PREPARED FOR

Mike Sternberg

CEO, Ford Motor Company

AI Transformation Roadmap

SPARK™ Executive Readout · The Holistyx Group · November 2026

25	7	$181.11M – $375.98M
Use Cases Evaluated	Use Cases Shortlisted	3-Year Value Range

About This Report

This report was produced via SPARKle - The Holistyx Group's accelerated AI-transformation methodology. It applies the full SPARK™ framework (see Appendix B for SPARK™ details) through an AI-assisted pipeline, delivering a complete engagement output in 24–48 hours. All financial estimates are Notional-level [Est.], for prioritization and board framing.

Discovery Summary

Here's what shaped this analysis. Take a moment to confirm it's right.

COMPANY

Company	Ford Motor Company	YOU TOLD US
Industry	Automotive — OEM (Vehicle Manufacturer)	YOU TOLD US
Headquarters	Detroit, Michigan, United States	YOU TOLD US
Company size	100,000+ employees	YOU TOLD US
Estimated annual revenue	$187.3B (FY2025, Ford Motor Company Form 10-K FY2025)	PUBLIC RESEARCH

STATED CHALLENGES

Primary challenge	Data is scattered and hard to use, and competitors are moving faster on AI.	YOU TOLD US
Additional context	Rising operational costs, difficulty hiring and retaining skilled staff, and difficulty finding the right use cases and building business cases and ROI.	YOU TOLD US
Success looks like	A 5% improvement in employee productivity and a reduction in warranty cost.	YOU TOLD US
Industry validation	Industry warranty costs reached about $58 billion globally in 2024 and exceed 4% of revenue for some OEMs (McKinsey, 2026).	PUBLIC RESEARCH

CURRENT AI POSTURE

AI maturity	Established (mapped to Scaling), with plant-floor predictive maintenance as the active initiative.	YOU TOLD US
Data availability	Less siloed, good quality, with standardized analytics (historical dashboards and reports).	YOU TOLD US
Change readiness	Cautious.	YOU TOLD US
Peer comparison	No public adoption rate is available for automotive OEMs. Industry precedents show predictive warranty and quality analytics delivering 5%–10% warranty cost reduction in automotive.	PUBLIC RESEARCH

WHAT WE FOUND

Research scope: Public financial data from your FY2025 Form 10-K, competitor positioning, market pressures, regulatory exposure, and automotive AI use case precedents.

Competitive signal	General Motors hired a chief AI officer and uses AI for plant equipment monitoring, quality control, and dealer ordering.	PUBLIC RESEARCH

Discovery Summary (continued)

WHAT WE FOUND (CONTINUED)

Industry context	Most OEM executives expect automakers to bear the majority of US tariff costs, pressuring margins while new vehicle sales sentiment softens.	PUBLIC RESEARCH

If anything above looks off, tell us. This summary is the foundation for everything that follows: the opportunities we've prioritized, the financial ranges we've estimated, and the roadmap we've built. A wrong number or an outdated fact here can throw off the analysis downstream, so it's worth a 30-second check before you read further.

Spot something that doesn't match your business? Reply to this email or contact us directly and we'll correct it and regenerate the affected sections at no charge.

SECTION 1 OF 6

Executive Summary

Confidence Qualification: This report was produced by the SPARKle AI engine. H1 (Foundation & Quick Wins, Months 1–4) at 80% confidence, H2 (Scale & Expansion, Months 5–8) at 55% confidence, and H3 (Transform, Months 9–12+) at 35% confidence, with no H3 use cases shortlisted. These are Notional-level estimates, directional for prioritization and board framing, and they are inferred from your client profile rather than a live discovery engagement. We recommend that all figures be validated in a Day 2 Technical Deep Dive before commitment.

AI Readiness Scorecard

You are scaling from a proven first initiative. Plant-floor predictive maintenance is live, and you rate data availability as 'Less siloed, good quality' with standardized analytics, which supports your strongest readiness scores. Your constraints are just as clear: you describe change readiness as cautious, AI governance as actively developing, and AI technical and prompt engineering skills as medium, while scattered data and difficulty hiring and retaining skilled staff remain on your challenge list. The pattern is strong direction with limited organizational capacity to absorb change. That is why the roadmap starts with assistant-style and automation pilots that people review and approve, and builds data foundations before heavier predictive work. Overall you score 3.3 of 5.0, moving from Developing toward Scaling.

