Know what every grid change will deliver before you make it

How Veloq Simulations helps operators test investments and optimize the grids they already run

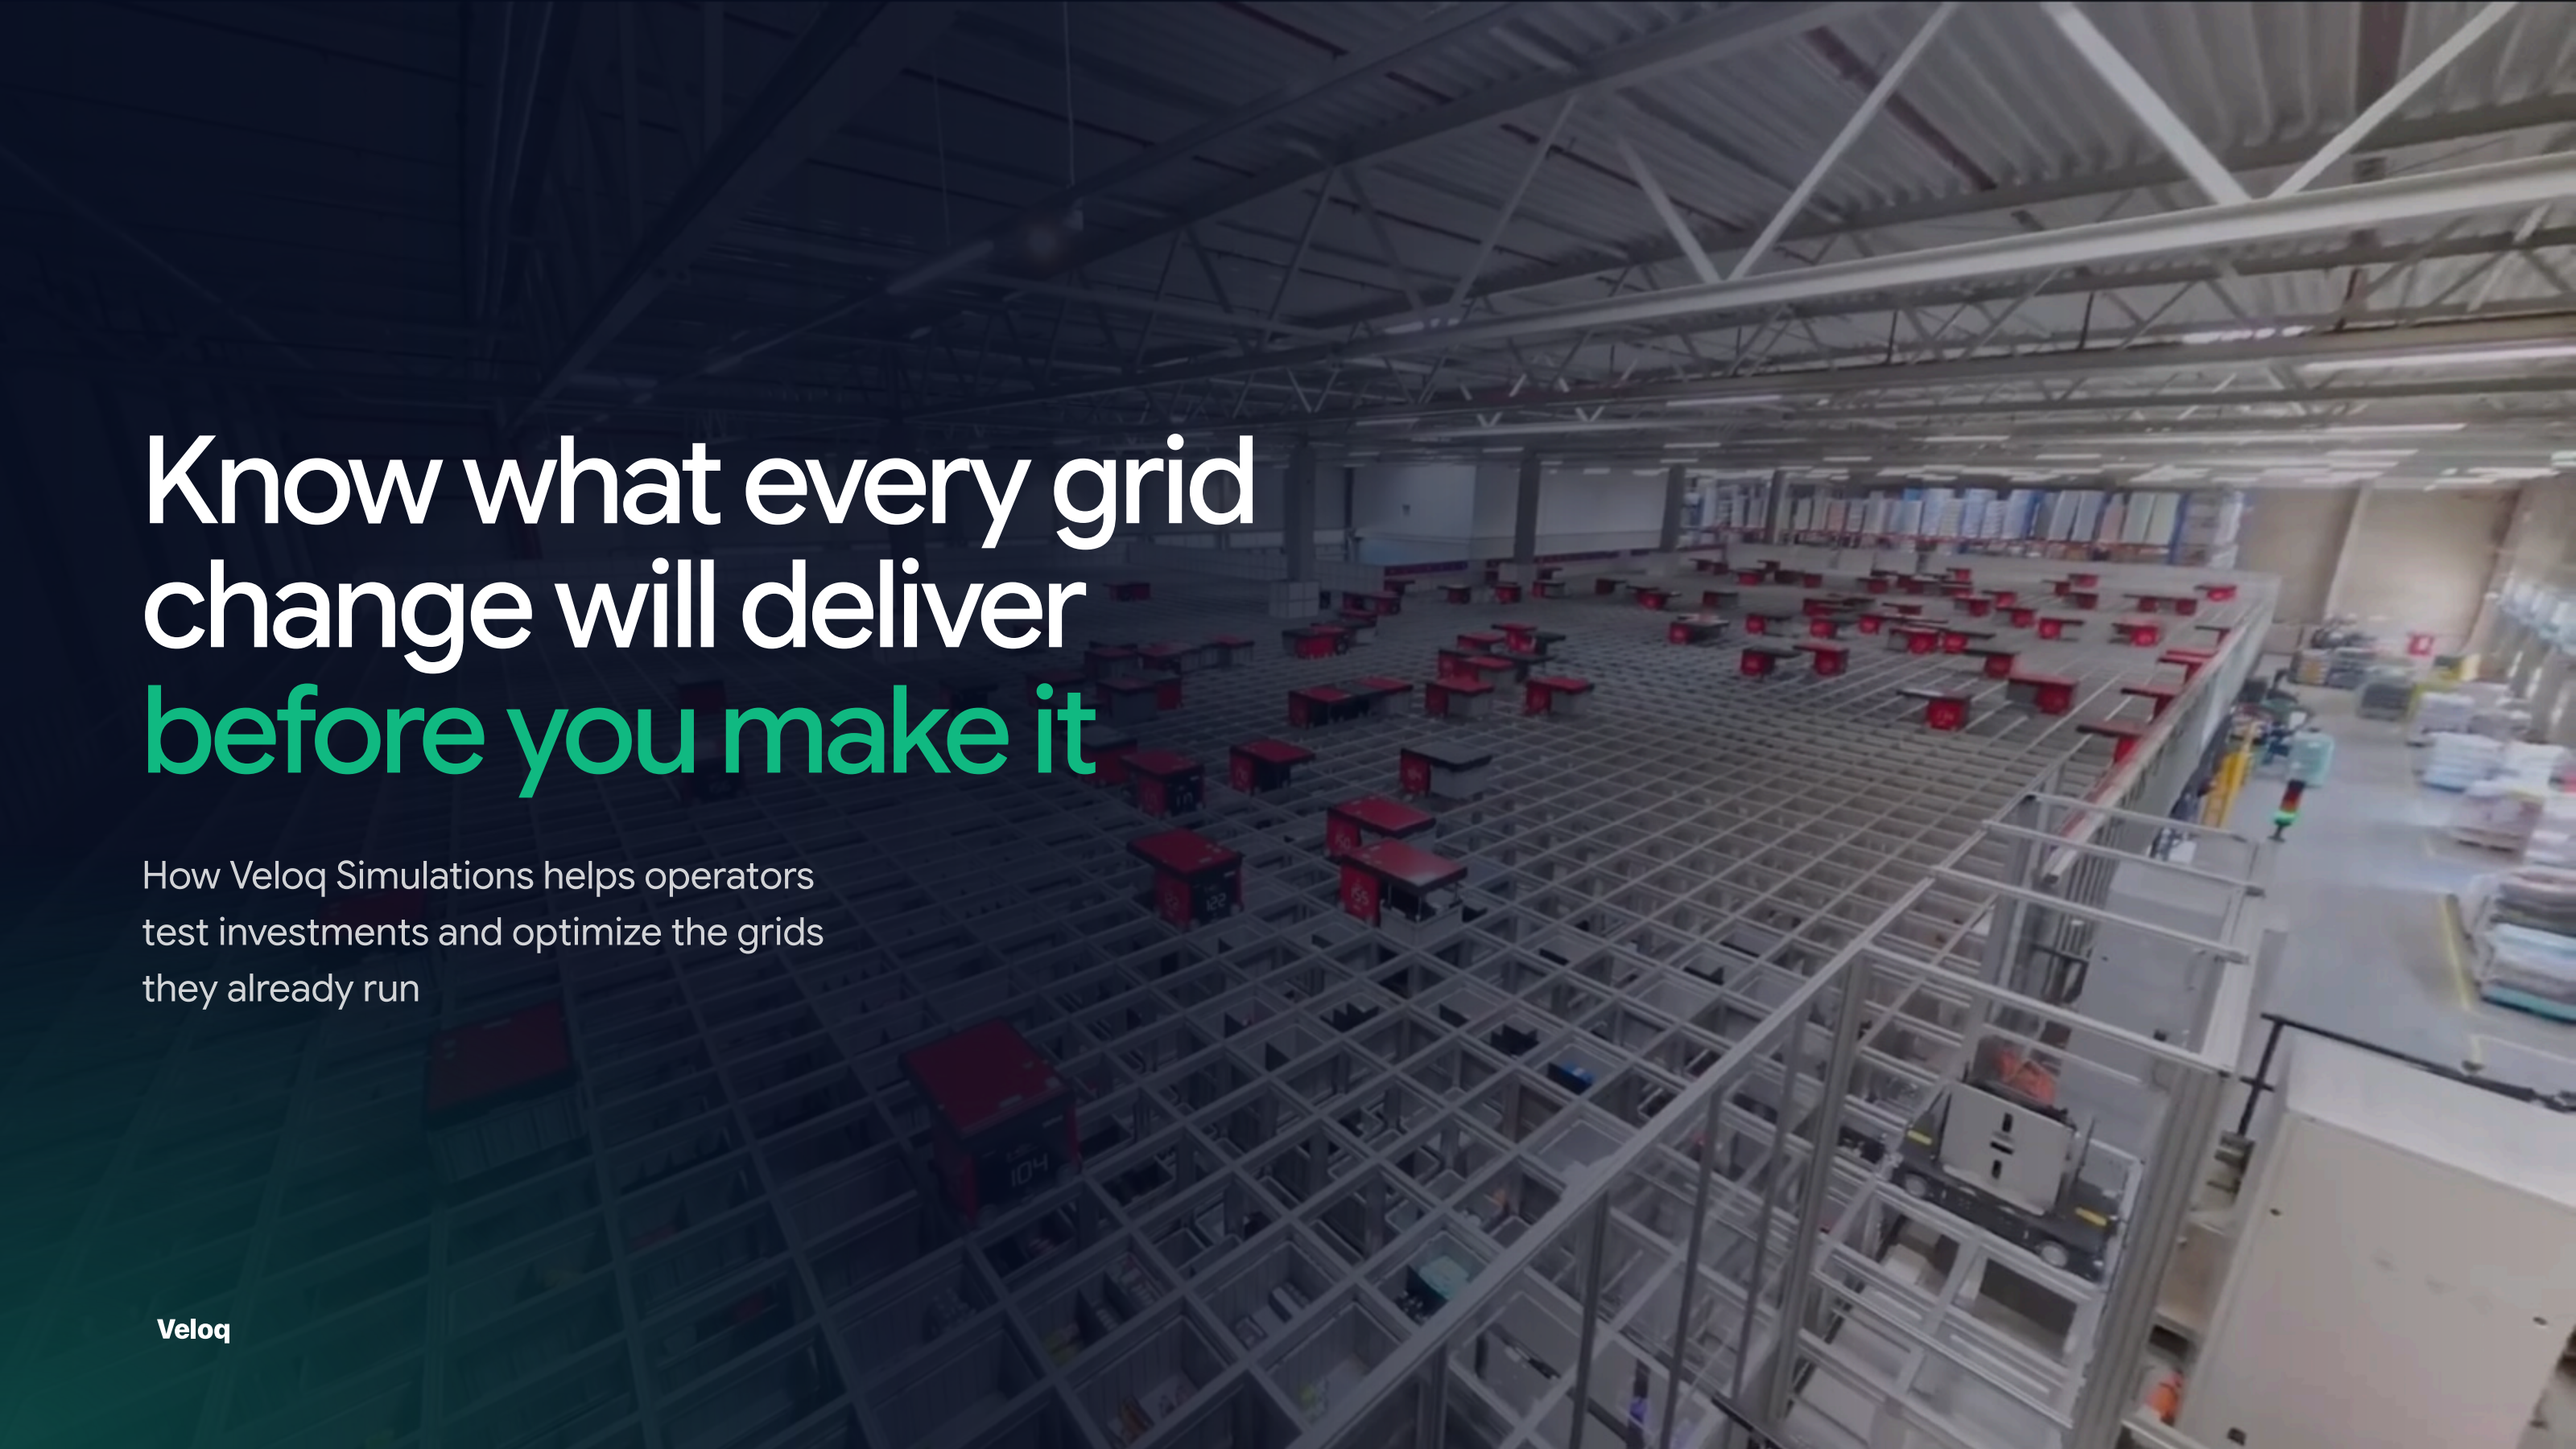

Veloq

SUMMARY

Robots and ports are the two largest spend items in an AutoStore grid, but they aren't the only decisions that shape its output.

Opening hours, port roles, assortment, and peak plans all decide how much a grid delivers, and each is hard to test on a live floor without putting service at risk.

