

ETHOS AUTOMATION / CASE STUDY

Prove It First

How a 300 part test run won the job before the purchase order arrived

Bin picking is a solved problem. Rack picking is not. This one was won by running 300 parts in front of the customer and writing down every failure.

AT A GLANCE

- 294 of 300** Parts placed successfully in the proof of principle run
- 0.3 mm** Placement accuracy, under one degree rotational deviation
- 1.1 - 1.3 m** Camera standoff window that made it work
- 150 + 150** Left hand and right hand parts in the run
- 4** Criteria logged by hand on every pick
- Before the PO** When the capability was proven

AUTOMATION

There is a term circulating in manufacturing for the new plants coming online with minimal human intervention: ghost factories. They are built from fresh capital, which means every bin and every rack in them is engineered for precision from day one. Parts arrive in the same position and the same orientation, thousands of cycles in a row, because the containers were designed alongside the automation that empties them.

Established plants do not have that luxury. They run hundreds of existing bins and racks, tied to line layouts that predate the robots, damaged and warped by years of service. Replacing that infrastructure to buy repeatability would cost more than it returns.

So the interesting engineering question in automation right now is not how to build a ghost factory. It is how to get ghost factory reliability out of equipment that was never designed for it.

This is a project about exactly that.

The Setup

A major Tier 1 automotive supplier came to Ethos Automation about automating a de-racking operation. They wanted stamped parts pulled off shipping racks by a robot, found by an AI vision system, and placed ready for the next operation. The question on the table was not price or lead time. It was whether the thing was possible at all.

So before any money changed hands, the customer shipped Ethos a robot.

They free-issued a FANUC ArcMate, two racks full of production parts and two empty racks, and sent them to Brantford. The arrangement was straightforward. Ethos would build a working cell on its own floor, run real parts through it, measure what happened, and write it up. If the study convinced them, a purchase order would follow. If not, both companies would have learned something cheaply.

That is an unusual amount of trust to extend to an integrator, and an unusual amount of exposure for the integrator to accept. It is also, arguably, the correct way to buy automation nobody has built before.

Why Racks Are Harder Than Bins

Bin picking is the better known problem, and the easier one. A bin is a box. Its walls can be described to the software as fixed obstacles, and a path planner can route around them.

Racks are worse in every dimension that matters. They vary from rack to rack. Their geometry is non-standard. And critically, they warp over time, so the shape the software expects is not the shape in front of the camera. The collision avoidance that works inside a bin cannot simply be pointed at a rack and switched on.

The vision came from Apera AI, whose system builds a 3D reconstruction from a stereo pair of 2D monochrome cameras, combining classical feature matching with AI depth estimation so that shadows, low contrast and partial occlusion do not leave holes in the point cloud. A second model, trained from scratch for the specific part using millions of simulated examples, then finds the part inside that reconstruction and scores how confidently it could be picked.

That confidence score matters more than it sounds. Parts in a shipping rack are not where a CAD model says they should be. They lean, they shift in transit, and they clump. A system that reports how sure it is can be told to skip the ones it is unsure about, which is a different and more useful behaviour than a system that simply reports a position.

The Part Is Not the Problem. The Rack Is.

Stamped components arrive stacked on horizontal rods inside a steel rack, roughly seventy five parts to a rod, packed tight. A robot reaching in has to fit between the rods, clear the rack's front lip, grip a part that may be touching its neighbours, and withdraw without dragging anything else out.

The Ethos team distilled the constraint into a single rule of thumb that appears in the study report:

If the operator cannot load or unload the part while keeping it flat, the robot, as currently designed, will be unable to pick it up.

That sentence is worth pausing on, because it defines the boundary of the system honestly. The robot is not more dexterous than a person. It is more consistent, and it never gets tired, but it inherits every defect in the rack that a human hand would work around: warping in the upper holding bar, warping in the lower bars, rust buildup that makes parts screech as they slide, welds protruding where the surface should be smooth, and parts clumping at the front lip.

Backing the Camera Off to See More

The camera sits on the end of arm tool, which makes it an eye-in-hand application: where the robot looks and where the robot reaches are the same problem.

That created an immediate difficulty. The parts are large enough that the whole component will not fit in the camera's field of view at close range. To capture a usable image the camera has to stay at least a metre back, which is the opposite of what you want for resolution.

The answer was to stop treating the capture as a single event. After the initial image, the camera moves incrementally forward, holding the part within a 1.1 to 1.3 metre window where the resolution is highest and the whole part is still in frame. Working within that band, the cell hits 0.3 mm placement accuracy with under one degree of rotational deviation.

The customer also asked that operators not have to tell the system how deep a new rack is loaded. So the robot establishes it itself: a depth estimation pass identifies the nearest part and sets a baseline distance, and mathematical offsets place the camera at the optimal range from there. Every pick after that starts from the best available vantage point rather than from an operator's estimate.

Building a Tool That Fits

The central mechanical problem was the end of arm tool, and the study report names it plainly: the tool had to be compact enough to navigate inside the part rack and strong enough to manipulate the parts reliably, without unintentionally picking up more than one.

Those two requirements pull in opposite directions. Magnetic holding force scales with magnet size. Reach into a crowded rack scales inversely with tool size.

Ethos works this problem at the front end rather than discovering it after fabrication. Robot, tool and the full motion path are modelled in Siemens Process Simulate before any hardware is cut, and the tool design is stress-tested across thousands of randomized rack configurations to see whether it still meets requirements under real-world variability. That is how you find out you need a reorientation pedestal, or a different tool geometry, while it is still a drawing.

