THE HOLISTYX GROUP

SPARKle™ Express Report

AI Transformation Roadmap

SPARKLE by THG

PREPARED FOR

Ford Motor Company

PRIMARY CONTACT

Mike Sternberg, CEO

INDUSTRY

Automotive — OEM (Vehicle Manufacturer)

HEADQUARTERS

Detroit, MI, United States

FOR YOUR LEADERSHIP

Ford Motor Company can turn scattered data and fast-moving competitors into a measurable productivity and cost advantage through five prioritized AI programs.

Competitors including General Motors, Toyota, and Volkswagen Group are scaling AI while tariff costs and quality exposure squeeze margin. Our five ranked use cases represent a directional three-year value of $300M to $564M against an estimated investment of $74M to $117M. The first three, led by the Engineering and Enterprise Knowledge Copilot, address your stated 5% productivity goal and your warranty and operating cost pressure directly.

PORTFOLIO VALUE (3 YEARS)	ESTIMATED INVESTMENT	AI READINESS
$300M to $564M	**$74M to $117M**	**3 / 5**
Top 5 use cases, directional	Across all five use cases	Developing

THE ASK

Approve a 90-day foundation program with a named executive owner, a productivity baseline, and a data scope covering engineering, warranty, and quality sources.

TOP THREE OPPORTUNITIES

1. Engineering and Enterprise Knowledge Copilot
2. Predictive Quality and Warranty Analytics
3. Warranty Claims Processing Automation

Figures are Notional-level directional estimates. See the methodology note in this report.

Executive Summary

Ford Motor Company is a global OEM with more than 100,000 employees, operations in six or more countries, and FY2025 revenue of $187.3B, up 1.2% year over year. You told us three things plainly: data is scattered and hard to use, competitors are moving faster on AI, and operating costs keep rising. The market context sharpens each one. Most OEM executives expect automakers to absorb the majority of US tariff costs, the December 2025 EV strategy reset produced a $19.5B write-down, and quality and recall costs remain a visible drag. You rate your AI maturity as Established. Our assessment is Developing (3 of 5): real AI capability exists, but fragmented data and quality systems that underdelivered limit how far it scales today.

The timing favors action. General Motors hired a chief AI officer in 2025 and applies AI to plant quality control. Toyota is deploying generative AI in vehicle design and AI agents for productivity. Volkswagen Group, with a footprint comparable to yours, reports more than 1,200 AI applications and an AWS-based factory cloud that is delivering cost savings. Industry AI adoption sits near 18%, so the leaders have not yet locked in an unassailable position, but the gap is widening with each program cycle. For a company of your scale, benchmark-level gains compound quickly: our five ranked use cases represent a directional three-year value of $300M to $564M against an estimated investment of $74M to $117M.

We ranked the use cases on your own priorities rather than on industry patterns alone. The first is an Engineering and Enterprise Knowledge Copilot ($75M to $150M, Medium effort), because it addresses scattered data directly, advances your 5% productivity target, and builds the data layer that later use cases depend on. The second is Predictive Quality and Warranty Analytics ($94M to $187M, High effort), the largest single value pool, aimed at warranty and recall cost. Because earlier AI-led quality efforts underdelivered, we design it with veteran engineers inside the loop rather than replaced by it. The third is Warranty Claims Processing Automation ($56M to $94M, Medium effort), which removes manual handling from a high-volume process and lowers operating cost without changing safety decisions. Supply chain and tariff risk sensing and a customer and dealer support assistant complete the portfolio.

The recommended first action is a focused 90-day foundation. Establish a measured productivity baseline against the 5% target, inventory the document, warranty, and quality data sources that the top three use cases need, and set a tiered AI use policy that separates internal productivity tools from safety-critical outputs. Pair this with early engagement of labor relations so the productivity case is framed around capacity and quality, not displacement. By day 90, the knowledge copilot should be in a controlled pilot and the quality analytics program should have a defined data model and a first platform scope.

AI Readiness

OVERALL READINESS

3 / 5

Developing

You rated your AI maturity as Established, which maps to Scaling (4); scattered data and AI-led quality systems that underdelivered place you one step lower today.

