


The Real Problem: Decision Overload

Building Decision Support System for Dark Store Managers



1. Context & Problem Definition

India's quick commerce market is worth roughly **\$5 billion in GMV as of 2025**, with platforms delivering groceries and essentials within 10-15 minutes across 100+ cities. As of late 2025, India has approximately **2,525 dark stores**, compact warehouses of **3,000-8,000 sq ft** that stock **6,000 to 25,000 SKUs** and exist only to pick, pack, and dispatch online orders. Blinkit, the market leader with roughly **45% share**, operates over **2,000 of these stores**, processes about **1.65-1.75 million orders daily**, and is on track to reach **3,000 stores by March 2027**.

A single dark store handles **500-700 orders on a normal day** and **over 1,000 during festivals**. Each order is expected to go from screen to rider in about **2 minutes**. The system works. At average load, it runs smoothly.

1.1. What the Numbers Say vs. What Is Actually Broken

The way most people describe this problem is straightforward: dark stores get too many orders during peak hours and don't have enough staff or space to handle them. That framing is not wrong. But it leads you to the wrong solution.


If the problem is capacity, then the answer is more pickers, bigger stores, or more riders. Those are supply-chain decisions, and they take months, cost money, and don't scale across 2,000+ stores in any practical way. More importantly, that framing misses what actually breaks first.

The store doesn't fail because it physically cannot handle the orders. It fails because **the one person responsible for keeping everything running, the manager, loses the ability to make good decisions under pressure**. They face 40+ unstructured decisions per hour during peak hours. Which orders to prioritize. Whether to pull someone from restocking to picking. When to call backup riders. Whether to skip staff breaks. None of these decisions are supported by data, recommendations, or early warnings. Every call is gut-based, made under noise and time pressure, with no way to see what's about to go wrong.

And when the pressure crosses a threshold, the manager doesn't just make worse decisions. They stop making decisions altogether. They drop their oversight role and start picking orders themselves. That is the single most damaging failure mode: the person who should be coordinating the system is now trapped inside it, unable to see anything else that's breaking.

This is a decision-support problem, not a capacity problem. That distinction changes what you build

1.2. Why This Problem Will Not Fix Itself



Illustrative hourly order distribution for a store processing ~650 orders/day

Dark stores experience two daily peaks: **morning (6-11 AM)** and **evening (5-11 PM)**. During evening peak, order volumes spike **30-50%** over the normal hourly rate. But the store's space, staff count, and tools stay exactly the same. Demand surges. Capacity doesn't. And the manager, with no forward-looking signals and no prioritization framework, is left to figure it out in real time.

This problem is not going to balance out on its own. It has a self-reinforcing loop.

Peak-hour chaos leads to picking errors and delivery delays. Delays lead to customer complaints and refunds. Complaints eat into the manager's bandwidth for the next shift. Staff who burn out during repeated bad peaks quit, which means new hires who are slower and make more mistakes. The next peak hits a weaker team, and the cycle repeats.

1.3. Business & Stakeholder Impact

Now layer on scale. Blinkit is expanding from 2,000 stores today to 3,000 by early 2027. Every new store inherits the same problem: the same fixed space, the same limited tooling, the same single manager trying to hold things together with WhatsApp and gut instinct. At 500 stores this was a local problem. At 3,000 stores with no decision-support layer, it becomes a systemic one. The math is simple: **if peak-hour instability costs X per store per day, it costs 3,000X across the network.**

Quick commerce runs on thin margins. Every delayed order, every refund, every mis-pick adds direct cost. Even a **5% drop in on-time delivery** in high-demand zones affects hundreds of orders daily across multiple stores.

The customer-facing impact is visible. On Voxya, a consumer complaint platform, Blinkit's recorded customer satisfaction score sits at **71.15%**, with only **74 out of 104 tracked complaints resolved**. That is a limited sample, but it signals a pattern: when operations break during peak, it shows up as unresolved complaints, delayed refunds, and trust erosion. On the

ground, a single high-demand store can have **roughly 60 riders waiting outside at any given time** during peak, a visible bottleneck between store output and delivery capacity.

1.4. Scope Definition: In and Out

Repeated peak-hour instability hits the business in six places at once:

Higher cost per order	Increased refund payouts	Rider idle time and clustering
Manager burnout and attrition	Picking errors and wrong deliveries	Erosion of customer trust and ratings

Top-line revenue can keep climbing, but if stores are consistently strained every evening from 6–10 PM, the unit economics that determine whether this model scales take a quiet, compounding hit.

The case study is clear on scope: we’re not redesigning the supply chain, fixing demand forecasting, or increasing store capacity. We’re focused on improving the store manager’s decision clarity during peak hours. Here’s where we draw the line:

Area	Scope	Reason
Real-time store health visibility for managers	In	Core to the decision gap identified
Peak-hour prioritized action recommendations	In	Directly reduces decision overload
Automated backup staffing triggers	In	Removes the biggest time sink (WhatsApp calls)
Rider allocation and pre-positioning	In	Addresses P0 pain point (rider shortage)
Expanding physical store size	Out	Infrastructure decision, not product scope
Demand forecasting models	Out	Separate data science workstream
Reducing growth or capping order volume	Out	Business constraint, cannot limit demand
Drastically changing SKU count	Out	Operational constraint per case study
Rider pay structure or gig worker policy	Out	HR and policy domain, not product

1.5. Problem Statement in Plain Language

During peak hours, Blinkit’s dark store managers face decision overload with no real-time support. They cannot see stress building before it becomes a crisis. They cannot tell which problem to fix first when five things break at once. And they have no system-level way to act, whether that means summoning backup staff, flagging a late rider, or escalating to a cluster

manager. The result: they react instead of anticipate, firefight instead of coordinate, and the store runs in a degraded state for 3-4 hours every evening.

2. Market Context

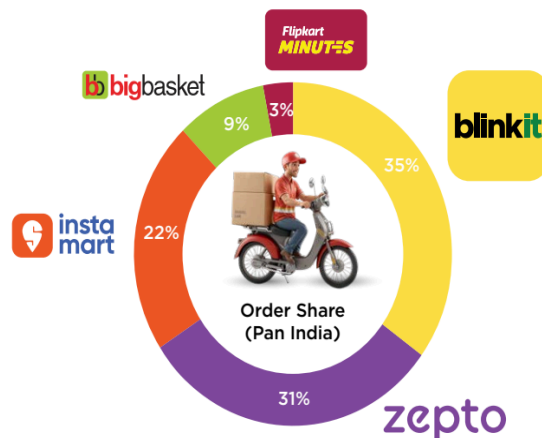
India doesn't wait anymore. A consumer who craves instant gratification, groceries in 10 minutes, snacks before the craving fades, has quietly rewritten the rules of retail. Quick commerce isn't a convenience upgrade; it's a behavioral shift at scale. And the infrastructure racing to serve it (dark stores, fulfillment networks, last-mile fleets) is being built faster than almost any retail category in history.

2.1 Market Size & Platform Position

The numbers are hard to ignore. India's quick commerce market is on track to hit **\$57 billion by 2030**, growing at a CAGR north of 40% one of the steepest growth curves in global retail. What started as a metro-Delhi-Bengaluru phenomenon is now pushing into Tier II and III cities, with dark stores projected to **triple from ~2,500 to 7,500** in five years.

Three players dominate the race: **Blinkit** (approx 46% share, 1,800+ dark stores), **Zepto** (29%), and **Swiggy Instamart** (25%) each spending aggressively to out-store and out-speed the other. Quick commerce already accounts for **67% of all online grocery orders** in India, and its footprint — 13 million sq ft of dark store space and growing — is expanding toward residential zones at a pace that makes traditional retail look slow.

