scope and a result is posted downstream for inventory and transport services to reconcile.

5. Event-Driven and Digital Manufacturing Infrastructure (Stream C)

EDA is built around events that are notifications of a meaningful state change. Events are published by producers, and consumers update state or start a workflow. Architectural literature depicts EDA as a basis for scalable and responsive integration [18,19] and event-driven data pipelines as a way to process continuously arriving data [20]. These are not empirical manufacturing studies, but design oriented. MRP has four capabilities that are provided by a durable event stream. It cuts down on polling by informing planning services of change. An ordered history provides auditability, allowing a planner to see which events resulted in a recommendation. Replay allows to rebuild projected state after a logic correction or a service recovery. Independent consumers can provide inventory projection, consolidation, alerting and analytics without tying each producer to each application. These benefits can only be realised through governance. Events should be about facts in the domain, not facts in the database. Identifiers, occurrence and processing time, version information, units of measure, source and correlation fields are required for schemas. Consumers need to be tolerant of compatible evolution, and reject incompatible changes of meaning explicitly. The duplicate and out-of-order events instead of as faults. The results of planning should be documented as the input-event boundary, the policy version, the parameter set and the constraint snapshot that were used in the planning computation. The manufacturing bridge is provided by digital-twin studies. Ouahabi et al. discuss the digital twins for dynamic production scheduling and emphasize the importance of real-time data management and model integration and disruption response [7]. Physical–digital synchronisation can be used for adaptive control in flexible systems, optimisation and integrated simulation [8–10]. These systems are similar to event-driven MRP, in that they rely on current state, and they have the same drawback: model fidelity. A quick, but wrong, model is worse than a slower, but correct, model. An organisational perspective is added through industry 4.0 and supply-chain-resilience studies. Ghobakhloo et al. align digital transformation with capabilities of resilience [21]. Iftikhar et al. combine the literature on digital innovation, analytics and supply-chain resilience [22]. Ivanov and Dolgui put disruption management in the context of digital supply-chain twins [23]. Frederico et al. define Supply Chain 4.0 in the context of digital capability and performance [24] and Belhadi et al. discuss resilience strategies in manufacturing and service supply chains following a major disruption [25]. In all this literature, visibility is a means to an end, not an end in itself, but it is a means to a preparedness, coordination or recovery decision. The synthesis indicates a directed hybrid structure. Orders, inventory, production status and shipments continue to be authoritative in the transactional systems and changes are appended to an event stream. Keeping decision-ready views: net-flow position, component availability, accumulated shipment volume, short-term capacity is done with a state-projection layer. A policy layer is used to determine thresholds and constraints. The damping layer is used to identify if a proposed action is material, stable and permissible within the active time fences. The actions that are accepted are performed and the results are returned as new events. Four tensions follow. Observation cadence is not decision cadence: state can be observed continuously without licensing continuous schedule change. Second, local optimality does not imply system feasibility; a locally economical consolidation threshold can be in violation of vehicle, dock, production or customer constraints. Third, responsiveness does not mean to be nervous — replanning is warranted only when the benefits of change are greater than the cost of change. Fourth, technical consistency is not business consistency: services can process events reliably with conflicting views of what is available, ordered or received. As shown in The three streams and where the tensions occur are summarised in figure 1.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

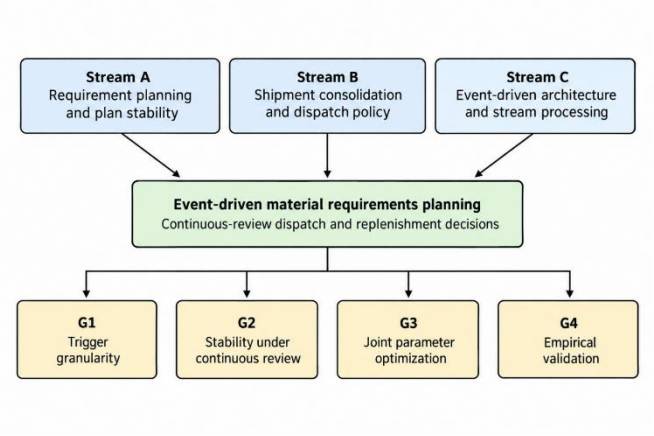

Figure 1 The figure depicts the three contemporary literature streams converge on event-driven material requirements planning

6. Industrial Implementation

6.1.Evidentiary status

The author participated in the planning platform at Tesla, Inc., in the design and implementation of the capabilities described in this section [role title, team and dates to be confirmed – Appendix A, item F3]. It is an architectural description of the account, with no details of volumes, costs, run times or performance results being disclosed. No superiority over the existing system is asserted. The validation in §6.9 was not a measure of cost or service improvement, and the absence of such a measure is the gap identified in §8.