Veloq Simulations changes the starting point. It is a discrete-event digital twin of each live grid, built from Rohlik Group's own production data. It replays a real day of picking and putaway, task by task, with real timestamps, real port hours, and real bin inventory. Then it changes one thing at a time (more robots, more ports, longer picking hours, a different SKU mix) and shows what that change does to output.

Simulations is a module of the Veloq core platform. It runs today across 14 AutoStore grids at 7 fulfillment centers, and rebuilds the inputs for all 14 grids in six minutes.

14
live AutoStore grids modeled

7
fulfillment centers

~6 min
to refresh every grid

Real days
replayed, not forecast demand

THE PROBLEM

The most expensive decisions get the least evidence

An AutoStore grid is sized before it is built. Vendors and integrators model expected demand, pick a robot count and a port layout, and the operator signs off. This gets a grid designed and installed.

The trouble starts afterward. Once the grid is live, the questions become: Should we add robots or ports before peak? Is the grid really full, or just badly balanced? Would longer picking hours buy more than new hardware? These are capital and service decisions, often worth more than the original sizing exercise, and no operator has publicly claimed to have a tool built to answer them against how their grid actually runs.

As a result, decisions fall back on rules of thumb, vendor quotes, and the experience of whoever is in the room. Sometimes that works. When it doesn't, the cost shows up as robots that sit idle behind a port bottleneck, or as a peak week the grid can't handle.



A working model of each grid in operation

Veloq Simulations builds a working model of each grid in operation and lets the operator ask it questions before spending money.

It replays real days.

Demand isn't forecast or averaged. The twin takes every pick and putaway task from a chosen production day and runs it through the modeled grid in the order and at the time it really happened.

It models the grid as it runs.

Robots, ports, port opening hours, bin positions and the real inventory in the grid on that day.

It changes one thing at a time.

Add robots, add ports, extend hours, change the SKU mix or the robotic-arm rules, and compare the result against the same day run as is.

It stays current.

Inputs are pulled from production by script and refreshed in minutes, so the twin tracks the grid as it changes rather than freezing at the design stage.

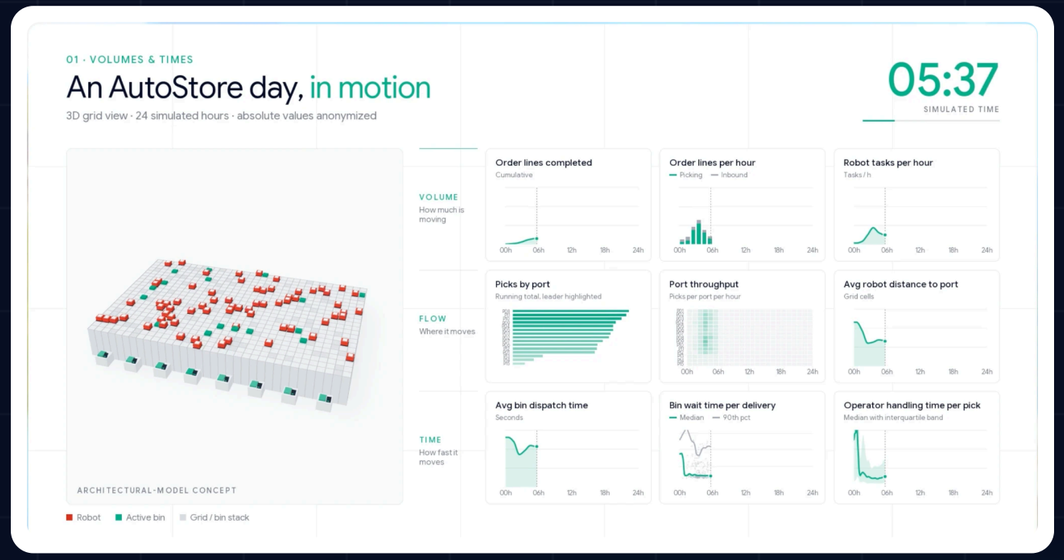

A simulated day in motion: volume, flow and time, hour by hour. Values anonymized.