Even so, the magnets brought a failure mode that only appears in metal. A magnetic gripper does not fail cleanly. It holds, mostly, until the part shifts far enough that the field cannot keep it, and then the part slides. So the team defined failure in measurable terms rather than by feel: a magnet failure was any case where the robot needed three or more attempts to pick a part.

300 Parts, Counted One at a Time

On 25 June 2025 the cell ran 300 parts, 150 left hand and 150 right hand. Every pick was logged by hand against four criteria: did the magnets engage, did the part touch the rod, did the part touch the front lip, and did it drop successfully.

The speed testing produced the most useful finding. Running the robot between 30 and 50 percent of maximum, the team found that at 50 percent parts occasionally slipped off the magnet during certain manoeuvres, and two were dropped outright. At roughly 45 percent the cell was stable and still comfortably inside the cycle time target. The programme was adjusted to slow specific moves rather than the whole cycle.

A second finding ran against intuition. Cycle time was not meaningfully affected by how deep into the rack the robot had to reach. The variation came almost entirely from failed picks that needed re-attempts. In other words, the cost of a difficult rack position is not time, it is reliability, and the two are not interchangeable.

Writing Down What Did Not Work

The most unusual thing about the Proof of Principle report is not the success rate. It is the section listing what was still wrong.

The report states directly that the recognition model was not finished, and then enumerates three specific limitations. When the rack is full and the camera is at its furthest distance, the system could identify the rack's own pole as a candidate part. One particular column consistently scored lower pick confidence than the others, sometimes low enough to be excluded from processing. And the system had no way to report that a rack was empty.

None of that had to go in a document written to win a purchase order. Putting it in is what makes the other numbers credible, and it gave both companies a shared list to work from rather than a surprise during commissioning.

All three were closed out, and how they were closed is instructive. The false pole detection was fixed at Brantford, working through Apera. The weak column was fixed on site, again jointly with Apera. The third was not fixed at all, because on inspection it did not need to be: rather than teach the system to recognise an empty rack, the cell simply counts parts and stops when it has picked the number it expects.

That last one is worth dwelling on. The instinct in AI-driven automation is to solve every problem with more model. The better answer here was a counter.

The Retry That Made the Difference

The single most interesting piece of engineering in this project is buried in the difference between two robot programs.

The first magnetic pick routine did the obvious thing. If a part was stuck, apply more magnetic force, pull it free, and carry it to the drop position. It worked, and it produced failed drops, because of something that is only obvious in hindsight.

As the robot programmer put it: the camera gives you a reliable position the first time, but once you have interacted with the part, your picking location is no longer the exact same.

A part that needs extra force to break loose does not come free in a predictable orientation. It comes free crooked, held on a magnet that is now gripping it somewhere other than where the vision system said it would. Carrying that part to the drop is carrying a known-bad grip across the cell and hoping.

The revised routine refuses to do that. When a part requires higher magnetism to dislodge, the robot lets go immediately, in the dislodge position, without attempting the drop. It then re-picks the part, which is now loose, from a fresh camera position with a known grip. The count of failed drops fell.

It is a small change in logic and a significant change in philosophy: rather than trying to rescue a compromised pick, throw it away and take a clean one.

Reading the Rack Instead of Assuming It

The other notable development addressed rack warpage directly.

Because the racks distort over time, the hooks holding parts at the front of each column are not where the drawing says. They shift column to column and rack to rack. Rather than treating that as noise to be tolerated, the team trained an additional AI model to find the hooks themselves.

The result is a routine that locates the hook at the front of each column and uses it as a live reference. After a part is picked, it is matched to its hook column, positioned relative to that hook, tilted around it, and merged into a single path for the drop. XY position and depth both come from what the camera actually sees rather than from what the rack is supposed to look like.

A warped rack stops being a source of error and becomes a measurable input.

From Study to Cell

The customer signed the Proof of Principle on 8 July 2025. The purchase order arrived the following day.

From there the project ran at pace. Design approval mid July. Free-issued fencing, slide gate, gate box, carts and dress pack arriving through July and early August. Pre-start health and safety review on 22 August. System acceptance testing and buy-off on 28 August, with the buy-off open issues list closed the same day. The full customer documentation package went out on 10 September.

The customer issued its own post buy-off open issues list on 10 September. Of the eleven items, Ethos closed items one through ten by 16 September. Item eleven fell outside the original scope and was quoted separately rather than absorbed quietly, which is the right way to handle a scope question between two companies that intend to keep working together.

Why This One Matters

The technical achievement is a robot reliably pulling parts out of a rack that was never designed for a robot, guided by a vision model that scores its own confidence and is allowed to be unsure.

The commercial achievement is different and probably more interesting. Ethos took on the risk of proving an unproven application, on its own floor, with the customer's hardware, before there was a contract. The customer took on the risk of handing over a robot and production parts to find out. Both parties got an honest report with the failures counted, and then the job proceeded on a shared understanding of exactly where the system was strong and exactly where it was still thin.

That is a better foundation than a specification that promises everything and discovers the truth at run at rate.

For manufacturers weighing the same question, the path is now well worn: a feasibility study using CAD or sample parts to establish technical fit and cycle time before any hardware is bought, then a pilot cell that proves pick and place under real conditions, then production. Most of the risk gets retired on the integrator's floor, which is where it belongs.

About Ethos Automation

Ethos Automation is a Canadian automation integrator based in Brantford, Ontario. We design, build, program and commission robotic and automated systems for manufacturers in automotive, aerospace, nuclear and general industry, from single cells to complete lines.

info@ethosautomation.com +1 (888) 384 6756 +52 222 286 6121 ethosautomation.com

70 Easton Rd, Brantford, Ontario, Canada N3P 1J5

Carril de San Cristóbal 1704-3, Chachapa, Amozoc, Puebla, Mexico 72990