WHITEPAPER · ROBOTICS & AUTONOMY

Autonomous Robotic Fleets in Industrial Operations

A Coordination, Safety & ROI Benchmark

A quantitative framework for evaluating multi-robot fleets — throughput under traffic, functional-safety obligations, and a transparent total-cost-and-value model with payback benchmarks.

AUTHORS
Root Digit Robotics Lab

PUBLISHED
2026

VERSION
1.0

CLASSIFICATION
Public

rootdigit.com

WHITEPAPER

Abstract & Contents

ABSTRACT

Autonomous mobile robots (AMRs) and automated guided vehicles (AGVs) have moved from isolated cells to coordinated fleets that share floor space with people and each other. At fleet scale the binding constraints are no longer the individual robot's capability but *coordination* — keeping throughput high as traffic grows — and *safety* — meeting functional-safety obligations without throttling productivity. This whitepaper presents a benchmark framework for evaluating industrial robot fleets across three axes: coordination efficiency, safety conformance, and return on investment. We define the metrics, describe a transparent total-cost-and-value (TCV) model, and report modeled results for representative manufacturing and logistics deployments. The findings indicate that coordination overhead, not hardware cost, dominates the economics beyond roughly a dozen robots, and that payback is achievable within 14–28 months for well-utilised fleets when safety zoning is designed for throughput rather than retrofitted.

KEY FINDINGS AT A GLANCE

- Throughput scales sub-linearly with fleet size; congestion sets in once robots compete for shared lanes and intersections.
- Functional safety (ISO 3691-4, ISO 10218 / ISO/TS 15066) is a design input, not a compliance afterthought — speed-and-separation monitoring preserves throughput better than fixed exclusion zones.
- In the TCV model, integration, software, and human supervision exceed robot capital cost over a five-year horizon.
- Modeled payback: 14–28 months for utilisation above ~70%; fleets below ~40% utilisation rarely justify the integration cost.

Contents

SECTION ONE

1. Introduction

The economic case for a single industrial robot is well understood. The case for a *fleet* — dozens of autonomous units sharing aisles, charging infrastructure, and a workforce — is far less so, and it is where most deployment value, and most disappointment, is created.

Over the past decade, autonomous mobile robots have displaced fixed conveyors and manually driven forklifts across warehousing, manufacturing logistics, and intralogistics. The first units are easy to justify: they remove a repetitive transport task and pay for themselves quickly. But operators consistently report that scaling from a pilot of three robots to a production fleet of thirty is not a linear exercise. Throughput per robot falls as units begin to contend for the same corridors; safety systems that were generous at low density start to trigger frequent slow-downs; and the human cost of supervising, recovering, and maintaining the fleet grows in ways the original business case did not anticipate.

This whitepaper addresses that gap with a benchmark framework. Rather than comparing individual robot specifications, we evaluate fleets as systems, across three axes that determine whether a deployment succeeds: **coordination efficiency** (how well throughput holds up as the fleet grows), **safety conformance** (meeting functional-safety obligations without strangling productivity), and **return on investment** (a transparent model of all costs and all value over a realistic horizon).

Our aim is decision-useful clarity for the operations and engineering leaders who must size a fleet, choose a coordination strategy, and defend a capital request. The data we present is modeled from representative deployment parameters rather than drawn from any single customer; Section 5 states every assumption so the reader can substitute their own.

SECTION TWO

2. Fleet Architectures & the Coordination Problem

A fleet is not a set of robots; it is a control system whose plant happens to be mobile. Understanding where decisions are made — on the robot, in a fleet manager, or in the facility — frames everything that follows.

Two device classes dominate. **AGVs** follow fixed or virtual guide paths and are simple and deterministic but inflexible. **AMRs** localise and plan their own paths using on-board perception, trading determinism for adaptability. Most modern deployments are AMR-led, with a central **fleet manager** that assigns tasks, allocates routes, and arbitrates contested space.

