



Axis Fulfilment ships ~487,000 orders across 18 months, spanning three distribution centres and four carriers. This deck traces where that pipeline holds, and where it's straining.



// LATEST-MONTH + Δ MoM TILES READ THROUGH MAY 2026 • THE LAST COMPLETE MONTH

ORDER VOLUME BY MONTH




Delivered Rate and ship-to-deliver time are the two SLA measures that matter once an order leaves the warehouse door. The pages ahead trace **each one to its root cause**.

>> Start with the pages that name a specific warehouse, carrier or zone: each is a single, addressable fix.

Of every order placed, ~85% reach **Delivered**. The rest sit In Transit or are Cancelled before completion.

COUNT COMPLETION • PLACED → DELIVERED



TERMINAL STATUS MIX

Status	Orders
Delivered	411,178
InTransit	50,446
Cancelled	24,602
Packed	351
Placed	25
Picked	18
Total	486,620

DELIVERED RATE

84.5%

CANCELLATION RATE

5.1%

// ORDER LIFECYCLE • MEAN MILESTONE DURATIONS • PICK → PACK → SHIP




The gap between Delivered and Placed is mostly **InTransit: orders still moving, not lost**, with a small, steady Cancelled tail.

>> Watch the terminal_status mix for a rising Cancelled share; it's the funnel's earliest leak signal.

// DEGRADATION • SYDNEY PICK-TO-PACK • SINCE MID-MAY 2026

Sydney's pick-to-pack time has run **~1.5× its baseline** since mid-May 2026. Melbourne and Brisbane haven't moved. This is a Sydney-specific shift.

MEAN PICK-TO-PACK MINUTES BY MONTH • BY WAREHOUSE



SYDNEY • PICK-TO-PACK

45.4
since mid-May · ~1.5× baseline

MELBOURNE • STEADY

30.0
unaffected

BRISBANE • STEADY

30.0
unaffected

Melbourne and Brisbane haven't moved. This reads as a **Sydney-specific operational shift**, not a company-wide slowdown.

>> Investigate Sydney's pick-to-pack workflow for the mid-May change before it compounds into ship-to-deliver SLA.

/axisanalytica

Ship-to-deliver performance isn't one number. It's a **carrier × warehouse matrix**. No carrier is best everywhere: the pairing is what matters, and NationalParcel inverts (best at Brisbane, worst at Sydney).

MEAN SHIP-TO-DELIVER HRS • CARRIER × WAREHOUSE


Carrier	Brisbane D.C.	Melbourne D.C.	Sydney D.C.	Total
EconomyShip	86.3	86.4	86.4	86.4
MetroExpress	71.9	50.2	48.0	54.8
NationalParcel	50.2	55.5	70.3	59.3
RapidPost	57.7	48.0	47.9	50.4
Total	63.3	56.6	60.3	59.6

RAPIDPOST • STRONGEST OVERALL

50.4
premium, strong everywhere

ECONOMYSHIP • SLOWEST OVERALL

86.4
uniformly slow, every warehouse





// FASTER • LIGHT VIOLET → DARK VIOLET • SLOWER

No single carrier or warehouse ranking tells the whole story. The **pairing is what matters**, which is exactly what a flat carrier-only or warehouse-only view would hide.

>> Route by carrier×warehouse pairing, not carrier rank alone, especially for NationalParcel at Sydney.

On the NSW Central Coast, **NationalParcel** runs ~2× its own network-wide ship-to-deliver time, the network's one clear carrier-specific service gap.

MEAN SHIP-TO-DELIVER HRS BY CARRIER • CENTRAL COAST (NSW 2250-2299)



NationalParcel's network-wide average sits in line with peers, so this reads as a **route problem, not a carrier-wide one**.

>> Route Central-Coast NationalParcel volume to Melbourne or Brisbane DCs, or renegotiate the CC leg.

Delivery cost isn't random. It's a monotonic ~4× gradient from Zone A to Zone D, driven by distance from the metro core.

MEAN DELIVERY COST BY ZONE (\$)



ZONE D • REMOTE

\$32.57

~4× Zone A

ZONE A • INNER METRO

\$7.75

the cheapest leg

TOTAL DELIVERY SPEND

\$6.24M

Zone C+D = ~53% of it

SHARE OF ORDERS

SHARE OF SPEND

ZONE A

ZONE B

ZONE C

ZONE D

0%

50%

100%

Remote delivery legs cost more to run: Zone D's premium reflects genuine **distance-to-doorstep economics**, not a pricing error.

>> Feed the zone gradient into delivery-fee tiers or minimum-order thresholds for Zone C/D.

// MODEL • WHY YOU CAN TRUST THESE NUMBERS

Four fact grains and **one accumulating snapshot** let this model answer revenue, SLA and stock questions from a single, denormalised order pipeline.

Every order is one accumulating-snapshot row.

fact_order_lifecycle carries milestone timestamps (placed → picked → packed → shipped → delivered) on one row per order, not one per event: a textbook Kimball accumulating snapshot.


Orders, lines and stock each get their own grain.

Header (fact_orders), line (fact_order_lines) and weekly inventory (fact_inventory_snapshot) sit at three deliberate grains, so revenue, SKU-level GMV and stock questions each query the table built for them.

Every figure traces to source.

Each number reconciles back through governed marts to the originating fulfilment CSVs: one auditable path from the answer-card to fact_orders / fact_order_lifecycle.

// THE ACCUMULATING SNAPSHOT, ONE ROW PER ORDER



// HEADER vs LINE PATH BY WAREHOUSE • DELTA 0 EVERY ROW

Warehouse	header path	line path	Δ
Melbourne Distribution Centre	\$83,895,785	\$83,895,785	\$0
Sydney Distribution Centre	\$73,914,008	\$73,914,008	\$0
Brisbane Distribution Centre	\$52,295,830	\$52,295,830	\$0

// FOUR FACTS, FOUR GRAINS

fact_orders

487K

one row per ORDER (header)

fact_order_lines

1.46M

one row per (order, SKU) LINE

fact_order_lifecycle

487K

ACCUMULATING SNAPSHOT per order

fact_inventory_snapshot

2.34M

one row per week/warehouse/SKU

// FOR THE TECHNICAL READER

CSV → fulfilment_ops marts (SQL views) → DAX measure → answer-card

