



Beyond OTIF: Five KPIs For Guiding Supply Chain Success by 2030

A Strategic Guide to Measuring What
Matters in AI-Driven Logistics



FOREWORD

This guide serves as **your launchpad** for reimagining performance measurement in an AI-driven landscape. Consider these KPIs and frameworks as starting points for innovation rather than rigid prescriptions. Your success will be shaped by the quality of your data foundation and the strategic insights you bring to their application.

As you embark on this journey of analytical transformation, we encourage you to customize these recommendations to your organization's context. The future of supply chain excellence is yours to define—**let this guide inspire your path forward.**

A Note on Implementation

These KPIs represent directional thinking for the future of supply chain measurement—ways to capture value that traditional metrics miss. They're designed to spark conversation about what becomes possible as AI transforms logistics operations.

We recognize that measuring these metrics requires robust data infrastructure and integration across systems. While these aren't capabilities that platforms natively report today, forward-thinking organizations with mature BI environments can begin exploring these concepts using their existing data.

Consider this guide a starting point for strategic discussions rather than an implementation manual. The goal is to evolve how we think about supply chain performance in an AI-driven world.



The Measurement Gap

Leaders aren't short on technology budgets—92 percent of companies say they will boost AI investment over the next three years—yet only one percent feel they have scaled it ([McKinsey & Company](#)). The gap isn't funding; it's measurement. We still point dashboards at yesterday's problems while AI quietly averts tomorrow's. If we can't see that value, we can't defend the business case, and the board will wonder whether AI is just another shiny object.

I've spent three decades watching supply chain metrics evolve from ledger lines to real-time control towers. The firms that outperformed in every era shared one habit: they measured what their competitors ignored—until it was too late to catch up. The next era is here. Let's start measuring it.

The Metric That No Longer Matters

I still love an on-time-delivery chart; it kept many of us employed. But let's be honest: when an AI agent auto-rebooks a vessel, re-slots a dock, and notifies the customer hours before anyone panics, the OTIF column registers "green" and the story ends there. The entire rescue—forecast, decision, execution—never shows up.

Traditional KPIs are rear-view mirrors. They tell you the truck was late, the inventory piled up, or the cost per mile spiked after you've already lost money, margin, or market share. In a world where machines correct course in milliseconds, lagging indicators become a comfort blanket: familiar, reassuring, and dangerously obsolete.

So let's retire the idea that a metric has proved its worth simply because it is precise. Precision about the past is not insight; it is nostalgia. What matters now is visibility into value created before it hits the P&L—the revenue you kept, the carbon you avoided, the manual hours you redeployed.

In the pages ahead, you'll find five north-star metrics built for that reality. Adopt them, adapt them, and above all make them your own. Because the next competitive edge won't come from squeezing another basis point out of cost per mile; it will come from proving, every single day, that your supply chain can see around corners—and act before anyone else even knows there was a bend.

Dov Shenkman

CEO and Founding Member
Atid Group

TABLE OF CONTENTS

-
- 1. Revenue-at-Risk Mitigation (RRM):** Dollars that were headed toward late-delivery write-offs but got rescued in-flight—tracking how much revenue the network “saved” before customers noticed
 - 2. Predictive OTIF Accuracy:** The gap between what the AI said would be On-Time-In-Full at order creation and what actually hit the dock—showing how reliably you can promise service days or weeks in advance
 - 3. Mean Time to Recovery (MTTR):** Average hours it takes the supply chain to get back on plan after a disruption—your real-world measure of resilience speed
 - 4. Autonomous Decision Rate (ADR):** Percent of exceptions the system resolves end-to-end without human keystrokes—proof of how far you've shifted planners from firefighting to higher-value work.
 - 5. Network Fuel-and-Carbon Efficiency Index (NFCEI):**
The percentage drop in fuel burn and CO2 per ton-mile versus historical norms—tying logistics cost cuts and Scope 3 gains into one ESG-ready number

Revenue-at-Risk Mitigation (RRM)

TURNING NEAR-MISSES INTO COMPETITIVE ADVANTAGE

Every supply chain professional knows the sinking feeling: a critical shipment is trending late, penalties are looming, and a key customer relationship hangs in the balance. Traditional metrics capture these failures after they happen. Revenue-at-Risk Mitigation captures something far more valuable—the failures that never occurred.

RRM quantifies dollars that were heading toward late-delivery penalties, lost sales, or customer defection but were rescued through proactive intervention. It's not about celebrating problems you fixed; it's about showcasing problems your customers never experienced.