Business Value — 4 / 5

A stated 5% productivity target, three clear priorities, and public FY2025 revenue of $187.3B give Ford Motor Company a well-defined business case.

Technology — 3 / 5

Data is scattered and hard to use, and AI-led quality systems underdelivered, which points to integration and data foundation work before AI can scale.

Human & Change — 3 / 5

CEO-level sponsorship is a strength, but change readiness and skills ratings were not provided and a large, global, unionized workforce raises the change demands.

SINCE YOUR SPARKLE FREE SNAPSHOT

The Free snapshot scored Ford Motor Company at 4 (Scaling) from industry patterns, while this report scores it 3 (Developing) because confirmed scattered data and AI-led quality systems that underdelivered place the company one step below its stated Established maturity.

The knowledge and data assistant from the snapshot is kept as rank 1 and rescoped around engineering content and your 5% productivity target, while document automation moved into warranty claims processing and the competitive signal agent was replaced, because Research shows warranty and quality cost and tariff exposure are larger and more concrete value pools than competitor monitoring.

Your first step: Phase 1 builds on the snapshot's priority challenge of scattered data, and company data narrows it to engineering, warranty, and quality sources rather than the whole enterprise.

AI Landscape: Automotive — OEM (Vehicle Manufacturer)

The automotive industry sits in the early-to-middle part of the AI maturity curve. Industry AI adoption is around 18%, which means most OEMs have pilots and embedded vendor features but relatively few have scaled AI across engineering, manufacturing, and service. The market itself is mature, with growth near 1.0% annually, so competitive advantage comes from cost position and speed of execution rather than from market expansion. That makes AI a margin and productivity lever as much as an innovation one, which fits the pressure Ford Motor Company is already managing.

Traction is concentrated in a handful of solution areas. Predictive quality and warranty cost reduction leads, supported by McKinsey analysis showing that warranty costs can rival R&D spend and that AI-led quality can reduce them materially [1]. Generative AI for R&D and software development follows, with McKinsey estimating meaningful engineer time savings in R&D processes [2]. Manufacturing vision inspection and digital twins, connected-vehicle fleet analytics, and AI-supported customer service and personalization complete the common list, and IBM's Institute for Business Value research on automotive in the AI era points to measurable productivity gains in customer support and marketing efficiency [3]. In aftermarket and service, McKinsey finds that AI can cut labor hours and nonproductive technician time [4].

Named competitors show where the industry is moving. General Motors hired a chief AI officer in 2025, uses AI in plant quality control, works with Nvidia on factory simulation, and is bringing Google Cloud-based assistants into vehicles. Toyota applies generative AI to vehicle design and engineering, uses AI agents for productivity, and added a new North American CTO. Volkswagen Group runs over 1,200 AI applications on an AWS-based factory cloud and attributes cost savings to it. Each of these peers has paired AI with an executive owner and an enterprise data or cloud platform, not just a collection of tools.

The most common implementation barriers appear in Ford Motor Company's own market pressures. Tariff cost absorption squeezes margin and competes with AI for capital. The EV strategy reset, including the $19.5B write-down in December 2025, forces tighter investment discipline and shifts attention to near-term returns. Quality is the sharpest lesson: AI-led quality systems underdelivered, and Ford rehired about 350 veteran engineers while remaining the most recalled US automaker. The takeaway is that AI in safety-adjacent work succeeds when it extends experienced engineers' judgment, and when the underlying data is trustworthy.

What separates early movers from laggards is rarely the model. Leaders consolidate data into governed, reusable layers before they scale use cases, name a single accountable executive, measure productivity against baselines, and design human review into regulated and safety-relevant outputs. Laggards run many disconnected pilots, each with its own data

extract, and struggle to show enterprise-level results. For Ford Motor Company, with scattered data confirmed and an Established self-rating, the shift from many capable pilots to a few scaled, measured programs is the decisive step.