But scale creates its own pressure. When a single dark store processes **1,200–2,000 orders a day** and evening peaks spike **30–50% above normal**, the operational margin for error shrinks dramatically. The market is scaling fast. The question is whether the operations running inside it can keep up.



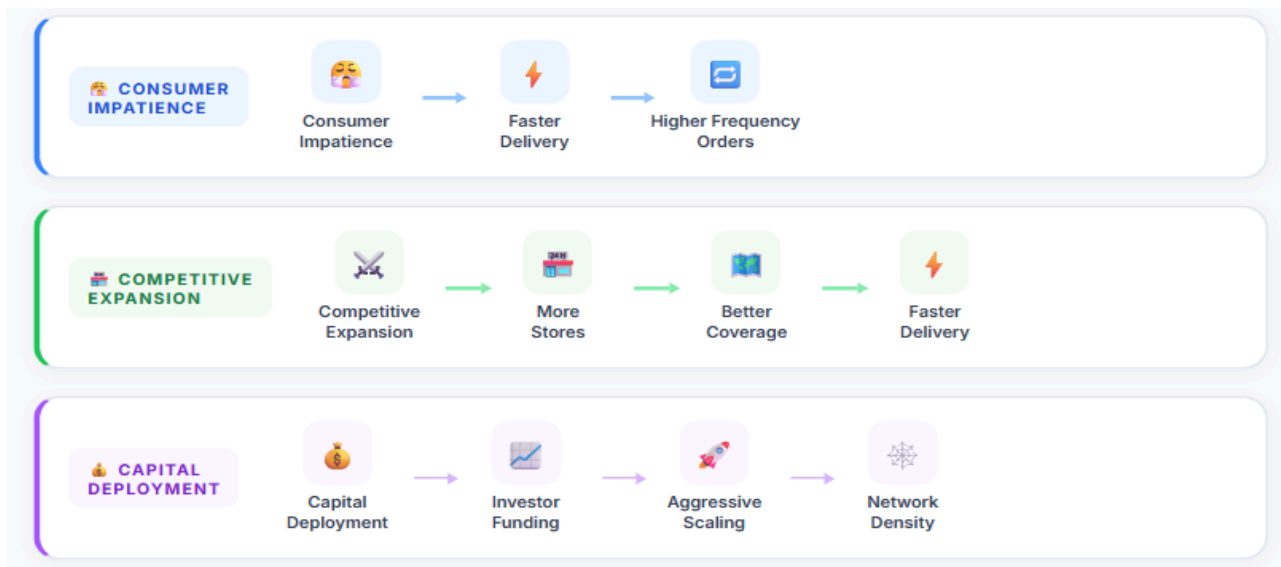
Quick commerce market share based on order volume (daily orders/ day) ¹³
Source: Moneycontrol, Business Standard

2.2 Current Trends

- **Category Diversification:** Shift to high-margin non-grocery (pharma, electronics, private labels) for better AOV; Tier-2/3 city penetration via 80+ expansions.
- **Tech/Operations:** Manual picking dominates but automation (AI picking, IoT) emerging for efficiency
- **Challenges:** High capex on setups/logistics, demand fluctuations, environmental impact from packaging/fleets, persistent losses despite revenue jumps (e.g., Blinkit 155% YoY)

2.3 Quick-Commerce Growth Loops:

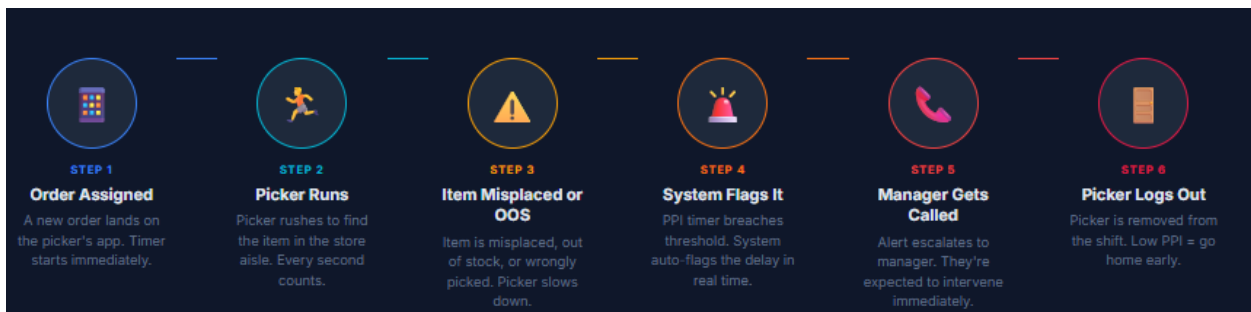
3 forces driving dark store expansion flywheel.



Insight: All three loops converge, and peak-hour pressure spikes. As order volume grows, the store footprint expands, and targets get tighter, managers hit overload during peak-exactly the failure mode decision support is meant to address.

2.4 PPI (Per Picking Item)

A single delay triggers a full system chain reaction - from picker to manager to lost shift.

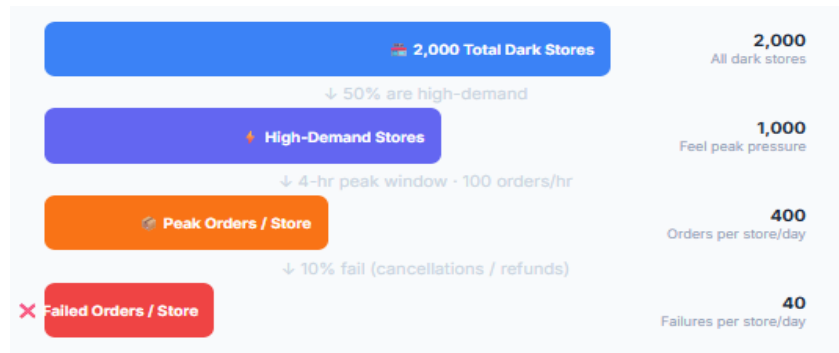


Insight: The problem isn't the picker, it's the lack of early warning. By the time a manager gets the call, the damage is done. Decision support that detects stress before PPI breaks down can stop this chain before it starts.

2.5 Who Loses if This Is Not Solved?

Market sizing: Peak-hour decision support for dark store managers. Let's quantify the peak-hour loss that could be addressable via manager decision support

Key assumptions (example scenario):



Peak-hour failures are not edge cases — they are a systemic, daily cost centre hiding in plain sight.

Level	Calculation	Annual Impact
Per Store / Day	40 failed orders × ₹200 avg refund	₹8,000 / day
Network / Day (2,000 stores)	2,000 × ₹8,000	₹1.6 Crore / day
Annualized (refunds only)	₹1.6 Cr × 365	~₹585 Crore / year

This is direct refunds only. Indirect costs (rider re-dispatch, customer churn, lower ratings, increased CAC, manager attrition) multiply the true cost 2–3x.

The Recovery Window

If Peak Assist Reduces Failures By...	Orders Saved / Day	Annual Savings
20% (conservative)	16,000 orders across network	~₹117 Crore
30% (target)	24,000 orders across network	~₹175 Crore

→ A 20–30% reduction in peak-hour failures recovers ₹100–175 Cr annually - far exceeding any investment in the tool. The loss pool is real, recurring, and waiting to be captured.