Dimension	Readiness Signal	Score
Strategic Alignment	Three clear priorities and a board push for ROI-defensible AI value.	4.0/5
Data Readiness	Good-quality, less siloed data and standardized analytics, yet scattered data remains a stated challenge.	3.0/5
Technical Infrastructure	Predictive maintenance is live; AI technical and prompt engineering skills rated Medium.	3.5/5
Organizational Readiness	Change readiness rated Cautious, with difficulty hiring and retaining skilled staff.	2.5/5
Governance & Risk	AI governance rated actively developing; exposure across 6+ countries adds regulatory complexity.	3.0/5
Executive Sponsorship	Competitive and cost pressure act as board-level forcing functions for AI investment.	4.0/5
Overall Readiness	*Developing ◪ Scaling*	**3.3/5**

Scale: 1 = Early stage | 2 = Experimenting | 3 = In progress | 4 = Optimized | 5 = Industry-leading

Strategic Implication

The single highest-leverage action Ford Motor Company can take in the next 90 days is:

Consolidate warranty claims and repair data

This step brings claim, repair-history, and coverage data from separate systems and regional formats into one clean dataset, which takes roughly seven weeks. Without it, Automated Warranty Claims Adjudication cannot reach its shadow-mode pilot, and Early Defect and Warranty Cost Prediction cannot be trusted. With it, you have a warranty cost baseline to measure against your stated goal of reducing warranty cost, and both use cases can proceed on one shared foundation.

Three Key Findings

01 Your 5% productivity goal is within reach

You told us you want a 5% improvement in employee productivity, and your engineering and finance teams are where it starts. The Engineering and Software Development Copilot carries a $73.03M–$121.72M 3-year value range, and the Finance and Shared Services Productivity Copilot adds $28.21M–$47.01M. McKinsey automotive benchmarks, plus an industry precedent of 15%–25% labor time reduction for field service copilots, show the direction of travel at your 100,000+ employee scale. Together these H1 pilots can show measured results by Months 2 and 3, giving you evidence before any H2 funding is released.

02 Scattered warranty data is your primary constraint

Your stated challenge is data that is scattered and hard to use, and it shows up most sharply in warranty. Claim, repair-history, and coverage data sit in separate systems and regional formats, which is the long pole for Automated Warranty Claims Adjudication and likely stretches its pilot to about 20 weeks. Industry warranty costs reached about $58 billion globally in 2024, per McKinsey research, and exceed 4% of revenue for some OEMs. Your warranty line is not separately disclosed, so the baseline you build now anchors the $19.48M–$77.9M value range.

03 Close the AI speed gap beyond efficiency

Competitors are moving faster: General Motors hired a chief AI officer and uses AI for plant equipment monitoring and quality control, while Honda applies AI to warranty claims. Your answer is a sequence, with $181.11M–$375.98M of 3-year value across seven use cases. One hypothesis is not yet addressed: In-Line Computer Vision Quality Inspection (UC-008) ranked outside the 20 highest-ranked of 25 use cases, so it was deferred on composite rank alone. Revisit it at the Month 8 gate, once warranty and plant data pipelines show progress, alongside Predictive Maintenance Expansion Across Plants as your enterprise-scale H2 destination.

BUILDING ON YOUR EARLIER RESULTS

Since Your Previous SPARKle Report

This readout builds on your earlier SPARKle™ Express report and Free snapshot. That work drew on industry patterns and benchmarks. This assessment adds your company's own financials, 25 evaluated use cases, and a phased roadmap. Here is what changed, and why.

Your AI readiness

Your earlier Express report scored readiness at 3 (Developing); this assessment scores 3.3 using a six-dimension scorecard and your discovery ratings for change readiness, data quality, and AI governance, which the earlier three-lens company view did not examine.

This assessment: 3.3/5 | Developing Scaling

Your opportunities

Four earlier opportunities made this shortlist: Engineering and Software Development Copilot, Back-Office Workflow Automation, Early Defect and Warranty Cost Prediction, and Predictive Maintenance Expansion Across Plants. Unified Enterprise Knowledge and Data Search, Agentic Sales and Service Assistant, and the Closed-Loop Product and Process Intelligence Hub ranked below the seven shortlisted use cases on this assessment's scoring, so they sit outside the shortlist rather than being discarded.

Your value estimates

Express showed a total value range of $363M – $790M by calibrating published benchmarks to company size. This report models each use case against your own disclosed financials, cost lines, and Three Horizons phasing, producing $181.11M – $375.98M.

This assessment: $181.11M – $375.98M portfolio value | 1.9x – 6.3x ROI

Where this readout differs from your earlier results, it reflects more information about your organization, not a change in method. Use the figures and priorities in this report for planning and funding decisions.