The Calculation That Changes Everything

$$\text{RRM} = \Sigma(\text{Order Risk} \times \text{Recovery Rate} \times \text{Blended Penalty Percentage})$$

Imagine, in one quarter, your business identified \$8M of orders destined to customers that were at risk of an OTIF penalty. You identified that 85% of those orders were about to avoid their OTIF penalty (a blended average of 7% across all suppliers), meaning that you avoided \$476,000 in OTIF penalties. That's \$476,000 in protected revenue—invisible to traditional KPIs but very real to your CFO.



P&L Impact: From Cost Center to Revenue Protector

The financial implications of RRM extend far beyond avoided penalties. When you prevent a late delivery to a key retailer, you’re protecting:

- **Immediate revenue:** Contract penalties typically range from 1-5% but can reach 10% for critical service agreements
- **Future business:** A single service failure can trigger competitive RFPs, risking millions in annual contracts
- **Pricing power:** Consistent reliability commands premium rates—typically 3-7% above market
- **Working capital:** Penalty disputes tie up payments for 30-90 days, impacting cash flow

Organizations tracking RRM report shifting executive conversations from “How much did we lose?” to “How much did we save?”—a subtle but powerful change in organizational mindset.

Starting Your RRM Journey

While fully automated RRM tracking remains aspirational, organizations can begin with simpler measurements:


- Track exceptions handled by AI agents versus those requiring human intervention
- Monitor shipments flagged as “at-risk” and their final outcomes
- Calculate potential penalties on your highest-value lanes as a baseline

The key is establishing a baseline understanding of your risk exposure, even if attribution remains imperfect initially.

HOW IT COULD LOOK

Revenue-at-Risk Waterfall

Shows, month by month, how many dollars were at risk of late-delivery penalties, what the network saved, and what still leaked out.



Predictive OTIF Accuracy

THE NEW CURRENCY OF TRUST


On-time-in-full delivery has long been the gold standard of supply chains. But there's a problem: traditional OTIF is a lagging indicator, telling you last month's performance when your customer needed certainty at order placement. Predictive OTIF Accuracy can change the game entirely.

This KPI measures the reliability of your AI's delivery predictions made at order creation versus actual performance. When you can tell a customer with 95% certainty that their order will arrive Tuesday at 2 PM—three weeks from now—and consistently deliver on that promise, you're not just providing a service. You're providing certainty in an uncertain world.

The Formula for Customer Confidence

Predictive OTIF Accuracy = (Correctly Predicted Outcomes / Total Predictions) × 100

When adopting new AI tools, it's crucial to assess their effectiveness and understand acceptable variability. Since AI tools won't achieve 100% accuracy due to inherent variability, it's important to benchmark your current On-Time In-Full (OTIF) exception resolution rate. This allows you to determine if your definition of success aligns with your human intervention rate or if you anticipate a higher level of performance, which is key to gauging program success.



P&L Impact: Certainty Commands Premium Prices

The financial benefits of predictive accuracy compound across multiple areas:

- **Revenue enhancement:** B2B customers pay premiums for guaranteed delivery windows, especially in JIT manufacturing environments
- **Inventory optimization:** When you can predict delivery within hours, customers reduce safety stock, creating stickier relationships
- **Cost reduction:** Accurate predictions eliminate emergency expediting, which typically costs 3-5x standard shipping rates
- **Capital efficiency:** Better predictions mean less inventory in the pipeline, freeing working capital

The real value comes from what customers do with this certainty—synchronizing production schedules, optimizing labor, and making commitments to their own customers based on your reliability.

Building Toward Predictive Accuracy

Start with what you can measure today:


- Compare initial ETAs at booking versus actual delivery times
- Track how often AI-suggested arrival times prove accurate
- Monitor variance between promised and actual delivery windows

As your data maturity grows, you can evolve from tracking simple on-time percentage to measuring the reliability of predictions made days or weeks in advance.

HOW IT COULD LOOK

Predictive OTIF Stacked Area

Illustrates weekly baseline OTIF performance and the extra lift delivered by AI—stacked so you can quickly see AI’s impact on your OTIF performance.



Mean Time to Recovery (MTTR)

RESILIENCE AT THE SPEED OF BUSINESS


Every supply chain fails. What separates leaders from laggards is how fast they recover. Mean Time to Recovery measures the average hours between disruption detection and return to normal operations. With the increasing frequency of supply chain disruptions, MTTR has evolved from an IT metric to a boardroom imperative.

Today's supply chains demand speed, as every hour of disruption can lead to days of downstream impact. While traditional metrics focus on disruption frequency or impact, MTTR prioritizes speed. AI plays a crucial role in identifying potential disruptions early, enabling businesses to implement proactive mitigation strategies or leverage AI agents for preemptive resolution.

Calculating the Cost of Time

$$\text{MTTR} = \Sigma(\text{Recovery Time per Incident}) / \text{Number of Incidents}$$

But the real insight comes from segmenting this calculation. A weather disruption might recover in 6 hours through standard rerouting. A supplier failure requiring alternative sourcing might take 48 hours. Understanding these patterns transforms MTTR from a metric to a roadmap for investment.