3. Secondary Research

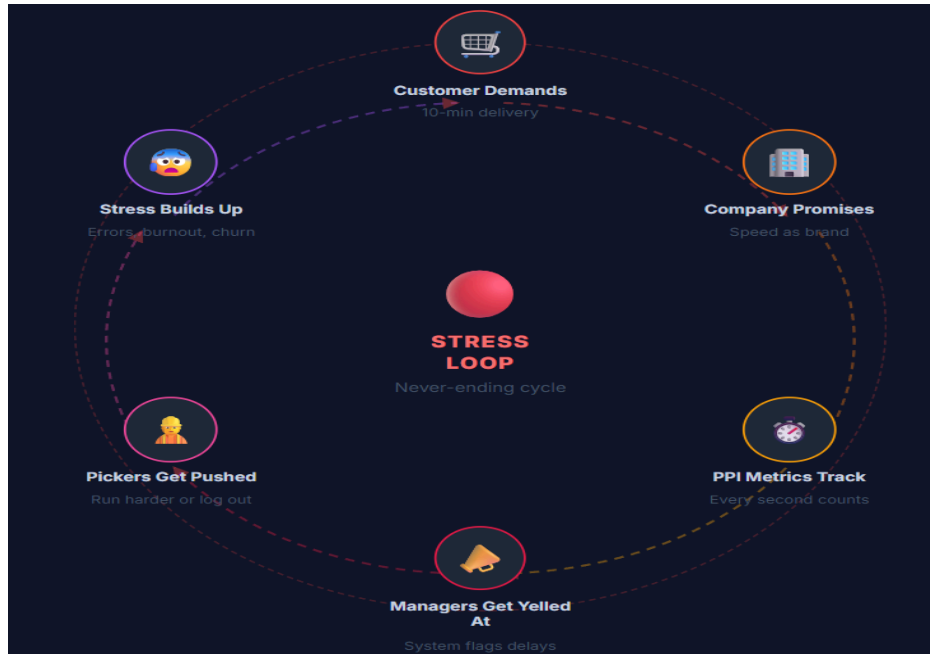
3.1 The Vicious Cycle Inside a Dark Store

Description: This self-reinforcing vicious cycle structurally disadvantages dark stores, turning peak hours into a predictable breakdown pattern rather than scalable throughput

Blinkit promises 10-minute delivery. Customers flood in. Orders pile up → WMS lags, dispatch freezes → Pickers burn out, then quit → Manager jumps in to pick, gets yelled at by angry riders → PPI metrics flash red. Stress skyrockets → More quit → leading to endless chaos.

Reinforcing elements:

- Stress builds churn among pickers/managers, worsening shortages
- Rider queues outside block flow, feeding back into delays and escalations.
- The loop closes as the never ending cycle of firefighting prevents fixes, sustaining chaos.



3.2 In-Store Challenges of Blinkit Dark Store:

- High Attrition and Operational Instability
- Manager Burnout and Overload
- Cash-Handling and Security Risks
- Rider Related Friction
- Inadequate Training and Ramp-Up
- Poor Infrastructure and Safety Conditions

3.3 Operational Bottlenecks During Peak:

Inside the Store: Where Throughput Breaks

Bottleneck	What Happens	Peak-Hour Impact	Data / Pattern
Picker Congestion	Too many pickers (including managers) in narrow aisles; paths collide, pickers stop and reroute.	Orders slow; riders start queuing while the store gets swamped.	Congestion adds 9.6%–13.1% to total picking time, depending on order size.
SKU Clustering & Inventory Stress	High-demand SKUs clustered in one zone → hotspots, crowding, more mis-picks, faster stockouts.	Partial orders, rework, and “missing item” escalations spike during peak throughput.	Physical-to-digital gaps (damage, shrink, FIFO/FEFO breaks) trigger partial orders and rework when it matters most.
Rider Wait Queues	Order readiness and rider arrival get out of sync; riders cluster outside blocking entry and operations.	Discipline issues (wrong drops, complaints) create loop-back work for already-stretched managers.	In high-density areas, 30%+ rider shortages create persistent, long-duration queues.
Dispatch Lag (“Packed but Stuck”)	Orders are ready but don’t leave due to QR glitches, manual WMS mapping, COD mismatches, reconciliation.	Extra exception-handling at handoff; customer escalations rise as orders miss SLA.	Exception-handling time at handoff grows under peak pressure and manual cash flows.
Order Batching Conflicts	Batching speeds up picking but crowds packing stations, increases sorting mistakes, and raises error rates.	Mid-shift HQ changes (promos, planograms, expiry checks) add unplanned work in a tight cycle.	Without early stress signals, store destabilizes before manager can intervene; they firefight instead of anticipate.

3.4 Common Tools & Metrics:

Tool / System	What It Tracks	Operational Purpose	Insight
Real-Time Performance Dashboard	Pick Rate (items/hour), Pick Time/order, Packing Time	Measures speed and throughput at individual + store level	Enables instant identification of slowdowns and performance gaps
Accuracy Monitoring System	Picking Errors, QC Failures, Wrong SKU scans	Tracks order accuracy before dispatch	Prevents customer churn due to fulfillment mistakes
Order Flow	Orders per minute, Peak congestion windows	Visualizes live demand spikes	Supports dynamic manpower reallocation in real time

Productivity Benchmarking	Individual vs Store Average, Best Performer Metrics	Compares associate performance	Encourages competitive productivity culture (can increase pressure)
Tracking System	Out-of-stock flags, Substitutions, Delay reasons	Identifies operational bottlenecks	Highlights inventory or layout inefficiencies
SLA Monitoring (Delivery Timer)	Order-to-dispatch time, Rider waiting time	Tracks end-to-end speed	Aligns store output with last-mile performance goals
Forecast vs Actual Dashboard	Predicted orders vs Real orders	Measures forecasting accuracy	Improves future staffing precision & cost control

3.6 Key Insights from Secondary Research

Store Manager (what they experience)	Root cause
Aisles feel crowded; pickers are stuck in each other's way	High order density + poor aisle-level congestion handling
Riders are queuing, sometimes for long stretches	Dispatch lag and rider-order mismatches due to imperfect rider-side algorithms and incomplete visibility into rider availability
Large orders get stuck at packing; small orders age in the queue	Suboptimal order batching logic that creates congestion at packing, not at picking, and no dynamic batching rules tuned for peak
Manager feels like everything is high-priority	No intermediate signal layer between "normal" and "SLA-breach"; only binary alerts ("SLA at risk") arrive too late
Manager jumps into picking or manually reorders	Absence of predictive signals and prescriptive actions (e.g., add 2 pickers or hold these batches) forces improvisation

3.7 Warehouse Operations End to End



Warehouse operations cover receiving, storing, and shipping goods, supported by systems like WMS and barcode scanning for accuracy and efficiency. They include layout and safety design, quality control, and staff training to maintain smooth, secure workflows.

Key metrics such as throughput, accuracy, and efficiency are tracked to optimize performance and reduce errors. Overall, warehouse operations ensure inventory is managed reliably from receipt to delivery

4. Primary Research

We interacted with a few dark store managers to understand more about them. We also spent time observing, shadowing, and speaking with the people who live through peak hours every day.

4.1 How We Did The Research

We visited four dark stores across two cities, spending roughly 45 minutes to an hour at each location. We observed operations during both regular hours and the lead-up to peak hours.