Why operators haven't had this: Design-stage simulation vs. operational simulation

Simulation of AutoStore grids isn't new. What is rare is simulation of a grid in operation, fed by its own production data and kept current.

This is because an operational twin is only as good as the data it replays: every pick, putaway, bin movement and port event, timestamped and checked against the grid it came from. Vendors and integrators build their models before the grid exists, and once it's handed over, they don't sit on the system that runs it day to day. For an operator, building that data feed is a project in its own right. It means extracting, cleaning and reconciling task-level records, agreeing on definitions, and keeping all of it current as the grid changes.

In Veloq's experience, this data work takes about two-thirds of the total effort, more than the simulation itself. Veloq Simulations starts with that part already done, because the platform that runs the grid is the one recording the data.



~2/3

of the effort in building an operational twin is data work. On the Veloq platform, it's already done.

Built to size a grid, or built to run one

	TYPICAL VENDOR / INTEGRATOR SIMULATION	VELOQ SIMULATIONS
Who builds it	The vendor or integrator, as part of the sale	Veloq, on the operator's own data
When it's used	Design stage, before the grid exists	Throughout the grid's working life
Demand input	Forecast or profiled demand	Real production days, task by task
Scope	One planned grid	Every grid in the network, on one pipeline
Refresh	Rebuilt for a new project	Minutes, from production data
Questions it answers	How big should the grid be?	What should we change next, and what will it buy?

No other AutoStore operator publicly reports running a simulation of its live grids, fed by production data, across a multi-country network. Published work falls into three groups: design-stage tools from vendors and integrators, academic models, and operators of proprietary grid systems simulating their own technology.

HOW IT WORKS

From the system that runs the grid to a decision

Veloq WCS → Veloq operations data → Automated export → Simulation model → Decision outputs

VELOQ PLATFORM DATA

Because the Veloq WCS already records every task, bin movement, and port event, the data the twin needs exists before anyone asks a question. A scenario layer then sits on top:

- Double-load stress days
- Fixed daily targets
- Extra robots or ports
- Changed opening hours
- A different SKU mix
- New robotic-arm rules

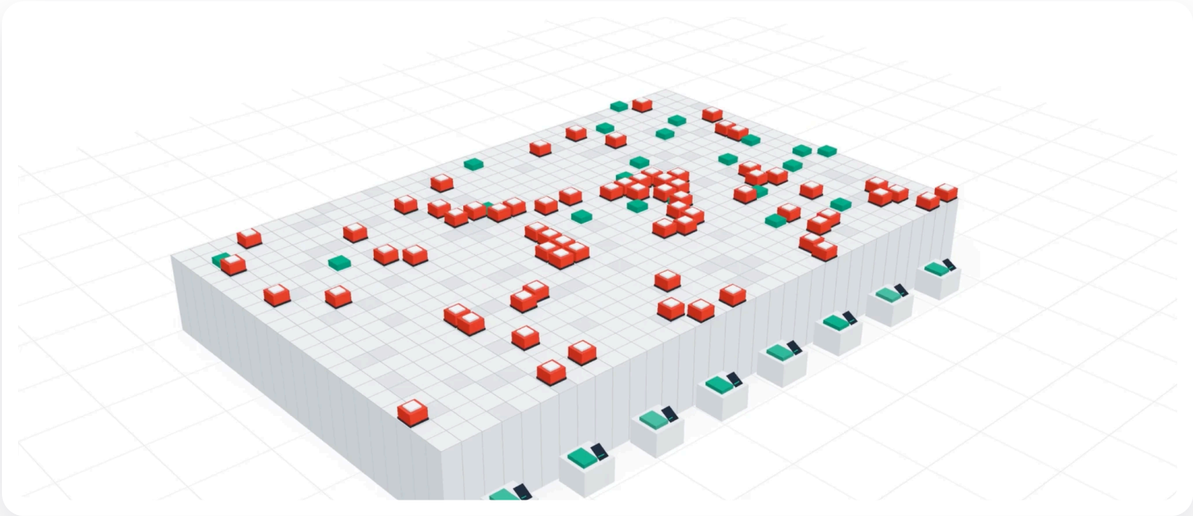

The modeled grid: robots, active bins and ports, replayed task by task.



