REEDKNAPP

RESEARCH & ADVISORY

AI & OPERATIONS · ROI BRIEFING

AI in Business: 5 High-ROI Use Cases

A briefing for C-suite leaders in product companies — where measurable value has actually been captured, with figures, named companies, Australian examples, and the lesson behind each.

> **The honest starting point.** Most AI spending does not move the P&L. MIT's 2025 research found that roughly 95% of generative AI pilots delivered no measurable return, and McKinsey reports that while ~88% of organisations now use AI in some capacity, only about 39% can point to enterprise-level earnings impact. Value is concentrated, not typical — and the difference is execution, not the technology. The five use cases below are where the evidence for captured value is strongest for companies that make and move physical products. Each measures a baseline with a direct line to cost or revenue.

1. Predictive maintenance

PRODUCTION highest evidence quality, fastest payback

30–50%	25%	6–12mo
less unplanned downtime	lower maintenance cost	typical payback

Sensors on critical equipment feed models that flag failures before they occur, turning unplanned stoppages into scheduled ones. Siemens reported a 30% cut in maintenance costs and up to 50% less downtime; an automotive deployment monitoring 10,000+ sensors reduced unplanned downtime by 35% at US$2.3M per year. **In Australia, Rio Tinto's AutoHaul** — the world's first fully autonomous heavy-haul rail network in the Pilbara — combines automation with condition monitoring across ~200 locomotives on 1,700+ km of track, increasing network throughput and eliminating 1.5M km of crew road travel per year. ◆ AU EXAMPLE

> LESSON Start where you already meter the baseline. Downtime and maintenance hours are tracked, so before/after attribution is clean, and finance can see the savings.

> LESSON A mid-sized manufacturer fitting smart sensors to its most critical machines saw results inside six months — you don't need a Rio Tinto budget to begin; you need the right ten machines.

2. Computer-vision quality inspection

PRODUCTION defect-escape and scrap map straight to warranty and material cost

99%+	US$2M	100%
defect detection vs ~76% manual	saved/yr (Intel wafer vision)	inspection vs sampling

Cameras and deep-learning models inspect every unit at line speed, catching defects that human inspectors miss in repetitive, high-volume work. A precision-parts manufacturer increased detection from 76% to 99.3% and cut quality-related returns by 91% within six months; Intel's wafer-vision system saves about US$2M a year at a single inspection step. **In Australia, FreshPack Foods** deployed an AI vision system (via Northside Design) that raised defect-detection accuracy to 99.7% from 93.7% for human inspectors, cut inspection time from 15 seconds to under one, and lifted usable production capacity from 78% to 94% of theoretical maximum — while removing recall risk.

> LESSON Quality data is a by-product worth keeping. The same images that pass/fail a unit reveal why defects occur — one manufacturer used them to cut the underlying defect rate a further 17%.

> LESSON Tune for trust. Operators override systems that throw false positives; validate to high recall AND low false-positive rates before full handover, or the line stops believing it.

3. Demand forecasting & inventory optimisation

SUPPLY CHAIN double-digit savings against continuously-measured KPIs

20–40%	~15%	5%
lower inventory cost	lower logistics cost	forecast-error cut (CCA)

Models ingest sales, weather, promotions and events to predict demand at the SKU and store level, trimming both stockouts and excess stock. McKinsey benchmarks AI-driven supply chains at roughly 15% lower logistics costs and 35% lower inventory levels; Australian retail deployments report 20–40% inventory cost reductions. **In Australia, Coca-Cola Amatil** cut its forecasting error by about 5% by modelling past sales, market trends, and seasonality — directly lowering inventory costs and improving availability. Woolworths uses AI for demand prediction, dynamic markdown of perishables, and labour scheduling to align staff with forecast footfall.

> LESSON Feed it local signals. Australian gains came from ingesting region-specific variables — weather (BoM data), heatwaves, sporting events — not generic models. The edge is in the inputs.

> LESSON A few points of forecast accuracy is a large number — acting on the entire inventory base, even a 5% error cut frees material working capital.

4. Personalisation & predictive scoring

MARKETING & SALES controlled-experiment evidence, not just vendor claims

5–15% incremental revenue (personalisation)	20–30% conversion lift (lead scoring)	US$100M Sephora personalisation uplift

Two linked plays: personalising what each customer sees, and scoring leads so finite sales effort targets the best prospects. The credibility comes from controlled tests — JPMorgan Chase found AI-written ad copy roughly doubled click-throughs compared with a human control. Sephora attributed US$100M+ in incremental revenue in a year to personalisation (11% higher conversion among engaged users), and Starbucks' "Deep Brew" tripled offer redemption rates compared with rule-based targeting.

> LESSON Insist on a control group. The defensible numbers come from A/B tests; standalone "+85%" claims without a control are marketing, not evidence.

> LESSON First-party data is the moat. Starbucks and Sephora won on years of loyalty data — start capturing and structuring yours now; the advantage compounds.

5. Generative design & simulation

PRODUCT DEVELOPMENT compresses the prototype cycle into compute

20–55% design-cycle compression	>3× engineering productivity (Synopsys)	up to 25% lower power in chip designs

AI explores vast design spaces and runs simulations that once required physical prototypes. GE used a machine-learning surrogate model to evaluate a million turbine-blade variants in 15 minutes — work that previously took days per iteration. A McKinsey R&D case reported a 55% reduction in design time and 48% less validation effort. Faster iteration both shortens time-to-market (typically 20–40%) and explores more options, lifting design quality.

> LESSON Frame the case around time-to-market and design quality, not headcount removed.

> LESSON Keep the engineer in the loop. AI proposes and optimises candidates; experts judge.

THE BOTTOM LINE

The winners are not those spending the most on AI, but those starting with **bounded, measurable use cases** where a clear baseline meets a direct cost or revenue line — then redesigning the workflow around it.

Selected sources: State of AI 2025; MIT Project NANDA, The GenAI Divide 2025; Gartner. McKinsey, Company examples: Rio Tinto/Hitachi Rail, FreshPack Foods/Northside Design, Coca-Cola Amatil, Woolworths, Siemens, Intel, JPMorgan Chase/Persado, Sephora, Starbucks, GE, Synopsys.

Figures are drawn from company disclosures, consulting research, and vendor case studies; vendor figures self-select for success — treat them as achievable upper bounds and verify them against primary sources before relying on them.

ABOUT THIS ARTICLE

Researched and written by an AI assistant working from live web research across company disclosures, consulting studies, academic papers and vendor case studies. A human editor set the brief, reviewed the output, and is responsible for final fact-checking. This is part of a recurring series; the prompt used to produce it is shown below so readers can see and reproduce the process.

THE PROMPT

“From a shortlist of five high-ROI AI use cases for product companies — predictive maintenance, computer-vision quality inspection, demand forecasting and inventory optimisation, personalisation and predictive lead scoring, and generative design and simulation — build a concise, objective briefing for mid-market C-suite executives. Research each use case for depth; prioritise Australian examples; name specific companies; give specific quantifiable results and the lesson learned. Keep it factual, credit sources, and flag vendor figures as achievable upper bounds to be verified before publication.”