What	Details
Stores Visited	3 Blinkit stores: Gachibowli, Khajaguda (Hyderabad) and Mukai Chowk(Pune) 1 Instamart store: Chandanagar, Hyderabad (for comparison)
Who We Spoke To	3 Blinkit Store Managers (primary), 1 Instamart Store Manager, Riders who were at the store: casual conversations and observation
Duration Per Visit	~45 min – 1 hour per store, including breaks and observation time
What We Captured	Voice recordings from interviews, Photos of the store floor, packing area, parking zone, and the live dashboard
Store Scale	High-demand stores up to ~2,000 orders/day on normal days, 3,000+ orders/day during peak

The goal was not just to ask questions; it was to watch how the store actually operates. A lot of what we learned came from things we observed that were never directly said.



Link to more photos and recordings - [Primary research](#)

4.2 What does a store manager even do?

The store manager is not just someone who watches over the store. They are the operational centre, which is the single point of contact for everything that happens from the moment an order is placed to the moment a rider leaves the gate with delivery.

On any given day, a store manager is juggling between

People	Operations	Systems
Pickers coming to report missing items (PNA)	Daily inventory audits, scrapping expired stock	Monitoring live orders on the dashboard
Riders arriving, depositing cash, or needing manual QR help	Managing shift overlaps and flex staff during peaks	Tracking PPI (Picking Per Item speed) in real time
Taking calls from delivery partners who are delayed	Ensuring packed baskets don't exceed the max-basket cap	Verifying packing via camera footage for dispute cases

The three KPIs that are held accountable for every single day:

KPI	What it means	Why it matters
Inventory Sanity	Does physical stock match what the system shows, and quality checks	Wrong data = wrong orders = refunds
PPI (Picking Per Item)	How fast are pickers collecting items per order? Directly impacts delivery time.	Slow PPI = riders waiting = delays
Filling Rate	% of ordered items are actually delivered without being marked as unavailable (PNA)	Low fill rate = unhappy customers

The manager's world is one of constant interruption, real-time decisions, and accountability for things that are often outside their direct control.

4.3 What Peak Hour Actually Looks Like on the Ground

Peak hours at a dark store are typically 6–10 AM in the morning and 6 PM onwards in the evening, apart from weekends, holidays, and special occasions. On the surface, it just looks like more orders are coming in. But inside the store, it feels like a completely different operation.

Order volume on a regular day: ~2,000 orders. During peak: 3,000+ orders.

That jump of 1,000 extra orders doesn't happen slowly. It hits fast, sometimes within 20–30 minutes, during events like IPL evenings, rain, or festive sales.

What we saw during our visits:

- Pickers move fast through aisles, sometimes missing shelves and coming back to the manager to report PNA (Product Not Available).
- Riders clustering at the store entrance, waiting for packed baskets
- The manager physically moves between the dashboard, the packing area, and the rider bay constantly.
- Staff are being pulled off their other regular tasks and redirected on the spot verbally.
- COD cash is being handed over by riders mid-rush, requiring the manager to stop and reconcile it.

“It's complete chaos, everyone is running on their toes.”

— Blinkit Store Manager, Hyderabad

One thing that stood out: during peak, the manager often stops being a manager and sometimes starts doing the work himself. He steps in to pick items, to keep PPI from dropping. This is a sign of how stretched the operation becomes where the manager responsible for making decisions ends up doing execution instead.

4.4 How Decisions Get Made Today

When peak hits, how does the manager decide what to do first?

“The short answer: gut feel and experience.”

There is no structured prioritization. No system that tells the manager, “this needs your attention first.” Decisions usually happen in the order that things come up first: whoever walks up to him, whatever notification appears, whatever he notices while walking the floor. Over time, experienced managers develop an instinct for what matters. But that instinct is built over months of making mistakes and getting corrected.


How learning happens:

Every day, managers have a debrief with their area manager. Things that went wrong during (delays, complaints, PNA spikes, rider wait times) are reviewed and discussed. Over time, this is how a manager learns to prioritize. But it is a slow process, and during the actual moment of peak, the manager is largely on their own.

What about the existing in-house Blinkit app?

The existing Blinkit dashboard is actually quite capable. Managers praised it. It shows inventory levels, live orders, rider proximity, PPI, surge signals, and more. But during peak, something interesting happens: most of those features go unused.

It's not that managers don't know the features exist. It's that reaching any specific action requires multiple taps and navigation steps and when you're being pulled in five directions at once, that effort adds up. The dashboard is built for information. Peak hour demands action.



“The app shows everything. But when things are busy, you just go by what feels right. You learn over time what to look at.”

— Blinkit Store Manager, Hyderabad

4.5 Key Friction Points We Observed

Below are the few friction points that came up most consistently across our store visits. These are not just inconveniences, each one directly interrupts the manager's ability to think and act clearly during peak hours.

Friction Point	What We Observed	Direct Quote / Insight	Impact
Rider follow-up calls	The manager sees on the dashboard how long riders have been out. During peak, getting new riders instantly during a sudden spike is not possible so calling existing ones back faster is the main lever. So he checks for people returning back and tries to coordinate riders returning from delivery.	"This is the most frustrating part. I can see on the app how long they've been out but I still have to call them manually."	Breaks manager focus. Every call is a context switch during peak.
Verbal staff reallocation	During peak, the manager walks up to staff doing inventory checks or quality audits and tells them verbally to stop and start packing orders. Has to update the app for every person.	"I just go to each one and tell them to stop whatever they're doing and come help with orders."	Time lost assigning manually. No option to assign all at once.
Dashboard overload during peak	The dashboard has all the information needed but each action requires multiple taps to reach. During peak, most features are not used. The manager relies on what's visible at a glance or what someone tells him.	"Everything is there in the app. But when it's busy, you don't have time to go looking for things."	Decisions get delayed or missed. Reactive instead of proactive.

4.6 What the Instamart Visit Added

To check if these patterns are specific to Blinkit or a wider industry challenge, we also spoke to a store manager at an Instamart store in Indore. The core experience was very similar to manual coordination, staff management by gut feel, and peak hour pressure.

The key takeaway from the Instamart visit: the operational chaos during peak is not a Blinkit-specific problem. It's a structural challenge for anyone running dark store operations at scale. What differs is the approach and there is room to do better.

4.7 Key Insights from Research

These are patterns we observed consistently across stores.

#	Insight	What it really means
1	The manager is skilled, but the system makes them work harder than needed	Experienced managers make good decisions. The problem is that the system forces them to spend time finding the right information and taking manual steps instead of acting immediately. The bottleneck is not the manager's judgment, rather it's the effort between judgment and action.
2	Peak hour is a fundamentally different operating mode	The jump from 2,000 to 3,000+ orders in the same space with the same staff is not just 'more of the same.' It changes what the manager needs to see, what decisions matter, and how fast those decisions need to happen. The current app is not designed for this mode.
3	Information is available. Prioritization is not.	The dashboard shows everything. But during peak, 'everything' becomes noise. What the manager actually needs is for the most critical 2–3 things to be surfaced automatically, with a clear action they can take without navigating.
4	The most frustrating actions are also the most frequent	Calling riders, reallocating staff, and checking basket status, all these happen multiple times every peak hour. They are all done manually today. Streamlining even two of these would meaningfully reduce the manager's cognitive load.
5	This is an industry-wide gap, not a Blinkit-specific bug	Instamart faces the same challenge. The tools exist. The people are capable. What's missing is a focused operational layer that is designed specifically for peak not as an afterthought.

“You learn over time what is important. But sometimes in the moment, it's just whoever comes to you first.”

