



WHITE PAPER

MULTI-AGENT ORCHESTRATION:

ENABLING SYSTEM-WIDE INTELLIGENT FLOW IN
THE MODERN WAREHOUSE EXECUTION SYSTEM



INTRODUCTION

Distribution and fulfillment operations today manage warehouse ecosystems of unprecedented complexity. Autonomous mobile robots (AMRs), shuttle systems, high-speed sortation, and dense storage have all become standard components of modern infrastructure — and each of them, in isolation, performs well.

Yet many operations still fall short of the throughput levels all that equipment should deliver. The constraint isn't speed or capacity.

It's coordination.

When people, robots, conveyors, and software each operate according to their own logic — without a shared, real-time view of the whole system — the result is an operation that reacts to problems rather than preventing them. Waves release too early and flood downstream zones. Labor plans built on yesterday's averages don't survive contact with today's order profile. Supervisors spend their shifts firefighting congestion that a better-coordinated system would never have allowed to form.

Closing that gap requires a specific capability inside the software layer that directs warehouse execution: multi-agent orchestration (MAO). Gartner, which now tracks multi-agent orchestration as a dedicated technology category, defines MAO platforms as:

“...intelligent middleware that integrates and orchestrates work between various business applications, heterogeneous fleets of operational robots and other automated agents (such as doors or elevators). These solutions take on tasks from connected systems or human instruction, then orchestrate and assign work, and monitor and coordinate the activities of diverse fleets of robots and other systems.”

Embedded inside a modern warehouse execution system (WES), multi-agent orchestration transforms execution from a reactive, equipment-by-equipment process into a coordinated, predictive one — continuously balancing labor, automation, and order flow against real throughput targets rather than static rules.





This white paper examines why today's warehouse operations need this capability, how multi-agent orchestration works inside a WES, and what the shift from reactive automation to predictive, system-wide execution looks like in practice.

THE HIDDEN CONSTRAINT IN MODERN FULFILLMENT

Over the past decade, the individual building blocks of warehouse automation have gotten dramatically better. Shuttle systems move product faster. Robotics platforms handle a wider range of tasks. Storage density has increased. Sortation and conveyance systems run at higher and higher rates. Taken in isolation, most of this equipment performs well against its own specifications.


And yet, many distribution operations still cannot reliably hit throughput targets when it matters most. A flash sale doubles inbound order volume mid-shift. An e-commerce spike hits at 3 p.m., two hours before next-day air cutoff. A single zone backs up and the ripple moves downstream faster than any supervisor can respond.

These aren't edge cases. They're the moments that define whether an operation actually performs. If every individual system is performing as designed, the obvious next question is: **Why isn't the operation keeping pace?** The answer, more often than not, is that nothing is actively, dynamically, and intelligently rebalancing the whole system as conditions change.

Why Equipment Isn't The Bottleneck Anymore — Flow Is

The answer, increasingly, is flow — the coordinated movement of work through zones, handoffs, and resources, rather than the speed or utilization of any single asset. An





operation can have excellent automation and still underperform if work is released too early and floods downstream zones, or too late and starves them. Labor can be fully staffed and still fall behind if it isn't synchronized with what automation is producing upstream.

This is a fundamentally different problem than equipment throughput. A conveyor that is technically capable of running at full rate does no good if downstream zones can't absorb the volume. A labor team that is fully staffed for an "average" day struggles when actual order profiles vary hour to hour. The system as a whole is only as fast as its worst-coordinated handoff.

The Cost of Siloed Control

Part of the reason this problem persists is architectural. In many operations, the warehouse management system (WMS) or enterprise resource planning (ERP) system owns order and inventory data, a warehouse control system (WCS) or WES layer owns automation routing, and labor management runs as a separate process — often built on static plans, spreadsheets, or supervisor judgment. Each layer optimizes its own slice of the operation without a shared, real-time view of the whole.

The result is a set of decisions that are individually reasonable but collectively uncoordinated. A conveyor zone releases cartons as soon as it is able to. A picking zone works through its own queue. Labor reacts only once congestion is already visible on a dashboard or upon receipt of a supervisor's radio call.

None of these decisions is wrong on its own. But without a layer coordinating them together, the operation ends up managing symptoms. Among them? Overtime spikes, congestion at consolidation and shipping, underused automation, and supervisors spending their day firefighting instead of planning.

Closing this gap requires more than faster equipment. It requires a way to coordinate everything at once, continuously, as conditions change. That is the role multi-agent orchestration plays inside a modern WES.

WHAT IS MULTI-AGENT ORCHESTRATION?

Multi-agent orchestration platforms have become significant enough as a category that Gartner now tracks them directly, publishing dedicated research on the space.

As detailed in this white paper's introduction, Gartner defines a multi-agent orchestration platform as software that acts as intelligent middleware. An MAO platform integrates and coordinates work between business applications, heterogeneous fleets of operational robots, and other automated agents, such as like doors, elevators, or conveyors. Further the MAO assigns and monitors that work in near real time.