Top 5 Prioritized Use Cases

Rank	Opportunity	Primary Benefit	Value Range [Est.]†	Est. ROI†	Business Impact	Feasibility	CEO Alignment
1	Engineering and Enterprise Knowledge Copilot	Cuts engineering search and document time across scattered enterprise data	$75M to $150M	5.0x to 6.0x	High	Medium	Direct, advances the 5% productivity target
2	Predictive Quality and Warranty Analytics	Reduces warranty and recall cost through earlier defect detection	$94M to $187M	3.8x to 4.7x	High	Low	Direct, reduces recall and warranty cost exposure
3	Warranty Claims Processing Automation	Removes manual handling from warranty claim intake, validation, and routing	$56M to $94M	4.7x to 5.2x	Medium	Medium	Direct, lowers operating cost through automation
4	Supply Chain and Tariff Risk Sensing Agent	Anticipates tariff and supplier disruptions to protect margin and continuity	$47M to $94M	3.4x to 3.9x	Medium	Low	Direct, defends margin against tariff cost pressure
5	Customer and Dealer Support Assistant	Raises support team productivity and lowers cost to serve	$28M to $39M	3.5x to 3.9x	Low	Medium	Indirect, supports productivity and cost to serve
PORTFOLIO TOTAL			**$300M to $564M**				

† Value ranges and ROI estimates are Notional-level [Est.] directional figures based on published industry benchmarks and company size. They are not financial projections. Investment estimates reflect typical implementation costs for this solution type and company scale. See methodology note and disclaimer below.

† NOTIONAL FINANCIAL ESTIMATES: WHAT THIS MEANS

Value Range [Est.] and **Est. ROI** figures are Notional-level directional estimates produced by SPARKle™ using published industry benchmark ranges calibrated to your company size, sector, and stated priorities. They are not financial projections, audit-ready figures, or capital commitment recommendations.

Notional estimates are designed for one purpose: to establish a credible order-of-magnitude range that allows leadership to prioritize which opportunities warrant deeper investment analysis. A use case showing a $15M–$40M value range with a 3x–8x ROI should be understood as: *this class of initiative has produced outcomes in that range for comparable organizations under reasonable adoption conditions*, not as a forecast of what Ford Motor Company will achieve.

Investment-grade projections, with Low/Mid/High scenario modeling, sensitivity analysis, and client-specific financial calibration, are produced through **SPARKle Pro** using the SAVA™ v2 financial modeling methodology.
Learn more about SPARKle Pro

METHODOLOGY NOTE: HOW THESE NUMBERS WERE DERIVED

Value ranges and ROI figures are directional, Notional-level [Est.] estimates built from published improvement benchmarks from McKinsey and the IBM Institute for Business Value [1][2][3][4], applied to scoped slices of Ford Motor Company's FY2025 revenue of $187.3B that represent the addressable cost or revenue base for each use case, and cross-checked against automotive precedents for warranty analytics, engineering copilots, field service, and customer support. Because Ford's warranty, claims, and engineering cost lines were not provided, the scoped bases (between 0.2% and 1.0% of revenue) are our assumptions and should be validated with finance before any budget decision. Investment estimates reflect typical implementation costs for each solution type at this company's scale, and all figures are in US dollars and are not financial projections.

SOURCES

1. McKinsey, When Warranty Costs Rival R&D Spend: Remaking Vehicle Quality with AI, 2026
2. McKinsey, Automotive R&D Transformation: Optimizing Gen AI's Potential Value, 2023
3. IBM Institute for Business Value, Automotive in the AI Era, 2025
4. McKinsey, AI Is Already Rewiring the Aftermarket and Services, 2026

Use Cases at a Glance

HOW TO READ YOUR USE CASES

THG supports every path, from guidance to partner delivery.

Type

Automation: handles repetitive, rules-based steps.

Assistant: helps your people do the work; they make the call.

Agent: carries out multi-step tasks across systems, with human checkpoints.

Effort

Low: a few weeks, using tools you already have.

Medium: 1 to 3 months, with some integration or data preparation.

High: 3 months or more, across multiple systems.

Path

Self-serve: your team can do it with off-the-shelf tools.

THG-led: needs outside design, prioritization, or change support.

Partner-led: needs a custom build or specialist implementation.

Value and time back

Value: how much the use case is likely to benefit your business: Low, Medium, or High.

Time back: a directional estimate of team hours saved each week.