Blinkit Store Manager, Kondapur, Hyderabad

5. User Persona

5.1 Persona 1 - Store Manager

Attribute	Persona - Store Manager
Name & Age	Amit Sharma, age 31
Location	BTM Layout, Bengaluru
Experience	Q-Commerce (Blinkit) - 3 years, store manager - 2 years
Background	In 2021, Amit began working as a picker at a Blinkit location. Because he was reliable, swift, and the most knowledgeable about the store, he was elevated to manager. Though no one informed him that the position would require him to combine cash receipts on Sundays or be on call at midnight for a WMS issue, he is happy about the promotion.
Entry Behavior	Opens WMS first thing in the morning, checks WhatsApp groups for overnight issues, and scans pending orders and staff attendance before the shift briefing. During peak, he is switching between three screens every few minutes.
A Typical Day	Amit reaches the store at 9 AM. By 10 AM he has already handled a staff no-show, approved a refund, and replied to twelve WhatsApp messages from his cluster manager. By 6 PM the real chaos begins. Orders spike, two pickers slow down, a rider is late, and someone from customer ops is calling about a complaint. He handles all of it simultaneously, mostly from memory and instinct.
Primary Goal	Get through peak hours without SLA breaches. Keep his team motivated. Do not get called after 10 PM.
Frustrations	Too little time, too many systems. The platform seems to have been designed for regular business hours rather than the critical window of 6 to 10 PM. is held accountable for delays that began without his control.
Key Pain Points	No single view of store health. Alerts give no context or suggested action. Cannot predict when a spike is coming. Forced to pick orders when staff fall short. Cash reconciliation eats into post-shift time.
Comfortable	Comfortable with apps and smartphones. Uses WhatsApp heavily for coordination. Finds the current WMS interface cluttered and slow.
Success Metrics	SLA breach rate during peak. Number of manual interventions per shift. Time spent on non-managerial tasks. Shift completion without overtime.

5.2 Persona 2 - Picker

Attribute	Persona - The Picker
Name & Age	Sumeet Kumar, Age 22
Location	BTM Layout, Bengaluru
Experience	Q-Commerce (Blinkit) - 1-year picker experience at the store
Background	Sumeet came to Bangalore from Mysore looking for a stable income. The picker job suited him because the timings were fixed and the pay was weekly. He is quick on his feet, knows the store layout by heart, and can hit 18 to 20 picks per hour on a good day. On a bad day, when the store is chaotic and everyone is shouting, his numbers drop, and he does not always know why.
Entry Behavior	Clocks in, checks his assigned zone on the picker app, and starts filling orders immediately. Does not spend time on setup or configuration. He expects the app to tell him exactly where to go and what to pick. If it does not load fast, he moves on instinct.
A Typical Day	At 4 PM, Sumeet starts his shift as the evening rush is just getting started. He manages it with ease for the first hour. But at 6:30 in the evening, orders begin to flood in more quickly than he can handle, like a tsunami striking the warehouse. While the software flashes a red-hot order from another location, Rajan occasionally shouts across to prioritise one zone. He follows his instincts and chooses what feels right. Pauses? They are unpredictable. On bad days, he works for four hours straight without taking a break.
Primary Goal	Hit his pick rate target, so he gets the weekly performance bonus. Finish his shift on time.
Frustrations	The app sometimes shows wrong shelf locations for items. He has to ask colleagues or check physically, which wastes time. Nobody tells him when an SKU is out of stock until he walks to the shelf and finds it empty. It feels like the system is always one step behind where the store actually is.
Key Pain Points	Steps are spent everywhere as Sumeet makes his way to a shelf, only to discover that the app is completely out of date or incorrectly located. He doesn't know about products that are no longer there until he's there and staring at empty spaces. He has mixed feelings because the app flashes a whole other set of priorities while managers shout one. Breaks occur at any time, and there is no tracking to ensure fairness. And the worst? Mid-shift, he has no idea how he's doing; there are no real-time statistics to tell him if he's on track.
Comfortable	Uses a smartphone mainly for WhatsApp and YouTube. Comfortable with simple app interfaces but gets confused with too many tabs or options. Needs the work app to be fast, simple, and clear.

Success Metrics

Picks per hour were maintained during peak. Pick the accuracy rate. Break compliance. Number of times the wrong shelf location is reported per shift.

6. Pain Points

6.1 Manager Side Pain Points

#	Pain Point	Severity	Why Prioritized	Frequency
1	<p>With multitasking overload, the manager is forced to handle operations, customer escalation + audits simultaneously</p> <p>Problem: The manager is the single point of contact for floor operations, inbound customer escalations, and compliance audits, all happening concurrently during peak.</p>	Critical	<p>During peak hour (between 6 and 10 PM), the single biggest cause of decision quality collapse during peak affects all other pain points downstream</p> <p>Impact on the User: Poor decisions in all three domains are caused by attention fragmentation. Errors in prioritization are a result of high cognitive load.</p> <p>Root Cause: All tasks arrive as equal-priority interruptions without any organized routing; there is no triaging or severity-ranking system in place.</p> <p>Quick win: This is a quick-win pain point</p>	Every peak hour and every day
2	<p>The manager becomes a picker when staff fall short during peak</p> <p>Problem: The manager takes on the role of a picker when the number of pickers is insufficient during peak hours, abandoning all supervisory and decision-making duties.</p>	Critical	<p>Impact on Users: During the most stressful time (peak hour), the store loses its sole decision-maker. Chaos in operations compounds.</p> <p>Root Cause: No procedure or instrument to initiate backup staffing prior to the management being compelled to substitute; no early warning mechanism for understaffing.</p> <p>Quick win: This is a quick-win pain point</p>	3 to 4 times per week
3	<p>Peak-time rider shortage, manual calling, and lack of auto-allocation</p> <p>Problem: Rider supply falls short of the needed capacity when demand rises. To cover gaps, the manager manually contacts riders or aggregator contacts; this is a laborious and unreliable procedure.</p>	Critical	<p>Directly results in SLA violations; under a limited time, the manual procedure is slow and unreliable.</p> <p>User Impact: Orders accumulate at handoff; SLA violations increase; managers' focus is shifted from retail operations to logistics coordination.</p> <p>Root Cause: No automatic pull mechanism to reallocate or surge-request riders before a serious shortage; no predictive rider demand signal.</p> <p>Long-term: This is a long-term solution</p>	Daily during peak hour

4	<p>Picker productivity: maintaining pick rate while ensuring breaks</p> <p>Problem: Without a system, the manager must manually track who is on break, who has been working nonstop, and whether pick rates are being met.</p>	Medium	<p>Removes the sole decision-maker from the supervisory position at the most critical moment, increasing the risk of cascade failure.</p> <p>User Impact: When too many pickers are on break at once, pick rates decrease, or breaks are missed (compliance + burnout risk).</p> <p>Root Cause: No real-time pick rate per person or system-level visibility into picker work/break cycles.</p> <p>Quick win: This is a quick-win pain point</p>	3 to 4 times per week
5	<p>Manager burnout: 12x14 hr shifts, on-call 24/7</p> <p>Problem: With no framework for delegation, no decision support, and no recovery time, dark store managers work in a structurally unsustainable manner.</p>	High	<p>Reduced productivity and declining decision quality have structural roots; when burnout worsens, all other problems worsen.</p> <p>User Impact: a high rate of loss of expertise, poor decision-making quality in subsequent shifts, and potential for safety and compliance issues. Root cause: Managers are the only decision node. No second-in-command system, no automation of routine decisions, and no shift boundary enforcement.</p> <p>Long-term: This is a long-term solution</p>	3 to 4 times per week
6	<p>COD Cash Reconciliation - Manual Mismatch Tracking</p> <p>Problem: WMS handles COD orders, but reconciliation of cash collected vs. expected is done manually, leading to errors, disputes, and audit delays.</p>	Medium	<p>High time cost and error risk, but effects are post-shift and do not immediately increase chaos during peak hours.</p> <p>User Impact: The manager cross-checks cash registers against WMS logs for 20 to 40 minutes each shift. Only at the end of the shift were errors found.</p> <p>Root Cause: The WMS does not display a real-time cash-expected vs. cash-collected delta. Only after the shift do discrepancies become visible.</p> <p>Quick win: This is a quick-win pain point</p>	Daily

7. SOLUTION SPACE

Peak Assist Mode, Decision Support System for Dark Store Managers

7.1. Solution Overview

Core Concept: Peak Assist Mode	
A dedicated operational mode within the existing Store Manager dashboard that activates during peak hours - stripping away non-essential metrics, auto-prioritizing tasks, and surfacing one-tap actions to shift the manager's role from "figuring out what to do" to "doing it."	