INSIDE THE TWIN

What it models, and how every input is checked

What the twin represents

- **Demand:** every pick and putaway task from a real day, with its real timestamp.
- **Grid:** layout, depth, bin positions and capacity, reconciled to the AutoStore console.
- **Fleet:** the active robot fleet, not the nameplate count.
- **Ports:** each port's type, role and real opening hours.
- **Handling:** pick and putaway times measured from the grid's own port events.

How every input is checked

Every input is pulled from production by script. Nothing is typed by hand, and every run can be repeated and audited. Before a grid is modeled, its data passes a standard set of checks:

- Task counts reconcile to production records.
- No orphan bins and no duplicate records.
- Grid capacity derived from the data matches the slot count in the AutoStore console.
- Structural invariants of the grid hold in every file.

WHY SIMULATE YOUR GRID?

Every change to a live grid is an experiment run on real customers

Add robots, and it takes weeks to learn whether they helped. Change port roles before peak, and you find out on the busiest day of the year. Floor results are noisy too: demand, staffing and assortment shift from week to week, so it's hard to tell what the change itself did.

Simulation moves the experiment into software. The same real day runs with and without the change, so the difference in output comes from the change and nothing else.

INVEST

It keeps capital out of hardware that won't raise output.

OPTIMIZE

It finds capacity the grid you already own can deliver with better rules, hours and port plans.

At Rohlik, operations teams use Veloq Simulations both ways: to size investments and to tune grids that are already running.



QUESTIONS YOU CAN ANSWER

The kinds of decisions it's built for

01

Robots or ports?

Will more robots raise daily output, or are ports and opening hours the real limit? Adding robots to a port-limited grid buys congestion, not throughput. The twin shows which constraint binds before the purchase order goes out.

02

Where does the grid break at double load?

Stress days built by folding real days together show where each grid runs out of capacity first. That tells you where peak investment should go, grid by grid, instead of spreading it evenly.

03

Should fast movers leave the grid for peak?

Moving top sellers out of the grid frees capacity, but it can also remove the most efficient work the grid does. The twin measures both sides.

04

How much can operating changes deliver before buying hardware?

Longer picking hours, pre-picking, earlier expedition cutoffs and different inbound timing all cost less than robots and ports. The twin shows how far they go, and where hardware becomes the only answer.

05

What hand-over rule should robotic-arm ports use?

When robotic picking shares a grid with people, the rule for which tasks go to the arm changes throughput for both. Test the rules against a real day's mix before you commit to one.

06

How full can the grid get?

More SKUs in the grid eventually cost throughput as bins dig deeper. The twin finds the point where assortment growth starts slowing the grid down.

07

How should ports be split on a peak day?

Picking, putaway and mandatory work all compete for the same ports. The twin tests different splits against a peak-day profile so the plan holds when volume arrives.