The platform is included here as it exemplifies the main thesis of the paper. Operational data is continually updated as events occur; replanning is authorised selectively, with explicit controls to prevent the loss of committed orders and schedule stability. The sections below follow the pipeline of Figure 2.

6.2.Event-triggered planning

The platform eliminated one time, one night planning run and replaced it with many incremental runs based on change. Demand events, such as new sales and service orders, changes to orders, and order cancellations, as well as forecast changes are ingested and applied to projected state immediately, and the same applies to supply events (purchase order confirmations, receipts, shipment and advance-ship-notice events, and supplier commit changes) and inventory events (on-hand changes, scrap, quality holds, and transfers between stocking locations). Planning run is authorised only when accumulated change for part group exceeds a materiality test and not for each event. A floor cadence also runs, so that a part group that is not getting any material change is refreshed on a known cadence and no plan is allowed to age forever. This is the selective-responsiveness principle of §3.2 realised in code: state is continuous, replanning is discrete and governed — Appendix A, item F4.

6.3.Scalable computation

Planning is broken down into part groups, where no bill of materials or network relationship crosses the group boundary within the planning horizon. Independent groups are done in parallel. Ordering is done on two axes within a group. Bill-of-materials level guarantees that the planned orders for a parent item are available before the dependent demand is netted, thus preventing the calculation of dependent demand from an incomplete parent plan. Network dependency tier allows for planning of a supplying location before the locations it supplies, and a transfer requirement will be visible to the source before that source commits its supply. The dependency graph is dynamically generated based on the current bill-of-materials and network definitions, and thus changes in structure (e.g. a new sourcing lane, a new bill) do not require re-partitioning. Thus, the dependency graph is a limit on achievable parallelism, not on hardware.

6.4.Multi-echelon planning

Demand is propagated across factories, distribution centres and service locations. A requirement at a downstream location is dependent on its supplying location, which can be met by either an internal transfer or an external purchase. Both are planned in the same pass, and the lead time, calendar and capacity for the internal lane is transferred to the transfer order and the supplier lead time, calendar and effective dates are transferred to the purchase requisition. Planning together means that if there is a shortage at a distribution centre, it can be solved by redirecting from a sister centre, where it can be done faster or cheaper, but not by raising a purchase order which is received after the requirement date. The echelon structure also defines what events are

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

relevant to what locations, thus an inventory change at one service location will not trigger a replanning run across the network.

6.5.Planning controls

There are several controls to assure stability and executability of the output, and all of these are applied before any output is written, so the damping described in §3.2 is in the execution path, not in a reporting filter. Frozen and firm horizons keep near-term commitments: In the frozen horizon the plan remains unchanged; in the firm horizon, the plan is proposed but must be approved before it can be released. Dynamic safety stock is not a static quantity, but is re-calculated based on the latest variability in demand and supply and protection follows the signal, not an annual assumption. Planned quantities are influenced by fixed and minimum order quantities and are converted to what can be purchased or produced. Coordinated ordering ratios ensure that parts which need to be ordered together are ordered in the required ratio and are not released as a set incomplete. Demand smoothing helps to reduce transient spikes that can trigger a step change in the bill of materials and into supplier schedules.

6.6.Multi-sourcing

If there are multiple approved sources for a part, the requirements are distributed either by percentage or supplier priority. Allocation honours supplier capacity limits and effective dates: if a supplier is being phased in, it will only receive volume from the date that it is effective and only up to the supplier's committed capacity; if the supplier is full, the volume will be passed to the next supplier in the ordering. Allocation is calculated on every authorised run, and thus a reduction in capacity or an agreement expiration will impact the next run, not the next overnight cycle — one of the more straightforward examples where event-driven execution affects the result, not just the timing.

6.7.Shipment consolidation

Eligible order lines are clustered together and released within minimum and maximum quantity windows, and pull-ahead is used to make up the remainder of the capacity from order lines scheduled a little bit later than the minimum window. The separation described in §4.3 is implemented literally: Eligibility screening, trigger evaluation and load selection are separate steps and a line that contributed to the accumulator is not guaranteed to be on the resulting load. Lines which are displaced by a maximum window or by capacity are carried over to the group as a residual and re-evaluated at the next evaluation.