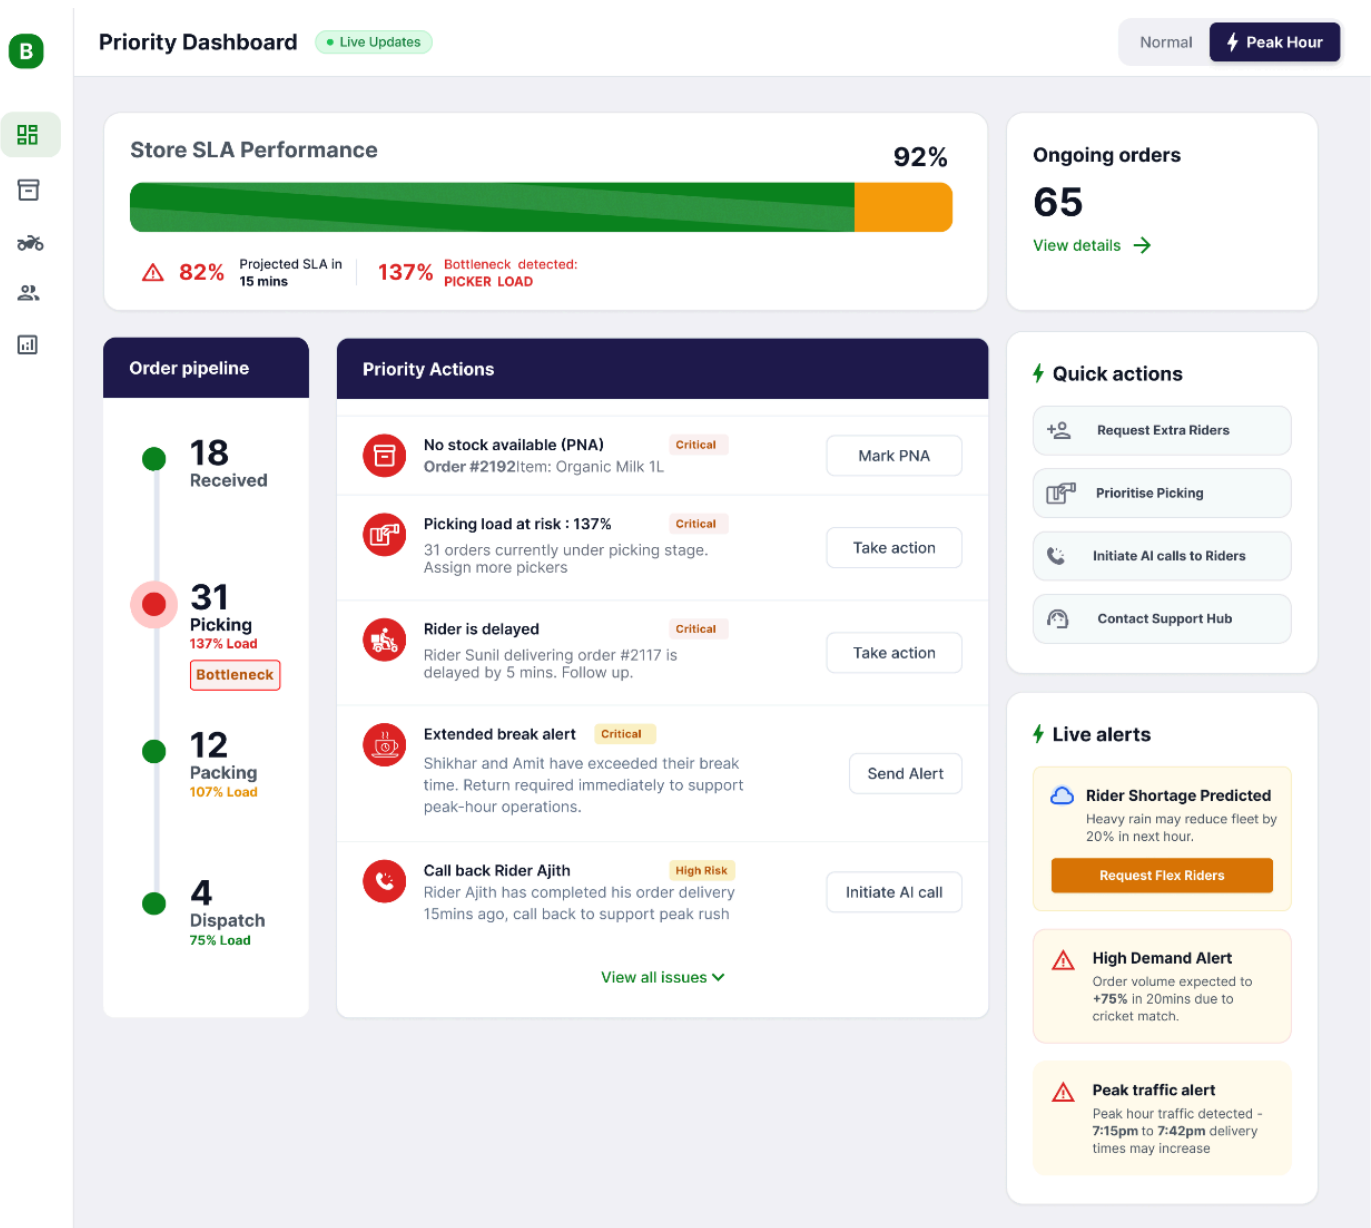
TODAY (Normal Dashboard)	PROPOSED (Peak Assist Mode)
Manager scans 15+ metrics to find issues	System surfaces top 3 - 5 prioritized actions
Manually calls riders, coordinates staff on WhatsApp	One-tap Quick Actions: Broadcast, AI calls, task re-prioritization
Reacts to SLA breaches after they happen	Predictive alerts flag breaches 8 - 10 mins before
Non-critical tasks (cash recon, audits) compete for attention	Non-critical tasks auto-deferred; only SLA-impacting items shown
Manager becomes picker when short-staffed	Instant backup pool broadcast fills gaps in 20 - 30 mins

Dashboard (Normal Mode)

[Prototype link](#)

The screenshot displays the 'Store Manager' dashboard in 'Normal' mode. On the left is a navigation menu with options: Priority Board, Inventory, Fleet Status, Staffing, and Reports. The main dashboard area features several key performance indicators (KPIs) at the top, such as Total Orders (1,245, +12%), Pending Orders (12, +2%), Avg Delivery Time (14m 32s, -15%), and Total Revenue (₹4,52,000, +8%). Below these are more detailed metrics like GMV (₹98,450, +18%), Velocity (5.8/min, +35%), SLA Losses (₹2,840, +42%), FBR Rate (94.8%, Target: 98%), and Queue Load (92%, CRITICAL). The 'Store Health Metrics' section includes an 'Order Pipeline: Bottleneck' chart showing 18 REVD, 31 PICKING, 12 RICKING, and 4 DISPATCHED orders, with a note about a bottleneck in the picking stage. The 'COD Tally' section shows pending COD orders for Vikram Singh (₹12,450) and Amit Kumar (₹8,200), with a total of ₹42,300. The 'Actions' section lists tasks like 'Staff Shortage - 2 Pickers Absent', 'Approve Write-off', and 'Cash Handover'. The 'Live Activity Feed' shows recent events like order deliveries and picker delays. The 'Rider Pipeline' section is divided into 'AT STORE & LOADING (12)' and 'EN ROUTE (SECTOR B) (8)', listing individual riders and their statuses (e.g., R1 Ramesh K. READY, R2 Amit B. STUCK IN TRAFFIC).