1 Engineering and Enterprise Knowledge Copilot

FROM YOUR FREE SNAPSHOT

TIME BACK **100+ hours a week**

TYPE Assistant | EFFORT Medium | PATH THG-led | VALUE High

Basis: Based on Company Revenue of $187.3B, using a scoped engineering and knowledge-work cost pool of 0.4% ($749M) and the Employees benchmark of 10% to 20% efficiency gain, with investment of $15M to $25M for an initial multi-domain rollout.

2 Predictive Quality and Warranty Analytics

TIME BACK **100+ hours a week**

TYPE Assistant | EFFORT High | PATH Partner-led | VALUE High

Basis: Based on Company Revenue of $187.3B, using a scoped warranty and quality cost exposure of 1.0% ($1.87B) and the Predict benchmark of 5% to 10% warranty cost reduction, with investment of $25M to $40M for data integration, models, and plant and service rollout.

3 Warranty Claims Processing Automation

TIME BACK **50 to 90 hours a week**

TYPE Automation | EFFORT Medium | PATH Partner-led | VALUE Medium

Basis: Based on Company Revenue of $187.3B, using a scoped claims and service administration cost base of 0.2% ($375M) and the Automate benchmark of 15% to 25% labor hour reduction, with investment of $12M to $18M.

4 Supply Chain and Tariff Risk Sensing Agent

TIME BACK **40 to 70 hours a week**

TYPE Agent | EFFORT High | PATH Partner-led | VALUE Medium

Basis: Based on Company Revenue of $187.3B, using a scoped supply disruption and tariff-exposed cost base of 0.5% ($936M) and the Predict benchmark of 5% to 10%, with investment of $14M to $24M.

5 Customer and Dealer Support Assistant

TIME BACK **20 to 35 hours a week**

TYPE Assistant | EFFORT Medium | PATH Partner-led | VALUE Low

Basis: Based on Company Revenue of $187.3B, using a scoped support and service cost base of 0.3% ($562M) and the Customers benchmark of 5% to 7% support productivity gain, with investment of $8M to $10M.

Use Case Deep Dive

1 Engineering and Enterprise Knowledge Copilot

THE OPPORTUNITY

Engineers, program managers, and plant and supply staff at Ford Motor Company work across data that you describe as scattered and hard to use. Specifications, requirements, test results, supplier documents, and past issue reports sit in separate systems and regions, so skilled people spend part of every week searching and reconciling instead of deciding. This gap sits directly behind your 5% employee productivity target.

HOW IT WORKS

A generative assistant connects to approved document and data repositories through a secure semantic search layer. An employee asks a question in plain language and receives a sourced answer, summary, or draft with links back to the original files. Access follows each person's existing permissions, so confidential engineering data stays within current rules. The assistant drafts and summarizes, and the employee reviews and decides.

WHY FORD MOTOR COMPANY, WHY NOW

Ford Motor Company named scattered data as a core obstacle, named productivity as its success measure, and has CEO-level sponsorship for this work. Toyota is already using AI agents for productivity, and Volkswagen Group reports more than 1,200 AI applications, so the speed gap you described is visible. This use case also builds the unified data layer that predictive quality and claims automation will need.

EXPECTED IMPACT

We estimate $75M to $150M of three-year value, a 5.0x to 6.0x return on $15M to $25M of investment. The range applies the Employees benchmark of 10% to 20% efficiency gain in R&D-type work [2] to a scoped engineering and knowledge-work cost pool. It is deliberately below the 20% to 30% efficiency gains seen in automotive engineering copilot precedents.

WHAT IT TAKES TO IMPLEMENT

1. Select two or three priority content domains, such as engineering requirements and supplier documents.
2. Inventory the sources, clean the highest-value content, and map access permissions.
3. Configure the retrieval layer and assistant with required source citations.
4. Run a controlled pilot with a defined engineering group and collect usage and time data.
5. Expand domain by domain based on pilot results.

The owner is the enterprise data and AI lead in IT, working jointly with an engineering operations leader. Prerequisites include access to the document repositories, identity and permission integration, a security review of any external model use, and early consultation with labor relations. This is a Medium effort initiative: expect roughly 10 to 12 weeks to a first release covering the initial domains, with expansion afterward.