6.8.Product lifecycle

The platform supports transitions which a steady-state MRP calculation does not represent. Part supersession is used to transfer the demand for a superseded part to its successor on a defined date, where the supersession is one way, any remaining inventory of the superseded part will be consumed before the successor is planned, thus the supersession will not leave any stranded stock. Use seed stock pre-positioning of inventory at service locations for parts with no demand history, planning for intended consumption. When the demand for a service part is sparse and intermittent, instead of time-phased netting, min/max replenishment is used, because a reorder-point policy is more stable than a forecast-driven plan when the forecast error is greater than the signal, and the demand for the service part is lumpy.

6.9.Validation and execution

Cutover was not treated as a release event, but rather an evidence problem. The new engine was operated side by side with the incumbent engine and outputs were compared part, location, quantity, due date and source. Differences were not simply tallied, but were categorized: differences that were the result of a conscious logic change were distinguished from differences that were not explained; and differences that were not explained were not accepted as noise but were resolved prior to cutover. The comparison period and scope and agreement threshold for the cutover decision should be confirmed – Appendix A, item F5. This was made possible by two execution properties. Writes are idempotent, meaning that if a write is replayed, then planned orders and transfers will not be duplicated. Run-level state tracking stores, per run, the identifier of the run, the input event watermark that it consumed, the versions of the parameter and policy that it was operating under, and the completion status. This means that if a run

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

fails, it can be retried without manual intervention, and any planned order can be traced back to the run and the inputs that generated it, making the auditability requirement of §5 real.

7. Reference Architecture and Operational Semantics

The implementation is generalised in a reference architecture as illustrated in figure 2. The decision flow is carried by nine stages and there are two cross-cutting bands that apply to all stages. Events published by the systems of record are ingested and normalised (1), projected into decision-ready state (2), and tested by a policy layer, which decides if the event should only be used to update state or also trigger a replanning run (3). Authorized runs are planning controls before any write (4), parallel computation (5) and propagation of requirements across echelons with multi-sourcing allocation (6). Eligibility is consolidated on the screen (7), the dispatched load is written (8), confirmations, exceptions and overrides are re-entered as events (9). Orchestration and validation are bands, not stages – they limit all stages. As shown in Figure 2 Reference architecture for governed event-driven MRP