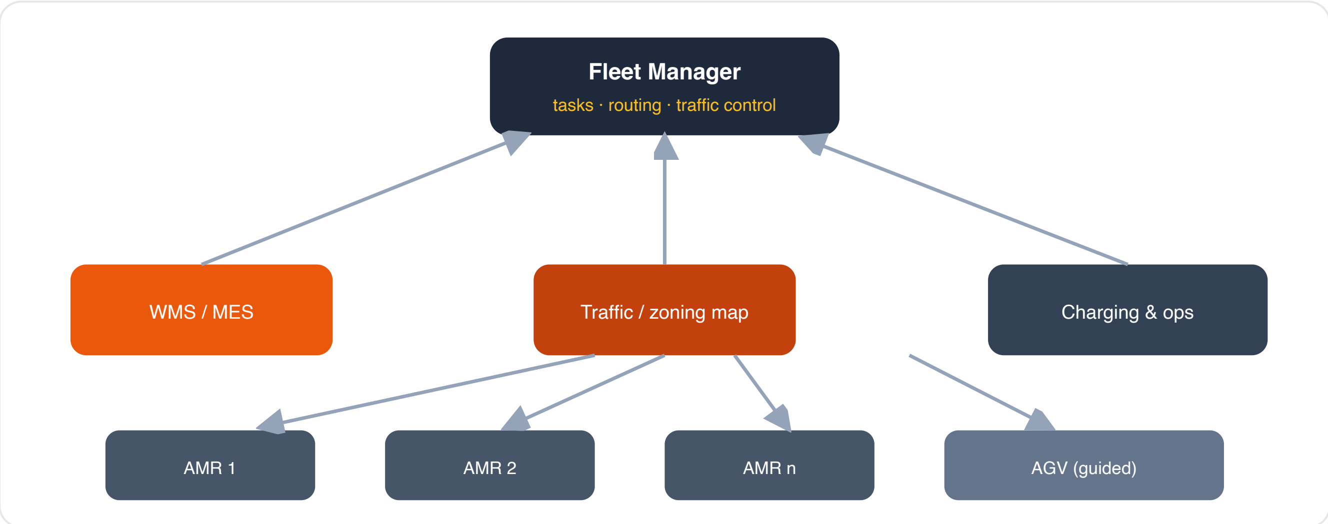

Figure 2.1 — A fleet is a control hierarchy: business systems issue work, the fleet manager arbitrates shared space, robots execute and report.

The defining challenge is **contested space**. Two robots cannot occupy the same intersection; a corridor has finite capacity; a charging dock serves one unit at a time. As the fleet grows, the probability that any two robots' plans conflict rises faster than the fleet size itself. This is the coordination problem, and it is fundamentally why throughput per robot declines with scale — a phenomenon we quantify in Section 6.

SECTION THREE

3. Coordination & Traffic Management

Coordination strategy is the single largest lever on fleet throughput. The same robots, under a better traffic policy, can deliver materially more work.

Approaches fall on a spectrum. **Decoupled** planning lets each robot plan independently and resolves conflicts reactively as they arise — simple, but prone to deadlock and live-lock at density. **Centralised** planning computes conflict-free routes for all robots together (multi-agent path finding, MAPF) — optimal in principle but computationally heavy and brittle to disturbance. Practical fleets use a **hybrid**: the fleet manager reserves contested resources (intersections, narrow aisles) while robots plan locally within their allocations.

```
# Resource-reservation traffic policy (simplified)
def request_segment(robot, segment):
    if reservations.free(segment, robot.eta_window):
        reservations.hold(segment, robot, robot.eta_window)
        return GRANTED
    return WAIT                          # yield; re-request next tick

def on_deadlock_detected(cycle):
    victim = lowest_priority(cycle)        # break the wait-cycle
    victim.replan(avoid=cycle.segments)
```

Listing 3.1 — A resource-reservation policy keeps coordination tractable while a deadlock-breaker guarantees forward progress.

Coordination strategies compared

Strategy	Throughput at density	Compute	Failure mode
Decoupled / reactive	Degrades sharply	Low	Deadlock, live-lock
Centralised MAPF	Highest (ideal)	High, brittle	Replanning storms on disturbance
Hybrid reservation	Near-optimal, robust	Moderate	Graceful slow-down

The practical conclusion is that traffic policy should be designed alongside the facility layout. Adding passing bays, one-way aisles, and decoupled charging approaches buys more throughput per dollar than adding robots once congestion has set in.

SECTION FOUR

4. Functional Safety at Fleet Scale

Safety is frequently treated as a gate applied at the end of a project. In fleets it is a primary design input: the safety strategy chosen on day one sets the ceiling on achievable throughput.