KEY RISKS & CONSIDERATIONS

Generated answers can be wrong, so source citations and user feedback loops need to be built in from the start. Engineering intellectual property requires strict controls on where data is processed. Workforce messaging matters in a large, global, unionized organization: frame the tool as time returned to higher-value work.

HOW TO GET IT DONE

- **Self-serve**: Internal IT could enable an off-the-shelf enterprise assistant on a single repository, though this would not resolve cross-system data fragmentation.
- **THG-led (recommended)**: THG designs the domain prioritization, governance, permission model, and change approach, with your teams configuring and operating the tool within the $15M to $25M estimate.
- **Partner-led**: A specialist integrator builds a custom retrieval platform across all engineering systems, which suits later expansion phases.

HOW TO MEASURE SUCCESS

- **Time spent finding information**: Sample engineers' weekly search and reconciliation time through a short survey and time study before launch.
- **Hours on document tasks**: Capture current hours on drafting and summarizing from workflow or timesheet data for the pilot group.
- **Weekly active use**: Record the pilot group size and baseline tool use so adoption is measured against a known starting point.

2 Predictive Quality and Warranty Analytics

THE OPPORTUNITY

Quality and recall cost is one of the most expensive pressures on Ford Motor Company, which remains the most recalled US automaker. AI-led quality systems underdelivered, and Ford rehired about 350 veteran engineers to restore judgment in the process. Warranty, plant inspection, supplier, and field service data are fragmented, so early signals of a defect are often found late.

HOW IT WORKS

Predictive models learn from combined warranty claims, plant quality records, supplier lot data, and service reports to spot patterns that precede a defect becoming a field issue. The system ranks emerging issues by likely cost and customer impact and shows the supporting evidence. Experienced quality engineers review each alert, confirm or reject it, and their decisions feed back to improve the model. The AI recommends, and the engineers decide.

WHY FORD MOTOR COMPANY, WHY NOW

Ford Motor Company faces tariff cost absorption and a $19.5B EV write-down, so every avoidable warranty dollar counts toward margin. General Motors is already using AI in plant quality control, which raises the cost of waiting. The earlier underperformance is a design lesson: keep experienced engineers in the loop and fix the data foundation first.

EXPECTED IMPACT

We estimate $94M to $187M of three-year value, a 3.8x to 4.7x return on $25M to $40M of investment. The range applies the Predict benchmark of 5% to 10% warranty cost reduction [1] to a scoped warranty and quality cost exposure. It sits at the lower end of the 5% to 20% reduction seen in automotive predictive warranty precedents, reflecting data integration work still ahead.

WHAT IT TAKES TO IMPLEMENT

1. Choose one or two vehicle platforms and a defined set of failure categories for the first scope.
2. Integrate warranty, plant, supplier, and service data into a governed quality data model.
3. Build and validate models against historical issues the quality team already knows.
4. Deploy an alert and review workflow for quality engineers with feedback capture.
5. Measure results on the first platforms, then extend to additional platforms and plants.

The owner is the head of quality, supported by the chief data officer. Prerequisites include access to warranty and plant systems, supplier data agreements, data science capacity, and agreement from engineering leadership that alerts feed the existing quality process. This is

a High effort initiative: plan for at least four to six months to a first operating scope, with scaling over additional quarters.

KEY RISKS & CONSIDERATIONS

Data lineage and master data consistency across regions are likely the biggest dependency, so allocate early effort to them. Safety-relevant outputs require documented human review and traceability aligned with US vehicle safety obligations through the National Highway Traffic Safety Administration, and with the NIST AI Risk Management Framework as a governance reference. Avoid overpromising: report precision on past issues before relying on forward predictions.

HOW TO GET IT DONE

- **Self-serve**: Internal analytics teams could prototype alerts on one data source with existing tools, which proves the idea but not enterprise value.
- **THG-led**: THG shapes the scope, governance, and engineer-in-the-loop operating model while your teams build, which fits a smaller first phase.
- **Partner-led (recommended)**: A specialist implementer integrates the data platforms and builds the models at the scale implied by the $25M to $40M estimate.