P&L Impact: Where Hours Equal Dollars

The financial impact of MTTR improvement is both immediate and compounding:

- **Direct costs:** Each hour of disruption typically triggers expedited shipping (at a potentially significant markup), overtime labor, and emergency sourcing premiums
- **Inventory impact:** Disruptions may force higher safety stock levels across the network, tying up working capital
- **Customer penalties:** Beyond contractual penalties, service failures impact long-term revenue through lost business
- **Market position:** Companies with faster recovery times report higher customer retention rates

Organizations can use advanced visibility and AI-powered predictive analytics to help anticipate disruptions and trigger proactive responses to slash recovery times. Real-time networks combined with machine learning models that recognize disruption patterns can enable faster mobilization of alternative routes, carriers, and inventory. The difference between 48-hour and 12-hour recovery can mean millions in protected revenue.

Making MTTR Actionable

Many organizations already track elements of MTTR:


- Time stamps when exceptions are identified
- Resolution time for different disruption types
- Escalation rates from automated to manual handling

The evolution is connecting these data points to understand your true recovery velocity and identifying which disruption types benefit most from automation.

HOW IT COULD LOOK

MTTR Violin Plot Area

Visualizes the spread of recovery times each month; thinner and shorter violins mean disruptions are being fixed faster and more consistently.



Autonomous Decision Rate (ADR)

THE METRIC OF SCALE

Here's an uncomfortable truth: your best supply chain planners spend most of their time on repetitive tasks a well-trained algorithm could handle. Autonomous Decision Rate measures the percentage of supply chain exceptions resolved without human intervention. The goal isn't to replace people, it's to focus them on higher-impact tasks that improve yield and productivity.

Research indicates that supply chain professionals spend a significant portion of their time handling repetitive exceptions. ADR tracks your progress in flipping this ratio, automating the routine so your team can focus on relationships, innovation, and strategy.

The Automation Equation

$$\text{ADR} = (\text{Autonomous Resolutions} / \text{Total Exceptions}) \times 100$$

But raw automation isn't the goal. The magic happens when you segment ADR by exception type and value. For example, you might automate 90% of routing exceptions but only 30% of supplier substitutions. This granularity reveals where human judgment adds the most value and where automation can scale your capabilities.



P&L Impact: Scaling Without Growing

The financial case for ADR improvement touches every line of the P&L:

- **Labor leverage:** Higher ADR frees planner capacity for strategic work, creating significant value without adding headcount
- **Decision speed:** Automated decisions happen in milliseconds versus hours, reducing cascade effects
- **Consistency:** Automated decisions follow optimal logic every time, eliminating exceptions caused by human error or inconsistency
- **Scalability:** High ADR operations handle significant transaction volume growth without proportional headcount increases

Some companies achieving high ADR report transformation beyond cost savings—their planners become strategic advisors rather than firefighters.

Practical ADR Measurement

Begin with straightforward metrics:


- Percentage of exceptions resolved without human touch
- Number of automated actions taken daily
- Time saved through automated resolution

Even simple ratios like “exceptions handled by AI” versus “total exceptions” provide valuable insights into your automation progress.

HOW IT COULD LOOK

ADR Stacked Area

Tracks exception workload: how much was resolved autonomously vs. by humans or escalation. The larger the gold layer, the higher the true ADR.



Network Fuel-and-Carbon Efficiency Index (NFCEI)

WHERE GREEN MEETS GOLD

Environmental metrics used to live in ESG reports. Now they drive operational decisions and investor valuations. The Network Fuel-and-Carbon Efficiency Index combines fuel consumption and carbon emissions into a single metric that speaks to both CFOs and chief sustainability officers.

NFCEI measures percentage improvement in fuel burn and CO₂e per ton-mile versus baseline. It represents the evolution toward more sustainable and intelligent supply chain ecosystems. More practically, it's where environmental responsibility stops being a cost and starts being a competitive advantage.

The Efficiency Equation

$$\text{NFCEI} = [\text{Baseline CCI} - \text{Current CCI}] / \text{Baseline Carbon Adjusted Cost Intensity} \times 100$$

Where Carbon Adjusted Cost Intensity (CCI) = $[\text{Fuel Spend} + \text{Carbon Cost}] / \text{Ton-Miles}$

Carbon Cost is a unit your business assigns to CO₂e, typically ranging from \$10/T for early-stage adopters and up to \$70-\$150+/T



Think of this as a way to blend cost and carbon into one harmonized equation.