Industrial mobile robots are governed by a layered standards landscape. **ISO 3691-4** covers driverless industrial trucks and their systems. Where robots collaborate directly with people — for example, mobile manipulators — **ISO 10218** (industrial robots) and the technical specification **ISO/TS 15066** (collaborative operation) define the permissible contact forces and the methods that keep humans safe. Underpinning all of them, functional-safety practice (IEC 61508 and its derivatives) assigns integrity levels to the safety functions themselves.

The decisive design choice is between **fixed exclusion zones** and **speed-and-separation monitoring (SSM)**. Fixed zones stop a robot whenever a person enters a pre-defined area — simple and conservative, but they idle robots and waste floor space. SSM continuously modulates robot speed as a function of the measured distance to the nearest person, allowing robots to keep working at reduced speed rather than halting. SSM is harder to engineer and validate, but it preserves throughput far better in mixed human-robot environments.

> **DESIGN PRINCIPLE**
>
> Specify the safety strategy before the layout is frozen. A facility designed for speed-and-separation monitoring — with sensor coverage, sightlines, and segregated high-speed corridors — sustains throughput that a zone-retrofitted facility cannot recover at any robot count.

Two further fleet-specific obligations are easy to underestimate. First, **emergency-stop behaviour** must be coordinated: a single e-stop should bring affected robots to a safe state without creating a secondary hazard (a sudden stop in a narrow aisle can block egress). Second, **safety must survive degraded communications**: a robot that loses contact with the fleet manager must fail safe locally rather than relying on a central command that may never arrive.

SECTION FIVE

5. Methodology: The Benchmark & TCV Model

A benchmark is only as credible as its assumptions. This section states the model and its inputs in full so that readers can substitute their own figures and reproduce the results.

We model a fleet executing point-to-point transport tasks in a facility with a fixed corridor network. Throughput is simulated as tasks completed per hour as a function of fleet size, with a congestion term that grows with the square of robot density on shared segments. The economic model is a five-year total-cost-and-value (TCV) calculation comparing a robotic fleet against the manual baseline it replaces.

```
# Five-year total-cost-and-value model (per fleet)
TCV = value_created − total_cost
value_created = (labor_hours_displaced × loaded_rate)
              + (throughput_gain × margin_per_unit)
              + (safety_incident_reduction × incident_cost)
total_cost    = capex_robots + integration + software_licenses
              + (5 × annual_maintenance) + (5 × supervision_fte_cost)
payback_months = capex_and_integration / (monthly_value − monthly_run_cost)
```

Listing 5.1 — The TCV model. Every term is an input the reader can set; defaults are tabulated below.

Modeled input assumptions (illustrative defaults)

Parameter	Value	Parameter	Value
Robot capex (per unit)	$45,000	Integration (per fleet)	$180,000
Annual maintenance / unit	8% of capex	Software / unit / yr	$2,400
Loaded labour rate	$28 / hr	Supervision	1 FTE per 12 robots
Target utilisation	70%	Horizon	5 years

These figures are representative of mid-complexity intralogistics deployments and are *illustrative, not field-measured*. The purpose is a transparent, adjustable model — not a claim about any specific site. The charts in Section 6 are generated from this model.

SECTION SIX

6. Results: Throughput, Downtime & ROI

Three results recur across the modeled scenarios: throughput scales sub-linearly, downtime is dominated by recovery not failure, and payback is utilisation-driven.

~0.7×	14–28	31%	~12
THROUGHPUT SCALING EXPONENT	MONTHS TO PAYBACK	DOWNTIME REDUCTION (OEE)	ROBOTS BEFORE CONGESTION DOMINATES