SECTION 2 OF 6

Use Case Portfolio Summary

Of the 20 highest-ranked of 25 use cases, 11 landed in Quick Win, 6 in Strategic Investment, and 3 in Experimental, with none deprioritized, so your portfolio is broadly feasible. Priority 1, cost reduction and operational efficiency, is addressed by Back-Office Workflow Automation and Automated Warranty Claims Adjudication. Priority 2, closing the AI speed gap, is addressed by the Engineering and Software Development Copilot and Predictive Maintenance Expansion Across Plants. Priority 3, talent attraction, retention, and productivity, is addressed by the Finance and Shared Services Productivity Copilot and the Recruiting and Onboarding Assistant.

Prioritization Matrix

STRATEGIC INVESTMENTS *High Impact | Lower Feasibility* #6 Predictive Maintenance Expansion Across Plants (H2)	**QUICK WINS** *High Impact | High Feasibility* **#1 Engineering and Software Development Copilot (H1)** **#2 Back-Office Workflow Automation (H1)** **#3 Automated Warranty Claims Adjudication (H1)** **#4 Finance and Shared Services Productivity Copilot (H1)** **#5 Early Defect and Warranty Cost Prediction (H2)**
RECONSIDER *No use cases. A strong result.*	**EXPERIMENTAL** *1 use cases deferred pending H1/H2 validation.*

Axis: Impact (strategic value, 1–10) × Feasibility (implementation readiness, 1–10). Feasibility scores are estimated [Est.] and would be confirmed through further validation, such as a Day 2 Technical Deep Dive.

SECTION 2 OF 6 · CONTINUED

Shortlisted Use Case Portfolio

Use Case	Horizon	Value Low	Value High	Investment	ROI Range
Engineering and Software Development Copilot	**H1**	**$73.03M**	**$121.72M**	**$7.0M – $9.8M**	**7.5x – 17.4x**
Back-Office Workflow Automation	**H1**	**$7.52M**	**$30.09M**	**$3.0M – $5.2M**	**1.4x – 10.0x**
Automated Warranty Claims Adjudication	**H1**	**$19.48M**	**$77.9M**	**$5.5M – $9.5M**	**2.1x – 14.2x**
Finance and Shared Services Productivity Copilot	**H1**	**$28.21M**	**$47.01M**	**$3.6M – $6.4M**	**4.4x – 13.1x**
Recruiting and Onboarding Assistant	**H2**	$19.39M	$32.32M	$9.0M – $15.0M	**1.3x – 3.6x**
Early Defect and Warranty Cost Prediction	**H2**	$16.74M	$33.47M	$14.7M – $22.0M	**0.8x – 2.3x**
Predictive Maintenance Expansion Across Plants	**H2**	$16.74M	$33.47M	$17.0M – $30.0M	**0.6x – 2.0x**
PORTFOLIO TOTAL		**$181.11M**	**$375.98M**	**$59.8M – $97.9M**	**~8.1x avg**

Sources: Industry-specific AI benchmark research (see Appendix A) | SAVA™ v2 Model — Ford Motor Company engagement

Section 3 of 6 — Top 3 Use Case Narratives

The following narratives detail the three highest-priority Quick Win use cases recommended for immediate H1 action. Each narrative follows the SPARK™ structure: Challenge / Future State / Implementation Path / Key Risks / Business Value / Pilot Scope.

SECTION 3 OF 6 · USE CASE 1 OF 3

1. Engineering and Software Development Copilot

Horizon: H1 · Solution Area: Employees · Months 1–2

The Challenge Today

Your engineering and software development teams spend significant time on repetitive coding, code review, documentation, and test creation. They work across large, complex vehicle and enterprise software codebases, so productivity gains from modern AI tooling are not yet captured consistently across the organization.

With competitors moving faster on AI, including General Motors' use of AI in plant monitoring and quality control, uneven engineering productivity slows your speed while skilled engineers stay hard to hire and retain. Without a consistent approach, the 5% employee productivity improvement you are targeting stays out of reach.

The AI-Enabled Future State

"Imagine if every Ford engineer had an AI pair-programmer that understood internal coding standards, drafted code and tests, explained legacy systems, and handled documentation, so the engineer could focus on design decisions and innovation rather than boilerplate and rework."

Your engineering teams get an assistant that suggests code, generates tests and documentation, summarizes code changes for review, and answers questions about internal codebases. A person reviews and approves all output, so speed rises while quality and security standards hold. Pilot results reach you by Month 2.

IMPLEMENTATION PATH

A hosted, enterprise-grade code-assistance service with developer-environment plug-ins and single sign-on, configured by your own engineering and IT teams, with no vendor mandated. The key build consideration is limiting repository access to pilot teams and adding security scanning, human review of generated code, and usage telemetry.

KEY RISKS

Security and intellectual-property review for exposing proprietary vehicle and software source code to a hosted AI service is the likeliest blocker, as approvals may run longer than technical setup. Starting the review in Week 1 and limiting the pilot to non-restricted repositories resolves it. Intellectual property protection is the governance consideration.