Notably, Gartner's own naming for this category evolved over time. The firm originally referred to these platforms as "multirobot integration/orchestration" tools before recognizing that robots need to communicate not just with each other, but with the broader physical environment around them.

Four Capabilities, One Platform

Gartner's research frames multi-agent orchestration around four interlocking capabilities:


- **Integration:** Connecting an increasingly diverse mix of robots and physical agents to a single platform and linking that platform upstream to core business systems, such as a WMS.
- **Orchestration:** Assigning work to the right resource and coordinating the interplay between systems, such as a robot triggering a conveyor, door, or elevator, rather than each system acting independently.
- **Optimization:** Applying algorithms and artificial intelligence (AI) to build the best possible work plan across a mixed fleet, accounting for real constraints and objectives instead of relying on static, one-size-fits-all rules.
- **Observation:** Continuously monitoring the fleet, work assignments, and projected needs, both to understand what has already happened and to anticipate what's coming next. This capability is something individual robot vendors' own analytics tools typically can't provide across a mixed fleet of automated assets.



Why This Matters Now

The urgency behind this category is not theoretical. According to Gartner's 2025 Supply Chain Technology survey research, the large majority of supply chain leaders have already deployed, or plan to deploy, their first intralogistics robot within the next two years, and most of those who have already deployed one plan to expand further.

Critically, most of that expansion is expected to involve multiple robot types from multiple vendors. This is precisely the condition that makes orchestration a necessity rather than a nice-to-have. Gartner also estimates roughly 100 intralogistics robotics vendors now compete across more than half a dozen categories and dozens of use cases, and projects



that by 2028 the large majority of warehouses and distribution centers will have deployed some form of robotics or automation.

As fleets diversify, simple point-to-point integration between systems stops scaling. Multi-agent orchestration is the layer built to keep pace by connecting a growing, heterogeneous set of agents and coordinating their work continuously as conditions change.

From Concept To Warehouse Floor

Inside a modern WES, multi-agent orchestration is the connective layer between the order and inventory data managed by a WMS and the physical systems — human and automated — that a WES directs on the floor.

Rather than a single rules engine issuing static instructions, the system is organized around specialized software agents representing flow zones, resources, buffers, and operational objectives. Each of these agents monitors its own conditions while continuously collaborating with the others toward system-wide goals, deciding in real time when and where work should be released, how priorities should shift as conditions change, and which resources — human or automated — should be assigned or reassigned to keep the whole system in balance.

That is a fundamentally different posture than the static, fixed-priority logic most legacy execution systems still rely on. Further, it sets up the outcome operations are ultimately after: not just automation, but predictive throughput. That is, the ability to see a constraint coming and act on it before it degrades performance, rather than reacting to it after the fact.



FROM LOCAL DECISIONS TO SYSTEM-WIDE INTELLIGENT FLOW

Traditional execution logic asks a narrow question, over and over: Is this resource busy right now?

A flow-intelligent system asks a different, more useful question:

Is work moving through the operation at the rate required to meet today's throughput, service, and peak-volume goals?

That shift — from monitoring utilization to monitoring flow against a target — is what separates reactive automation from intelligent flow.

Dynamic Optimization Replaces Wave-Based Execution

Much of today's execution logic is still organized around waves. Work is released in batches, processed, and then the next batch is released. This creates an artificial rhythm of peaks and valleys. A flood of work hits downstream zones immediately after release, then those same zones sit underutilized as the wave winds down and the next one is prepared.

A flow-intelligent WES with multi-agent orchestration functionality replaces that batch rhythm with continuous metering. Work is released, priorities are reshuffled, and resources are reassigned in real time as conditions evolve, rather than in large, disruptive batches.

Just as important, this happens proactively. The system doesn't wait for a dashboard to show a queue building up before acting. Rather, it identifies the conditions that predict congestion and redirects work before that congestion actually forms.

Continuous Monitoring and Predictive Insight

Delivering on that requires constant comparison between modeled and actual performance across every active flow zone, with throughput rates, queue depths, and execution timing all tracked in real time. From there, the system can continuously recalculate projected completion times and required processing rates against service-level and peak-day goals, flagging where the operation is



on pace and where it is not.

This is the key difference between a reactive operation and a predictive one. Reactive systems tell a supervisor that a zone is already backed up. A predictive, flow-intelligent system tells them a zone is on track to back up in the next hour, while there is still time for a corrective action to matter.

Coordinating a Unified Mix of People and Automation

None of this works if labor, conveyors, robotics, and fixed automation are still managed as separate silos. A flow-intelligent WES treats them as a single, unified resource network. One system routes tasks to whichever combination of human and automated resources makes sense at that moment, with the explicit goals of minimizing idle labor, preventing congestion at handoff points, and avoiding both overload and starvation across zones. Integration with vision systems for real-time resource detection also plays a role here, supporting both


coordination and safety as people and machines share the same physical space.