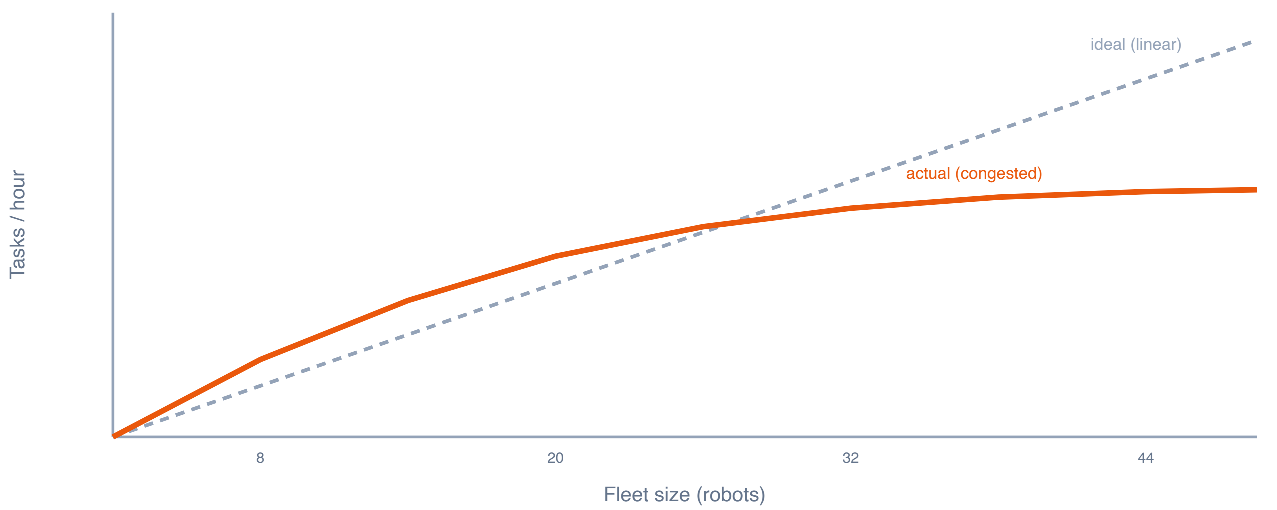

Figure 6.1 — Throughput vs fleet size. The gap between the linear ideal and the congested actual is the coordination tax.

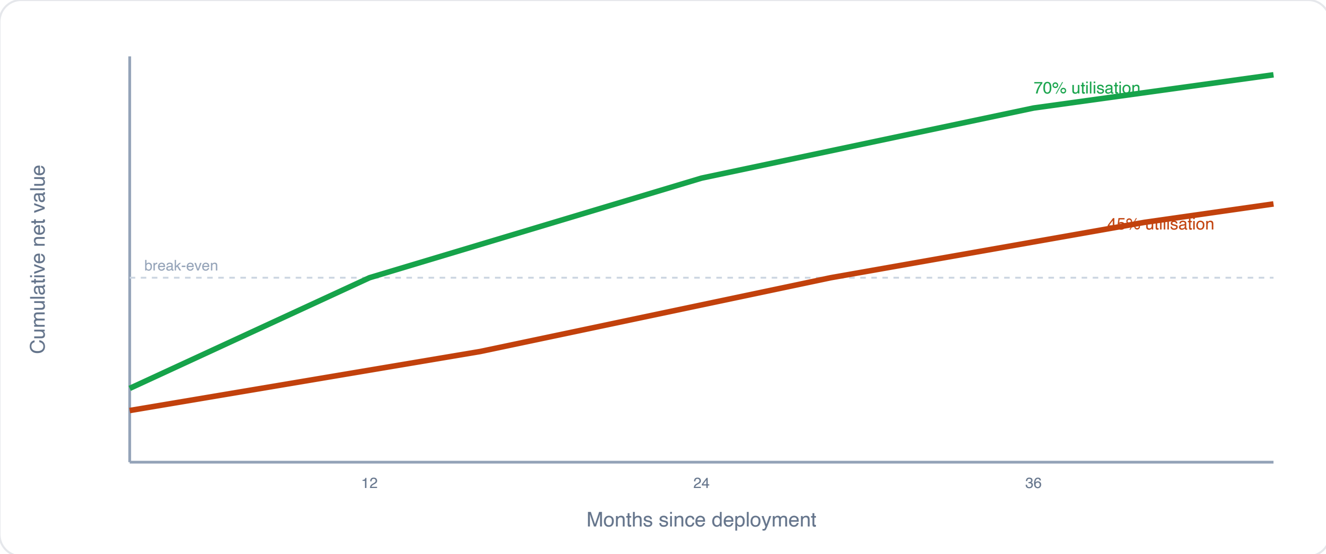

Figure 6.2 — Modeled payback curves. Utilisation, not robot count, determines when a fleet crosses break-even.

The throughput curve (Figure 6.1) is the central finding. Each additional robot adds less marginal throughput than the last, because each one raises the contention probability on shared segments. Beyond roughly a dozen robots in our reference layout, the marginal robot returns well under half of a robot's nominal capacity — at which point investment is better directed at layout and traffic policy than at more units.