HOW TO MEASURE SUCCESS

- **Warranty cost per vehicle**: Establish the trailing baseline from finance and warranty systems for the chosen platforms.
- **Time to detect an emerging issue**: Measure the average days from first field signal to quality action on recent issues.
- **Alert precision**: Score how many alerts quality engineers confirm, starting with a back-test on historical cases.

3 Warranty Claims Processing Automation

THE OPPORTUNITY

Ford Motor Company handles a high volume of warranty and service claims, supplier submissions, and related documents across dealers, regions, and systems. Much of the work is manual: reading claims, checking them against coverage and repair history, requesting missing information, and routing exceptions. You flagged rising operating costs as a core challenge, and this process is a repeatable place to reduce them.

HOW IT WORKS

An automation layer reads incoming claims and attachments, extracts the key fields, and checks them against policy rules, repair codes, and history. Clean claims move forward automatically, and claims with missing data trigger a templated request to the dealer or supplier. A language model summarizes complex or unusual cases so a claims specialist can decide quickly. Rule-based steps run on their own, and judgment calls stay with people.

WHY FORD MOTOR COMPANY, WHY NOW

Ford Motor Company is under cost pressure from tariffs and a reset EV strategy, and a process automation program can fund itself in operating cost terms. Document-heavy work was among the pain points identified in your discovery. The same claims data also strengthens the predictive quality program, so the two initiatives reinforce each other.

EXPECTED IMPACT

We estimate $56M to $94M of three-year value, a 4.7x to 5.2x return on $12M to $18M of investment. The range applies the Automate benchmark of 15% to 25% labor hour reduction [4] to a scoped claims and service administration cost base. Field service precedents showing about a 10% first-time-fix gain support the direction on service productivity.

WHAT IT TAKES TO IMPLEMENT

1. Map the current claims workflow and identify steps that follow clear rules.
2. Define the data fields, coverage rules, and exception categories with the claims team.
3. Connect the automation layer to the claims and repair history systems.
4. Pilot on one claim type and region, with specialist review of every exception.
5. Extend to additional claim types as accuracy and cycle time targets are met.

The owner is the head of warranty and service operations, with IT integration support. Prerequisites include API or file access to claims systems, documented coverage rules, dealer and supplier communication templates, and approval from finance and legal for

automated decision thresholds. This is a Medium effort initiative: expect roughly two to three months to a first operating pilot.

KEY RISKS & CONSIDERATIONS

Rules vary by region and program, so start with one claim type and expand deliberately. Automated approvals should have value thresholds and sampling audits to guard against error and fraud. Claims specialists should be involved from the start, with the goal framed as redeploying time to complex cases and dealer support.

HOW TO GET IT DONE

- **Self-serve**: Your operations team could automate one document intake step with off-the-shelf tools, which gives early learning but limited integration.
- **THG-led**: THG supplies process design, rules definition, and change management while your IT team handles the integration.
- **Partner-led (recommended)**: A specialist implementer integrates with the claims platforms and builds the automation, aligned with the $12M to $18M estimate.

HOW TO MEASURE SUCCESS

- **Cycle time per claim**: Pull the current average and 90th percentile processing time from the claims system.
- **Share of claims handled without touch**: Sample recent claims to establish how many could have passed without manual intervention.
- **Specialist hours per claim**: Capture current labor hours per claim from workload reports before launch.

Quick-Win Recommendations

1 Baseline the 5% Productivity Target

Define what a 5% employee productivity improvement means in measurable terms: hours on document search, drafting, reporting, and approvals for two or three representative functions. Use short surveys and existing workflow data to set today's baseline. This is a quick win because it needs no new technology and makes every later use case measurable. Timeline: 3 to 4 weeks.

2 Inventory Existing AI Deployments

You rate your AI maturity as Established, so deployed models and tools already exist in engineering, manufacturing, supply chain, and analytics teams. Catalog what is running, who owns it, which data it uses, and what value it has shown. The inventory exposes duplication, shows which proven use cases can scale, and identifies where fragmented data is slowing results. Timeline: 4 to 6 weeks.