Dashboard (Peak Assist Mode)



Factors we considered while designing

- Glanceable - Every screen element must be parseable in under 3 seconds
- Action-first - Every surfaced item has a CTA; no information-only cards during peak
- Auto-prioritized - System ranks tasks by SLA impact, not chronological order
- Stress-aware - Reduced cognitive load; suppress noise, amplify signal

7.2. Solution Architecture

The solution has two pillars that work together to shift the manager from reactive firefighting to guided execution:

	Pillar 1: Prioritized Actions	Pillar 2: Quick Actions
What	Real-time, auto-ranked task queue sorted by SLA impact severity	One-tap shortcuts for high-frequency peak-hour operations
Why	Eliminates "what should I handle first?" decision paralysis	Eliminates "how do I do this quickly?" execution friction
How	System ingests order pipeline, rider status, staff capacity, and inventory signals to rank actions	Pre-built action flows triggered from a persistent sidebar panel
Example	"No rider available - Order #2185 placed 12 mins ago" → [Take Action]	"Request Extra Riders" → one tap broadcasts to rider pool

7.3. Detailed Feature Breakdown

7.3.1. Peak Mode Activation & Intelligence Layer

B Store Manager

Normal

Peak Hour

Toggle: Normal mode

B

Priority Dashboard

• Live Updates

Normal

⚡ Peak Hour

☰

📄

🔊

👤

Store SLA Performance

92%



82%

Projected SLA in 15 mins

137%

Bottleneck detected: PICKER LOAD

Ongoing orders

65

[View details →](#)

Toggle: Peak mode

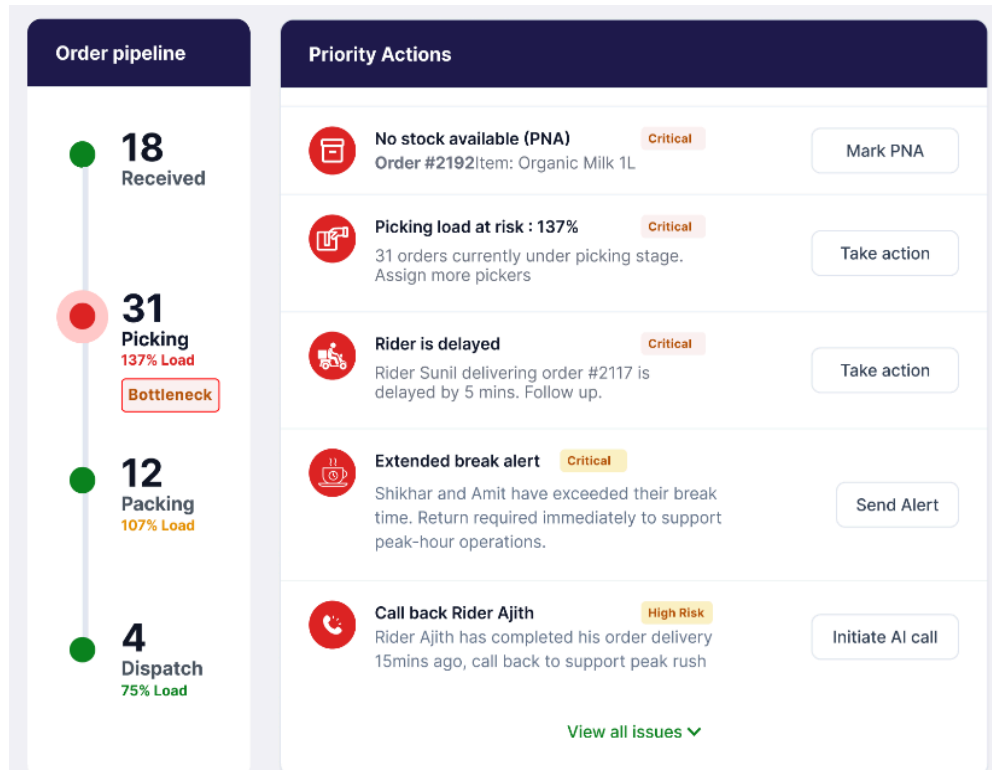
Peak Assist can be activated in two ways:

Trigger	How It Works	Manager Experience
Auto-Detection (Recommended)	The system monitors velocity, queue load, and SLA trend. When 3+ early stress signals are detected, it recommends activation.	Banner: "3 Early Stress Signals Detected - Peak Assist Recommended" with [Activate] button
Manual Toggle	Manager taps Normal/Peak Hour toggle based on experience (e.g., IPL evening, heavy rain).	Dashboard instantly switches to peak layout. Non-critical modules suppressed.

What changes when Peak Assist activates: Non-critical tasks (cash reconciliation, write-offs, audit items) are auto-deferred and marked DEFER. Dashboard collapses to show only: SLA Performance, Active Orders, Prioritized Actions, Quick Actions, Order Pipeline, and Rider Pipeline.

7.3.2. Prioritized Actions (The Brain)

A live, auto-updating task queue that answers the manager’s #1 question during peak: "What needs my attention RIGHT NOW?"



Priority Classification Engine

Priority	Tag	Examples	CTA
P0	CRITICAL	No rider available (order aging >10 min), PNA on high-demand SKU, SLA breach predicted in <10 min	[Take Action] - opens contextual resolution flow
P1	SLA RISK	Staff shortage (2 pickers absent, capacity down 15%), Picking bottleneck (31 orders averaging 6m 20s)	[View Pool] [Broadcast Now] - dual CTA for fast resolution
P2	OPERATIONAL	Inventory write-off approval, equipment issue, QR handoff glitch	[Review] - can be handled between peaks
P3	DEFER	Cash handover, reconciliation, non-urgent approvals	[Confirm] greyed out; auto-queued for post-peak

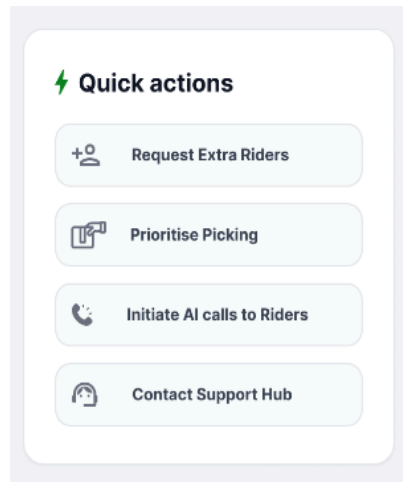
Prioritization Logic (Weighted Scoring)

Each actionable item is scored in real-time using:

- SLA Impact (40%) - Will this cause a delivery breach in the next 10 minutes?
- Order Value at Risk (25%) - How many orders / what GMV is affected?
- Time Sensitivity (20%) - How long has this been unresolved?
- Resource Availability (15%) - Can this be resolved with current staff/riders?