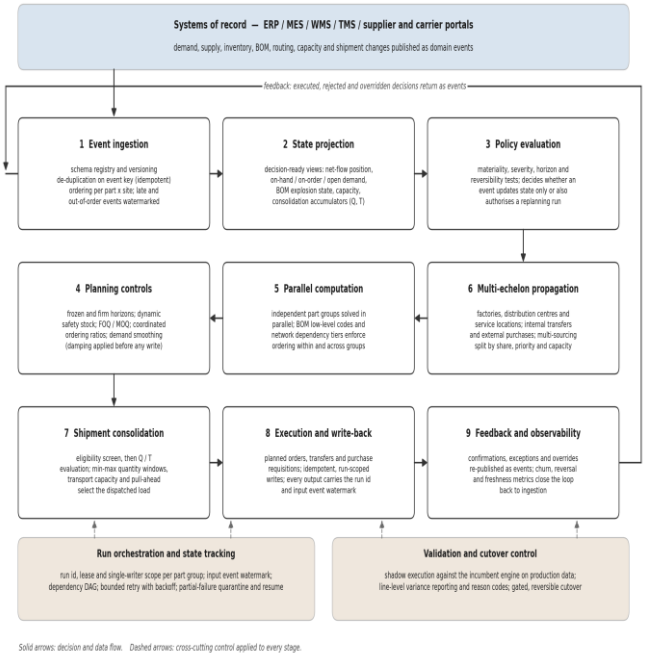

Figure 2 Reference architecture for governed event-driven MRP

The maximum number of runs that can be authorised for overlapping scope is 2 runs per first run. The architecture deals with this by assigning a single-writer lease per part group: a run gets leases for its groups, a second run on a leased group waits for the lease or is combined with the pending run, rather than running concurrently. Groups with different and non-overlapping scopes operate in parallel, without coordination. Waiting or merging of the production system should be confirmed – Appendix A, item F7. The run keeps track of the watermark of the input event stream it consumed. Only if that watermark is still current for the scope are results written, otherwise the results are discarded or rebased, depending on whether the events are material under the stage-3 policy test. This means that a slower, older run does not overwrite a newer plan, which is a failure mode that doesn't occur in batch execution as only one run is available, and is therefore easy to overlook when migrating from batch. Bounded exponential backoff is used for transient failures. Since writes are idempotent on the run identifier and the output key, a retry after a partial write will not duplicate orders. Retry budgets are limited: a run that uses up its budget is failed explicitly and is reported, not left in some indeterminate state that a planner has to find out from the output. Some part groups can be successful and others can be unsuccessful. Successful groups are committed; failed groups are quarantined with their inputs and an error classification retained; and groups that depend on a quarantined group are held, rather than planned on incomplete upstream results. Quarantined scope is picked up by the next authorised run. The other option, failing the entire run, was turned down as it would mean that the blast radius of a single malformed part is the entire plan, which is the exact opposite of what event-driven decomposition is supposed to do. The author would like to offer two suggestions for this design before it is submitted: 1) that the wait-versus-merge choice above be made automatic on the next run, or by planner action, or 2) that the distribution of run durations be observed to determine whether this choice is ever binding in practice [Appendix A, items F8 and F9].

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

8. The research gaps and agenda

The three streams are each developed and are not connected with each other. The remaining gap, narrowed as it is by the implementation reported above, is specific: co-ordinated replenishment and dispatch under governed, event driven planning, and most importantly its assessment under industrial conditions. The implementation shows that it is possible to create such a system at scale and that it can generate the same output as an incumbent engine on production data. It does not show that governed event-driven execution provides better cost, service or stability results than periodic execution, as no controlled comparison of results was performed. This is the key missing piece of evidence and the four priorities below reflect this. G1. Set trigger granularity and trigger overshoot. Comparing transaction level, aggregated and exception level triggers. The simulation should quantify threshold overshoot over review intervals, demand processes, vehicle capacities and event-processing delays, and should explicitly specify Q and T as defined in §4.2 to be able to compare results between studies. Outcomes should be: transport cost, fill rate, order waiting time, split shipments, comp. load, number of plan changes. G2. Stability: Continuous review. Nervousness controls should be checked on a lot of little events not a few big events. Experimentally evaluate time fences, hysteresis, debounce windows, change-cost penalties and severity-based escalation. Azzamouri et al. [26] give the natural baseline (no damping controls) for the effect of continuous versus periodic buffer assessment under fixed parameters; an open question is how this effect behaves when damping controls are added. Stability measures should reflect not only changes in inventory cost but also changes in the quantity and timing of orders, set-up churn, and planner interventions and revisions in supplier messages. G3. Joint parameter optimisation. Normally, the buffer zones, safety stocks, consolidation thresholds, dispatch time limits and event-trigger tolerances are designed independently, but literature indicates that replenishment, routing and dispatch are interrelated. Service targets, capacity, uncertainty, schedule-change cost and processing cadence should be taken into account in joint optimisation. Parameters that are optimised under periodic review should not be simply assumed to be valid for continuous monitoring, as the review period is also one of the parameters being optimised. G4. Controlled industrial evaluation. This is the main focus. The missing study is the one where the execution model is the only variable, and the planning logic, data and parameters are identical between both. One bounded decision (outbound consolidation is the obvious choice) can be enabled for a treatment group of parts or destinations while similar products or destinations continue to be controlled periodically by an event-driven decision service running in shadow mode. Evaluation should cover multiple demand cycles and include cost, service, stability, planner workload, exception volume and exception failure recovery. The instrumentation described in §6 demonstrates that the instrumentation needed to do such a study (run-level state tracking, line-level output comparison, reason-coded differences) is feasible in production, which eliminates the most common practical objection to conducting such a study. Future studies should also report transparently the architecture: event definitions, delivery guarantees, ordering scope, handling of duplicate and late events, schema versions, replay procedure, computation latency, policy parameters and human override rules. These are essential to ensure that results can be reproduced and not confused with the impact of other digitalisation taking place simultaneously.