3 Issue a Tiered AI Use Policy

Publish a short policy that separates internal productivity uses, which need light review, from engineering, quality, and safety-relevant outputs, which require documented human sign-off. Include rules for confidential engineering data and external generative tools, referencing the NIST AI Risk Management Framework. This gives teams clarity to proceed and reduces intellectual property exposure. Timeline: 4 to 6 weeks.

4 Run a Warranty and Quality Data Source Review

Convene quality, warranty, plant, and IT leaders to list every data source that touches field defects, with owners, refresh rates, and known quality issues. The output is a ranked list of integration priorities for the predictive quality program, grounded in the lesson that earlier AI-led quality systems underdelivered. Timeline: 6 weeks.

5 Open the Workforce Dialogue Early

Brief labor relations and employee representatives on the productivity goals before pilots begin, and frame the aim as removing low-value document work and improving quality. Early transparency lowers the risk of resistance in a large, global, unionized workforce. Timeline: initial briefings within 4 weeks, ongoing thereafter.

90-Day Action Plan

1 Phase 1 (Days 1-30): Foundation

This phase sets the measurement baseline, the governance rules, and the data scope. It builds on the priority challenge that scattered data is hard to use, and company data narrows it from an enterprise-wide data problem to the specific engineering, warranty, and quality sources that the top three use cases depend on.

- The chief data officer and enterprise IT lead complete a data source inventory for engineering documents, warranty claims, plant quality, and supplier records.
- Finance and HR define the productivity baseline for two or three functions against the 5% target.
- Legal, security, and the AI governance lead publish the tiered AI use policy and confirm US regulatory touchpoints, including vehicle safety obligations and state privacy laws.
- Labor relations holds initial briefings on the AI program and its goals.

Success criterion: a signed-off data scope, a documented productivity baseline, and an approved AI use policy by day 30.

2 Phase 2 (Days 31-60): Build

This phase turns the scope into working foundations. The focus is on connecting the first content domains for the knowledge copilot and designing the quality data model, with business cases confirmed against finance's real cost bases.

- The enterprise data and AI lead configures the retrieval layer and permission model for the initial two or three engineering content domains.
- The head of quality and the data team define the quality data model and select the first vehicle platforms and failure categories for the predictive analytics program.
- Warranty and service operations maps the claims workflow and identifies the first claim type for automation.
- Finance validates the scoped cost bases behind each use case and refines the investment estimates.

Success criterion: the knowledge copilot is configured and ready for a pilot group, and the quality data model and claims scope are approved by their business owners.

3 Phase 3 (Days 61-90): Launch

This phase puts the first capabilities in front of users and prepares the larger commitments. The knowledge copilot enters a controlled pilot, while the longer predictive quality and claims programs move to partner selection and detailed planning, consistent with their

longer delivery timelines.

- The engineering operations leader launches the knowledge copilot pilot with a defined group and begins tracking usage and time data.
- The head of quality and procurement issue requirements for implementation partners on the predictive quality and claims automation programs.
- The program office reports pilot results against the baseline to the CEO and executive leadership.
- Executive leadership decides on funding for the next wave, including supply chain and tariff risk sensing and the customer and dealer support assistant.

Success criterion: a running pilot with measured early results, selected delivery partners for the two larger programs, and a funding decision on the next wave by day 90.

Need board-ready numbers?

This Express report ranks your top opportunities with directional value ranges and a 90-day plan. **SPARKle Pro** is built for funding a real initiative: a full prioritized portfolio with an impact and feasibility matrix, investment-grade SAVA™ v2 financials with low, mid, and high scenarios, a 12-month Three Horizons roadmap, and an interactive portal to explore use cases and financials.

Explore SPARKle Pro

Ready to start on a use case? THG supports every delivery path in this report, from guidance for your own team to THG-led delivery or a vetted partner. Reply to your report email or write to **info@theholistyxgroup.com**.

This report was prepared by The Holistyx Group using the SPARKle™ AI-assisted analysis platform. The insights and recommendations are based on the information provided and publicly available industry data. Financial estimates, where included, are Notional-level directional figures based on published industry benchmarks and company size. They are not financial projections and should not be treated as guarantees or capital commitment figures. Investment-grade financial modeling is available through SPARKle Pro. The Holistyx Group recommends validation of all recommendations through appropriate internal review processes before implementation.