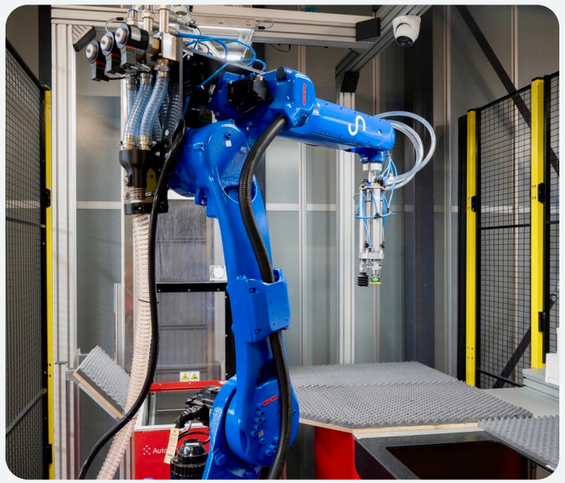

02 · UTILIZATION

How hard the grid works, hour by hour

Robot traffic heat on the grid · port and robot utilization in %

10:38
SIMULATED TIME

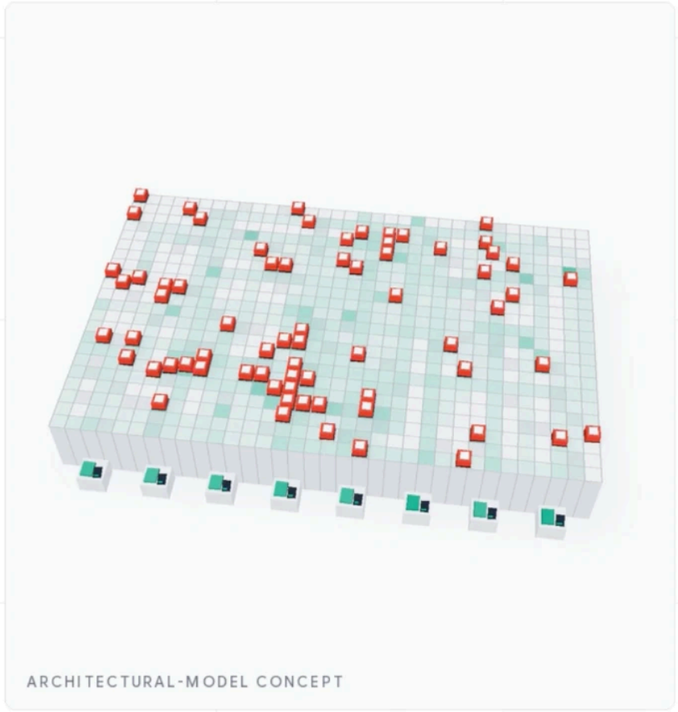

PORTS
How busy each port is

ROBOTS
Fleet health and bin delivery

FLOW
How work arrives

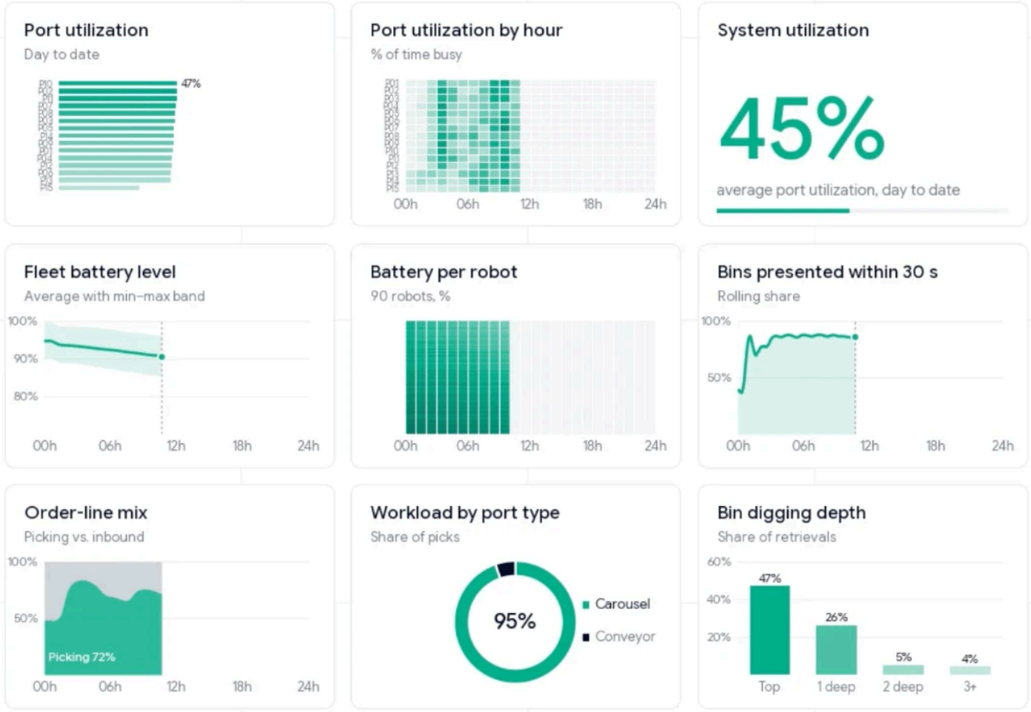

How hard the grid works, hour by hour: robot traffic, port and robot utilization, bin delivery and digging depth. Values anonymized.



Veloq

HOW TO USE IT

Four rules for trustworthy answers

1

Compare options on identical inputs.