9. Managerial Implications

It is not a good idea for manufacturers to start enterprise-wide continuous replanning. A better option is to choose a decision with high information latency, distinct outcomes and a small blast radius. Shipment Consolidation is a good candidate as there is trigger logic, transport cost, waiting time and utilisation that are all measurable. The events can be instrumented and the policy state and recommendations can be tested on a staged pattern as described in §6.9 before switching on the automated dispatch. There needs to be governance before scaling. Events should have an owner and a clear purpose, each automated action should have a

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 10 October 2026
Page No: 5828 – 1203
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0754

clear measurable goal, a constraint set, a fall-back and an audit trail. Time fences should be used to cover supplier and production commitments. Planners need to be able to see the reasons for the system to act and the inputs. Operational dashboards should differentiate between state freshness and the quality of the decisions, and track for duplicates, late events, model drift and manual overrides/reversals. The benefits are lower safety stock, fewer stockouts, improved vehicle utilisation, lower premium freight and quicker disruption response. Integration, event governance, model maintenance, computation, extra set-ups, expediting and planner workload/supplier schedule changes are all relevant costs. Event-driven execution is only warranted where the value of acting sooner is greater than the value of acting later for the decision at hand (not because the technology is available [18–22].

Conclusion

The elements of an event-driven MRP have been introduced in recent literature but not to the point of an unconditional replacement of periodic planning. MRP and DDMRP studies indicate that three MRP design decisions – location of decoupling buffers, representation of capacity constraints and frequency of buffer status re-evaluation – impact planning performance and can be altered independently of the speed of information flow. Logistics research shows that replenishment and consolidation/routing are linked and that the results of the policy depend on the time and the unit in which the thresholds are evaluated. Architectural and digital-twin research indicates that operational state can be synced and decisions recomputed rapidly, and resilience research indicates that visibility is only beneficial if it influences a decision. The defensible model is a hybrid and is governed. Events are continuously updated, policy rules are used to determine which changes are material, damping ensures that committed orders are not impacted, and bounded decision services only recompute the scope of changes impacted by the changes made, while periodic planning continues to provide broad coordination. This model is realised in §6, where its manufacturing scalability, event triggered run authorisation, parallel computation with dependency ordering, multi-echelon propagation, explicit planning controls and shadow validation before cutover are all realised. What it doesn't provide—and what the literature doesn't provide—is a controlled comparison of governed event-driven execution versus periodic execution, with the same planning logic. In the absence of this evidence, the most that can be done is to show that event-driven MRP is technically feasible and theoretically sound, but whether it is economically feasible is still a research question.

References

[1]. Butturi MA, De Rosa G, Balugani E, Gamberini R. Understanding the Demand Driven Material Requirements Planning scope of application: a critical literature review. Proceedings of the 32nd DAAAM International Symposium. 2021. doi:10.2507/32nd.daaam.proceedings.067.

[2]. Luna-Victoria-León GA, Collao-Díaz MF, Corzo-Chávez JA, Meza-Ortiz RN. El plan maestro de producción y la incertidumbre en el sector manufactura: una revisión de la literatura. Ingeniería Industrial. 2022;43. doi:10.26439/ing.ind2022.n43.6108.

[3]. Sáez P, Herrera C, Parada V. Reducing nervousness in master production planning: a systematic approach incorporating product-driven strategies. Algorithms. 2023;16:386. doi:10.3390/a16080386.

[4]. Xu G, Guan Z, Yue L, Mumtaz J. An efficient production planning approach based on demand driven MRP under resource constraints. International Journal of Industrial Engineering Computations. 2023;14(4):743–758 [volume, issue and page range to be verified]. doi:10.5267/j.ijiec.2023.5.003.

[5]. Araz OU, Ilgin MA, Eski OY, Araz C. Fuzzy demand-driven material requirements planning: a comprehensive analysis of fuzzy logic implementation in DDMRP. International Journal of Production Research. 2024 [volume and page range to be added]. doi:10.1080/00207543.2024.2328770.

[6]. Fernandes NO, Guedes N, Thürer M [author list to be confirmed against the published record — Appendix A, item F10]. Demand Driven Material Requirements Planning: using the buffer status to schedule replenishment orders. In: Rocha Á, Adeli H, Dzemyda G, Moreira F, Poniszewska-Marańda A, editors. Information Systems and Technologies: WorldCIST 2023. Lecture Notes in Networks and Systems. Cham: Springer; 2024. doi:10.1007/978-3-031-45642-8_41.

[7]. Ouahabi N, Chebak A, Kamach O, Laayati O, Zegrari M. Leveraging digital twin into dynamic production scheduling: a review. Robotics and Computer-Integrated Manufacturing. 2024;89:102778. doi:10.1016/j.rcim.2024.102778.

[8]. Ullah A, Younas M. Development and application of digital twin control in flexible manufacturing systems. Journal of Manufacturing and Materials Processing. 2024;8:214. doi:10.3390/jmmp8050214.