For example:

Input	Baseline FY 23	Current Q2 FY 25
Fuel consumed (diesel)	1,000,000 gal	850,000 gal
Pump price	\$3.20 / gal	\$3.70 / gal
Fuel Spend	\$3,200,000	\$3,145,000
CO ₂ e emitted*	10,000	8,500 t
Carbon price (internal)	\$50 / t	\$50 / t
Carbon Cost	\$500,000	\$425,000
Ton-miles moved	40 million	38 million

*Diesel factor ≈ 10 kg CO₂e per gallon.

Then we quantify the change in the CCI:

- **FY23:** $(\$3,200,000 + \$500,000) / 40,000,000 = \$0.0925/\text{ton-mi}$
- **FY25:** $(\$3,145,000 + \$425,000) / 38,000,000 = \$0.0939/\text{ton-mi}$
- **Change:** $[(\$0.0939 - \$0.0925) / \$0.0925] \times 100 = -1.5\%$

Despite a 15% reduction in CO₂e emissions, the blend of increased ton-miles and higher pump prices resulted in a 1.5% rise in cost-and-carbon intensity per ton-mile compared to FY23. This is reflected in the NFCEI of -1.5%, indicating that emission reductions have not yet led to overall cost savings.

The power of NFCEI lies in its inclusiveness. Every optimization—route efficiency, load consolidation, mode shifting—improves the metric. This creates organization-wide alignment around a single number that benefits both planet and profit.

P&L Impact: Sustainability as Strategy

NFCEI improvement delivers insights across multiple financial dimensions:

- **Direct savings:** Fuel represents a significant portion of transportation spend; improvements directly impact the bottom line
- **Carbon economics:** With rising carbon prices and regulations, emission reductions translate to real financial value
- **Customer requirements:** Increasing numbers of RFPs include sustainability criteria; strong NFCEI wins business
- **Investment access:** ESG-focused investment continues to grow, with emissions performance becoming a key screening factor

Leading companies treat NFCEI as a strategic differentiator, not a compliance checkbox.



Simplified Sustainability Tracking

While the full NFCEI calculation may seem complex, start with basics:


- Fuel spend per mile or per shipment
- Miles per gallon improvements over time
- Basic CO2e calculations using standard emission factors

Many organizations already track these elements for ESG reporting—the innovation is connecting them to operational decisions and cost optimization.

HOW IT COULD LOOK

CCI + NFCEI Combo

Gold bars show the carbon-adjusted cost per ton-mile; the blue line tracks percent improvement vs. January, turning fuel + carbon savings into one easy number.



Achieving Autonomous Excellence

Companies implementing these five north-star KPIs don't just improve their metrics—they transform their market position. The shift from reactive to predictive, from cost-focused to value-creating, from efficiency to resilience creates competitive moats that compound over time.

Organizations embracing this shift report remarkable transformations:

- **Strategic elevation:** Supply chain leaders get invited to revenue discussions, not just cost reviews
- **Talent magnetism:** Top performers seek employers offering strategic impact, not tactical execution
- **Innovation acceleration:** Teams freed from firefighting generate more improvement initiatives
- **Valuation premiums:** Markets increasingly price in supply chain capabilities as competitive differentiators

These outcomes don't require waiting for perfect data or technology—they begin with a decision to expand your measurement mindset.

From Theory to Practice: Your Next Steps

These KPIs intentionally push beyond today's standard metrics. They're designed to help you envision how supply chain measurement must evolve as AI capabilities expand.

Start Where You Are:

- Identify which KPI addresses your most pressing challenge
- Assess your current data capabilities and gaps
- Begin with proxy metrics that directionally indicate progress
- Use these concepts to guide technology investments and data strategy
- Build confidence with quick wins, then scale

Remember: The goal isn't perfect measurement today—it's preparing for the metrics that will matter tomorrow. As AI agents become more sophisticated and data integration improves, these theoretical KPIs will become practical necessities.



THE FOURKITES + FEDEX DATAWORKS ADVANTAGE

Two foundational capabilities are essential for the effectiveness of these north-star KPIs: comprehensive visibility and predictive intelligence. FourKites and FedEx Dataworks offer both, establishing an ecosystem where advanced metrics transition from theoretical concepts to practical applications.

FourKites provides real-time execution across all modes, carriers, and geographies. Their platform monitors millions of shipments daily, generating the data density required for precise predictions and empowering a Digital Workforce to act upon them. FedEx Dataworks contributes the analytical power, leveraging AI models trained on decades of logistics data, pattern recognition across billions of shipments, and optimization algorithms that convert insights into actionable steps.

Initiating this transformation doesn't necessitate a complete overhaul. Organizations can commence with focused pilot programs, demonstrate value in specific lanes or segments, and subsequently scale based on proven outcomes. The crucial step is to begin—because while these KPIs signify the future of supply chain management, that future is materializing more rapidly than many organizations anticipate.