WHY LABOR HAS TO BE PART OF THE ORCHESTRATION MODEL

Warehouse execution system technology has historically evolved to orchestrate equipment, with routing logic, sequencing, and wave releases built with automation in mind. Conversely, labor planning grew up almost entirely outside of that layer, typically handled through a standalone labor management system, static reporting, or a supervisor's own judgment on the floor. The result is

that no single system has traditionally coordinated demand data from the WMS, automation capacity from the WES, and labor availability together, in real time.





Why Labor Optimization Belongs Inside the WES

That gap matters because the WES already sits in the best position in the technology stack to fix it. It has visibility into inbound orders, wave plans, batching logic, automation capacity, and downstream constraints. By leveraging all this data, it can generate a real-time labor recommendation, rather than rely on a static staffing plan built the week before.

That reframes the operational question from “where do we need people right now?” to one that’s infinitely useful. “Where will we need capacity next, and how should people and automation be balanced to maintain flow?”

Moving Beyond Static Staffing Models

Most staffing models today are still built around averages, such as average picks per hour, average order size, average volume per shift. Real operations don’t run on averages, however. Order profiles, inventory complexity, individual performance, and automation speed all vary constantly, often hour to hour.

A flow-intelligent WES continuously compares expected to actual throughput across both manual and automated zones, refining labor recommendations as real conditions unfold rather than relying on a plan built on last month’s average day.




Smoothing Flow at Handoff Points


Handoff points — particularly consolidation and shipping — are where labor and automation misalignment shows up first and hurts the most. When automation outpaces labor, or the reverse, congestion builds quickly and unplanned overtime follows. Continuous metering of work release, paired with real-time labor recommendations, keeps release timing, resource availability, and staffing in sync at exactly the points where operations tend to break down.

and unplanned overtime follows. Continuous metering of work release, paired with real-time labor recommendations, keeps release timing, resource availability, and staffing in sync at exactly the points where operations tend to break down.

Turning Labor Constraints into a Coordinated Advantage

Put together, this capability lets an operation answer a different set of questions than it could before. Among them, where is capacity misaligned with demand right now? Which areas are at risk of a bottleneck or of sitting idle? How should





**WHAT TO LOOK FOR IN A
FLOW-INTELLIGENT WES:**

- **Unified resources** —
Connect labor and automation.
- **Continuous execution** —
Manage work in real time.
- **Predictive analytics** —
Forecast needs before they arise.
- **Real-time labor guidance** —
Optimize staffing dynamically.
- **System compatibility** —
Extend existing WMS/WCS.
- **Multi-vendor integration** —
Support mixed automation fleets.
- **Phased deployment** —
Minimize disruption and risk.

CONCLUSION

Automation alone no longer differentiates one operation from another. Indeed, nearly everyone has access to fast conveyors, capable robotics, and dense storage. What separates high-performing operations now is coordinated flow: the ability to see the whole system at once, predict where it’s headed, and continuously rebalance people and automation to stay ahead of constraints instead of reacting to them. That is what intelligent flow delivers, and multi-agent orchestration is the architecture that makes it possible.

This capability is built directly into **DATUM, Designed Conveyor Systems’ (DCS) proprietary warehouse execution system**. DATUM coordinates labor, robotics, and fixed automation as a single execution model; delivers predictive throughput through continuous monitoring and dynamic optimization; and extends real-time labor recommendations into day-to-day execution, with interactive labor planning on its ongoing roadmap. It is backed by DCS’s recognized position across consulting, systems integration, and orchestration software — including a named citation as a Representative Provider in Gartner’s own research on this category.

And because DATUM is designed for phased, low-disruption deployment, the path to getting there doesn’t have to be an all-or-nothing bet. Operations can adopt orchestration incrementally, layering capability onto existing systems without a costly cutover or extended downtime.

For operations weighing what comes next after their current round of automation investment, the conversation worth having isn’t about adding more equipment. It’s about whether that equipment, and the people working alongside it, are actually being orchestrated as one system. To see how DATUM applies to your operation, visit designedconveyor.com/software or connect with the DCS team for a conversation.





ABOUT DCS

Headquartered in Franklin, Tenn., Designed Conveyor Systems (DCS) has almost 45 years of experience serving major clients in multiple industries by providing material handling, full-scale warehouse operations, and conveyor design solutions that are custom crafted for their needs. DCS does not sell ready-made conveyor systems but builds relationships that empower collaboration to craft custom warehouse designs together. DCS utilizes consulting, engineering design, project management, installation services, and client support to ensure our customers can keep their promises to deliver on time.

DISCLAIMER

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.



4020 Aspen Grove Drive
Suite 200
Franklin, Tennessee 37067