The queue re-ranks every 60 seconds. Resolved items disappear; new alerts surface at appropriate rank.

7.3.3. Quick Actions (The Hands)



A persistent right-side panel of one-tap shortcuts for the most frequent peak-hour actions. These replace the manual, multi-step workarounds managers currently rely on (WhatsApp groups, phone calls, verbal instructions).

Quick Action	What It Does	Replaces (Today)	Time Saved
Request Extra Riders	Broadcasts need to rider pool via delivery partner API; first-to-accept flow	Manager calls riders individually or messages WhatsApp group	~10 - 15 min → 1 tap
Prioritize Picking	Re-assigns staff from low-priority tasks (inventory sanity, quality audit) to picking/packing queue	Manager verbally re-assigns each person; tasks pile up untracked	~5 - 8 min → 1 tap
Initiate AI Calls to Riders	AI voice call (via Sarvam AI) to riders with pending pickups - automated ETA check & nudge	Manager personally calls 5 - 10 riders to check status	~15 - 20 min → 1 tap
Broadcast for Backup Staff	Notifies pre-registered flex picker pool with shift details; first-to-confirm gets assigned	Manager calls known part-timers via personal phone	~20 - 30 min → 1 tap
Contact Support Hub	Escalates to central ops with auto-attached store context (pipeline state, SLA %, active issues)	Manager calls support, verbally explains situation	~5 - 7 min → 1 tap

7.4. Moonshot: Autonomous Micro-Fulfillment Zones



Bold Idea - Robotic Pick-Pack for High-Velocity SKUs

Convert 15-20% of the dark store's floor area into an Autonomous Micro-Fulfillment Zone (AMZ) — a vending-machine-style automated storage and retrieval system that handles the top 200 high-velocity SKUs (which typically account for 60 - 70% of all orders). Human pickers handle only the long-tail. During peak, the AMZ absorbs the surge without adding headcount.

7.4.1 Why This Fits the Problem

Peak-Hour Bottleneck	How AMZ Solves It
Picking is the #1 bottleneck (31 orders stuck, 6m+ avg)	Robotic retrieval picks high-velocity items in <30 seconds vs. 3–4 min manual pick. Eliminates the picking queue for majority of order items.
Staff shortage during peak - manager becomes picker	AMZ handles 60 - 70% of picks autonomously. Remaining staff focus on long-tail items only. Manager never needs to step in as picker.
Picking errors rise during rush	Automated retrieval has near-zero pick error rate. Wrong item / missing item complaints drop significantly for high-velocity SKUs.
Rider wait time increases due to slow dispatch	Faster pick → faster pack → faster handoff. Rider idle time drops proportionally.
Fixed physical space limits throughput	Vertical automated storage uses 3–4x less floor area than shelving for the same SKU count. Frees floor space for packing stations.

7.4.2 How Does It Work?

Step	Component	Description
1	SKU Selection	AI identifies the top 150 - 200 SKUs per store based on order frequency and peak-hour demand patterns. These are loaded into the AMZ. (Think: Atta, milk, Maggi, Coke, eggs - the items in almost every order.)
2	Automated Storage	Vertical carousel or shuttle-based system (similar to AutoStore / Dematic models adapted for Indian scale). Compact vertical bins replace horizontal shelving - fits within existing store footprint.
3	Order Splitting	WMS auto-splits each order: AMZ-eligible items are picked robotically and sent to a merge station. Long-tail items are picked manually. Both converge at the packing station.
4	Merge & Pack	Packer receives pre-picked AMZ tote + manually picked items simultaneously. Packs and dispatches. Total pick-to-dispatch time drops from 6 - 8 min to 2 - 3 min for most orders.

7.4.3 Impact Projection

Metric	Today (Manual)	With AMZ
Avg. Pick Time (top 200 SKUs)	3 - 4 min per order	<30 seconds
Peak-hour throughput capacity	Limited by picker headcount	2 - 3x higher for AMZ SKUs
Picking errors (top SKUs)	3 - 5% during peak rush	<0.5% (automated)
Staff dependency during peak	Need 6 - 8 pickers at full load	Need 2 - 3 pickers (long-tail only)
Manager stepping in as picker	Frequent during peak	Eliminated

7.5. Supporting Modules (Visible in Peak Mode)

Beyond the two core pillars, Peak Mode retains a minimal set of health indicators that provide situational awareness without adding noise:

Module	What It Shows	Why It Matters in Peak
SLA Performance Bar	Real-time % against 98% target with color-coded health	Single metric that tells manager if things are OK or not
Order Pipeline	Orders at each stage: Received → Picking → Packing → Dispatched, with bottleneck alert	Instantly shows where orders are stuck (e.g., "31 in Picking, avg 6m 20s")
Picker & Packer Efficiency	PPH (Picks/Hr) and Orders/Hr per person with break alerts	Identifies underperformers and overdue breaks at a glance

7.6. MVP Scope & Phasing

Phase	Scope	Key Deliverables
MVP	Peak Mode toggle, Priority Actions queue (rule-based ranking), Quick Actions (Request Riders, Prioritize Picking, Contact Support)	Manual activation, basic priority scoring, 3 quick actions, SLA bar + order pipeline
V1.1	Auto-detection trigger, AI rider calls (Sarvam AI), Backup staff broadcast	Stress signal detection, 2 additional quick actions
V2.0	Predictive SLA Guardian (ML-based), order-level breach predictions, pre-positioned rider suggestions	ML model trained on historical patterns, per-order risk scoring, proactive rider re-routing

7.7. Success Metrics

Metric	Current Baseline	MVP Target	V2 Target
Peak-Hour SLA %	~88 - 90%	93 - 94%	96%+
Avg. Manager Decision Time	~3 - 5 min per issue	<1 min per issue	<30 sec (auto-resolved)
Peak-Hour SLA Loss (₹)	High (refunds + comps)	Reduce by 20 - 25%	Reduce by 40 - 50%
Manager Actions/Hour (Peak)	~8 - 10 (slow, manual)	15 - 18 (guided, faster)	20+ (semi-automated)
Manager NPS / Satisfaction	Low (burnout reported)	Measurable improvement	Sustained positive trend

7.8. Trade-offs & Limitations

Trade-off	What We Chose	What We Accepted
Simplicity vs. Completeness	Show only 3 - 5 priority items during peak	Managers may miss lower-priority items; mitigated by post-peak queue
Automation vs. Control	System recommends; manager confirms with one tap	Slightly slower than full automation, but maintains manager trust & accountability
Rule-Based MVP vs. ML	Start with weighted scoring rules for prioritization	Less accurate than ML predictions initially; faster to ship & iterate
Deferred Tasks	Auto-defer cash reconciliation & audits during peak	Cash reconciliation delay risk; mitigated by auto-pause label + post-peak reminder
Rider Pool Dependency	Quick Action broadcasts to existing delivery partner fleet	Works only if rider pool exists; no impact in areas with structural rider shortage

Guiding Philosophy

Peak Assist Mode is not a new product - it is the right lens on the existing system at the right time. The data and actions already exist across multiple screens. We are collapsing them into a single, stress-optimized control layer that transforms the manager from an overwhelmed firefighter into a guided operator.