[9]. Ojstersek R, Javernik A, Buchmeister B. Optimizing smart manufacturing systems using digital twin. Advances in Production Engineering & Management. 2023;18. doi:10.14743/apem2023.4.486.

[10]. Lin WD. An integrated digital twin simulation and scheduling system under cyber-physical digital twin environment. Proceedings of IEEE IEEM. 2022. doi:10.1109/IEEM55944.2022.9989841.

[11]. Wei B, Çetinkaya S, Cline DBH. Inbound replenishment and outbound dispatch decisions under hybrid shipment consolidation policies: an analytical model and comparison. Transportation Research Part E: Logistics and Transportation Review. 2023;175:103135. doi:10.1016/j.tre.2023.103135.

[12]. Memon MA, Shaikh A, Sulaiman A, Alghamdi A, Alrizq M, Archimède B. Time and quantity based hybrid consolidation algorithms for reduced cost products delivery. Computers, Materials & Continua. 2021 [volume and page range to be added]. doi:10.32604/CMC.2021.017653.

[13]. Malmberg F, Marklund J. Evaluation and control of inventory distribution systems with quantity-based shipment consolidation. Naval Research Logistics. 2022 [volume and page range to be added]. doi:10.1002/nav.22090.

[14]. Alarcon Ortega E, Malicki S, Doerner KF, Minner S. Stochastic inventory routing with dynamic demands and intra-day depletion. Computers & Operations Research. 2024;163:106503. doi:10.1016/j.cor.2023.106503.

[15]. Alvarez A, Cordeau JF, Jans R, Munari P, Morabito R. Inventory routing under stochastic supply and demand. Omega. 2021;102:102304. doi:10.1016/j.omega.2020.102304.

[16]. Yang C, Lee Y, Lee C. Data-driven order consolidation with vehicle routing optimization. Sustainability. 2025;17:848. doi:10.3390/su17030848.

[17]. Yosephine I, Shin YI, Lee C. Integrated shipment consolidation and vehicle routing problem for profit maximization in e-commerce supply chains with dynamic demand. Journal of the Korean Society of Supply Chain Management. 2024;24:1–16. doi:10.25052/kscm.2024.05.24.1.1.

[18]. Nalam SMVV. Event-driven architecture: the backbone of real-time enterprise integration. International Journal of Computational and Experimental Science and Engineering. 2025. doi:10.22399/ijcesen.4178.

[19]. Kiran VNS. Event-driven architecture: building responsive and scalable systems. International Journal of Science and Research. 2021. doi:10.21275/SR24716231109.

[20]. Kundavaram VNK. Event-driven data pipelines: a cloud-based approach to real-time data processing. International Journal of Scientific Research in Computer Science, Engineering and Information Technology. 2024.

doi:10.32628/CSEIT24106183.

[21]. Ghobakhloo M, Iranmanesh M, Foroughi B, Tseng ML. Industry 4.0 digital transformation and opportunities for supply chain resilience: a comprehensive review and a strategic roadmap. Production Planning & Control. 2023 [volume and page range to be added]. doi:10.1080/09537287.2023.2252376.

[22]. Iftikhar A, Ali I, Arslan A, Tarba SY. Digital innovation, data analytics, and supply chain resiliency: a bibliometric-based systematic literature review. Annals of Operations Research. 2022 [volume and page range to be added]. doi:10.1007/s10479-022-04765-6.

[23]. Ivanov D, Dolgui A. A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. Production Planning & Control. 2021;32:775–788. doi:10.1080/09537287.2020.1768450.

[24]. Frederico GF, Garza-Reyes JA, Anosike A, Kumar V. Supply Chain 4.0: concepts, maturity and research agenda. Supply Chain Management. 2020;25:262–282. doi:10.1108/SCM-09-2018-0339.

[25]. Belhadi A, Kamble S, Jabbour CJC, Gunasekaran A, Ndubisi NO, Venkatesh M. Manufacturing and service supply chain resilience to the COVID-19 outbreak: lessons learned from the automobile and airline industries. Technological Forecasting and Social Change. 2021;163:120447. doi:10.1016/j.techfore.2020.120447.

[26]. Azzamouri A, Fenies P, Fontane F, Giard V [author list and ordering to be confirmed — Appendix A, item F10]. Impact of the continuous and periodic assessment of a buffer replenishment on the DDMRP method. International Journal of Production Research. 2023;61(13):4489–4508 [volume, issue and pages to be verified]. doi:10.1080/00207543.2022.2109219.