The twin is built to answer "A beats B, and by roughly how much." Running every option against the same real day removes the noise that makes before-and-after comparisons on the floor unreliable. That is how capital decisions should be made.

2

Stress with real days, not invented ones.

Double-load scenarios are built by folding real days together, so the shape of demand through the day stays true.

3

Set fixed targets from whole real orders.

When testing a daily volume target, the twin tops up demand with complete real orders, spread across the hours of the day, so the basket mix stays realistic.

4

Agree the definitions first.

An order line is one distinct pick task, not one item. Tasks never span grids. Hourly counts are never summed into a daily total. Small definitional choices like these decide the answer more often than the model does.

RUNNING IN PRODUCTION TODAY

Live across 14 grids

Veloq Simulations currently models 14 live AutoStore grids at 7 fulfillment centers for Rohlik Group. A full refresh of the inputs for every grid takes about six minutes, around 30 seconds per grid. Operations teams use it to test robot, port, peak, and assortment decisions before they are made.

~6 min
full refresh, every grid

~30 s
per grid

WHERE IT FITS IN THE VELOQ PLATFORM

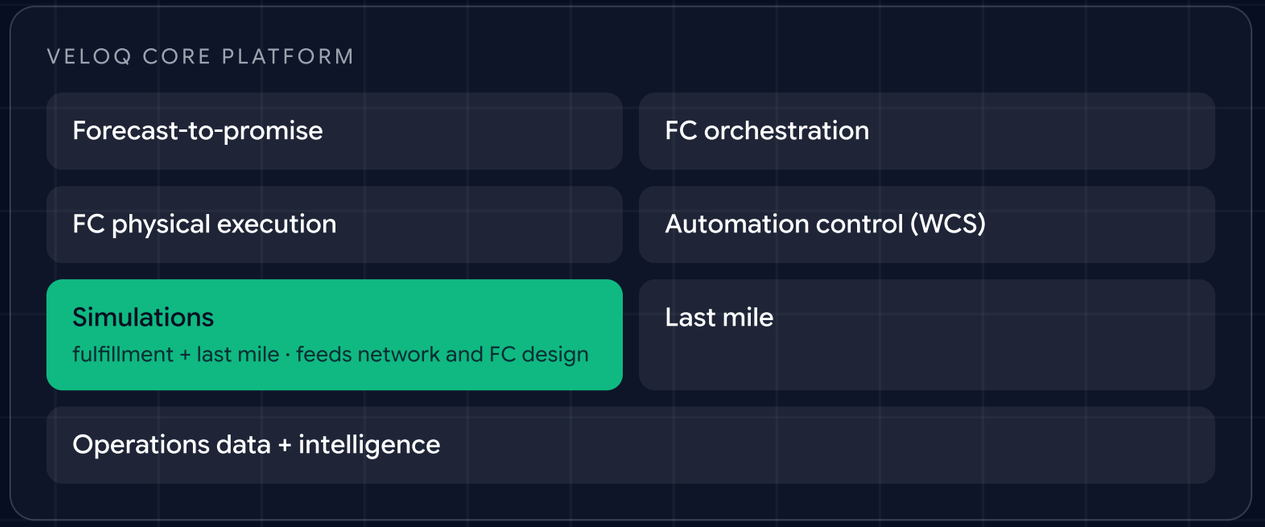

Simulations is part of the Veloq core platform alongside Automation control, FC physical execution and Operations data and intelligence, and feeds network and FC design. The same platform that runs the grid every day supplies the data the twin replays, which is why it can stay current without a separate data project.

The scope covers fulfillment and last mile, so the same approach extends from the grid to the wider operation.

MODEL YOUR OWN GRID

Every grid has its own binding constraint

Some are port-limited, some run out of hours, some are fuller than they look. The way to know which applies to yours is to run your own days through a model of your own grid.

Veloq can build that model from your production data and whichever WES runs your grid today. You'll see where your grid's limits sit, which changes would move output, and what the Veloq platform behind those answers could do for the rest of your operation.

Contact us Talk to Veloq about modeling your grid

Veloq

Profitable same-day e-grocery.
Automated.