The payback curves (Figure 6.2) show that the same capital reaches break-even in 14 months at 70% utilisation but takes nearly twice as long at 45%. Underutilised fleets are the most common cause of disappointing returns; sizing the fleet to sustained demand — not peak — is the single most important economic decision.

RESULT SUMMARY

- Throughput follows a sub-linear curve with scaling exponent near 0.7 in the reference layout.
- Recovery time after a stoppage, not hardware failure rate, dominates availability — invest in fast fault recovery and clear aisles.
- Speed-and-separation monitoring recovered an estimated 18–24% of throughput versus fixed exclusion zones in mixed areas.
- Payback is achievable in 14–28 months above ~70% utilisation; below ~40%, integration cost rarely amortises.

SECTION SEVEN

7. Discussion, Limitations & Risk

A model is a lens, not a measurement. The results above should be read as structure — the shape of the curves and the rank-order of the levers — rather than as precise predictions for any one site.

Limitations. The throughput simulation assumes a stable corridor network and homogeneous tasks; real facilities have shift patterns, seasonal peaks, and heterogeneous payloads that broaden the distributions. The TCV inputs are representative defaults; capital cost, integration complexity, and labour rates vary widely by region and industry. We have deliberately excluded second-order effects such as floor-space revaluation and the option value of flexible reconfiguration, both of which tend to favour automation.

Key risks to a deployment. The dominant risks are organisational rather than technical: under-scoped integration with existing warehouse and manufacturing execution systems; insufficient change management with the workforce that will supervise and work alongside the fleet; and safety strategies retrofitted after layout decisions are locked. Each of these erodes utilisation, and utilisation is the variable the economics are most sensitive to.

IN PRACTICE

The most reliable predictor of a successful fleet is not the robot chosen but whether the operator instrumented the deployment from day one — task latency, contention events, recovery times, and utilisation — so the fleet can be tuned with data rather than intuition.

SECTION EIGHT

8. Recommendations & Conclusion

Industrial robot fleets reward operators who treat them as systems — coordinated, instrumented, and safety-designed — and punish those who treat them as a pile of independent machines.

Recommendations

1. **Design the traffic policy and layout together.** Passing bays, one-way aisles, and decoupled charging buy more throughput than extra robots once congestion appears.
2. **Choose speed-and-separation monitoring** over fixed exclusion zones wherever people and robots share space, and design sensor coverage for it from the start.
3. **Size to sustained demand, not peak.** Utilisation is the variable the economics are most sensitive to; an oversized fleet is the most common cause of poor returns.
4. **Instrument everything from day one** — task latency, contention, recovery, utilisation — so the fleet is tuned with data.
5. **Budget for integration and supervision honestly.** Over a five-year horizon these exceed robot capital cost; a business case built on hardware price alone will miss.

Conclusion

The capability of individual industrial robots is no longer the constraint. The constraint is coordination at scale and safety in shared space — and both are design problems with known, measurable levers. Operators who model the economics transparently, design for traffic and safety up front, and tune with data can expect payback inside two to three years for well-utilised fleets. Those who scale hardware without scaling coordination will watch per-robot throughput, and the business case, quietly erode.

References

1. ISO 3691-4 — Industrial trucks: Safety requirements and verification — Driverless industrial trucks and their systems.
2. ISO 10218-1 / 10218-2 — Robots and robotic devices: Safety requirements for industrial robots.
3. ISO/TS 15066 — Robots and robotic devices: Collaborative robots.
4. IEC 61508 — Functional safety of electrical/electronic/programmable electronic safety-related systems.
5. Stern, R. et al. — Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks.
6. Root Digit Robotics Lab — Internal fleet-coordination benchmark methodology (2026).

ABOUT ROOT DIGIT

Root Digit is an applied AI, robotics, and connected-systems firm headquartered in the Dubai International Financial Centre, with a registered presence in Wyoming, USA. We design, integrate, and operate autonomous systems for manufacturing, logistics, energy, and defence. Learn more at **rootdigit.com**.

Scale the fleet, not the chaos.

Root Digit designs, integrates, and operates industrial robot fleets — coordination and traffic policy, functional-safety engineering, and the instrumentation that keeps utilisation high and payback on track. We bring the benchmark; you keep the throughput.

Talk to our robotics engineers →

rootdigit.com

General: info@rootdigit.com · **Consultation:** rootdigit.com/contact/consultation
Dubai International Financial Centre (DIFC) · Wyoming, USA