Business Value

15%–25% labor time reduction from technician copilots in industrial and automotive OEM settings	0.65 capture rate applied to McKinsey automotive improvement benchmarks for Established AI maturity	7.5x – 17.4x
Research precedent: field service technician copilots	SAVA anchor; McKinsey automotive benchmarks	3-Year ROI Range

Continued on the next page

SECTION 3 OF 6 · USE CASE 1 OF 3 · CONTINUED

Pilot Scope

Deploy the copilot to a limited set of software engineering teams working on non-restricted repositories, covering code suggestions, test generation, documentation drafting, and code-change summaries. Measure productivity against a pre-pilot baseline and capture developer feedback to inform a broader rollout decision.

Success metric: Measured reduction in time-to-complete for a defined set of routine engineering tasks (such as pull-request cycle time) against baseline, with a developer satisfaction score of at least 70% favorable.

SECTION 3 OF 6 · USE CASE 1 OF 3 · IMPLEMENTATION PLAN

1. Engineering and Software Development Copilot

Owner: Director, Engineering Productivity

TYPE **Assistant** | EFFORT **Low** | PATH **Self-serve** | TIME BACK **3,000 to 6,000 hours a week**

How to Get It Done

Plan on about 8 weeks, most of it spent on security review, repository scoping, and developer onboarding before a pilot with selected engineering teams.

1. Weeks 1–3: complete security, IP, and data-protection review and set repository scope for pilot teams.
2. Weeks 3–5: configure single sign-on, developer plug-ins, telemetry, and guardrails for pilot repositories.
3. Weeks 5–6: onboard pilot developers, capture baseline productivity metrics, and run short enablement sessions.
4. Weeks 6–8: run the pilot, gather usage and feedback data, and prepare the go/no-go readout.

PREREQUISITES

- Security and IP review approved for pilot repositories
- Single sign-on and developer-environment access available to pilot teams
- Pre-pilot productivity baseline captured, such as pull-request cycle time
- Acceptable-use policy defined for generated code

Delivery Options

Self-serve RECOMMENDED
Your engineering and IT teams can configure and roll out an off-the-shelf coding assistant themselves, with no custom build or model training required.

THG-led
Outside design, prioritization, or change support would add rollout and enablement help, though the configuration itself needs no custom build.

Partner-led
A custom build or specialist implementation would be needed only if you later want deep integration with proprietary internal engineering tooling.

How to Measure Success

KPI	How to capture today's baseline
Pull-request cycle time versus baseline for routine engineering tasks	Record current pull-request cycle time for the pilot teams over the weeks before launch.
Developer satisfaction score, percent favorable	Survey pilot developers before launch to capture current tooling satisfaction.

SECTION 3 OF 6 · USE CASE 2 OF 3

2. Back-Office Workflow Automation

Horizon: H1 · Solution Area: Automate · Months 1–4

The Challenge Today

Your finance operations, procurement, and administrative processing rely on repetitive, rules-based tasks involving manual data entry, reconciliation, and routing across multiple systems. This consumes skilled staff time and slows cycle times across a global enterprise.

Selling, administrative, and other expenses were $10.8B in FY2025 per your Form 10-K, and rising operational costs are already on your challenge list. Manual back-office effort is a direct drag on the cost reduction and efficiency mandate you described.

The AI-Enabled Future State

"Imagine if routine back-office work, such as invoice matching, data reconciliation, and request routing, simply happened in the background, with exceptions flagged to the right person automatically, and teams spent their time on analysis and decisions instead of keystrokes."

Automated workflows handle repetitive, rules-based steps across systems, extract and validate data from documents, and route exceptions to people with context attached. Cycle times shrink, error rates fall, and audit trails are captured automatically, freeing capacity for analysis and decisions.

IMPLEMENTATION PATH

A cloud-based workflow automation and intelligent document-processing capability connected to your existing finance and procurement systems, built with specialist integration help, with no vendor mandated. The most important build consideration is integrating the source systems and running the workflow in parallel with the manual process before switching over.

KEY RISKS

Integrating with multiple enterprise back-office systems and getting clean, consistent source data for the chosen process is the longest lead item, as access approvals and data cleanup tend to slip. Securing access and cleaning data in Weeks 3–8 resolves it. Financial controls and segregation of duties is a second consideration.

Business Value

3-Year Value Range: $7.52M–$30.09M · **ROI Range:** 1.4x – 10.0x · **19% support-function cost reduction from automated invoice processing in automotive** (Research precedent: automated invoice processing) · **$10.8B selling, administrative, and other expenses, 5.8% of revenue** (Ford Motor Company Form 10-K FY2025)

Continued on the next page

SECTION 3 OF 6 · USE CASE 2 OF 3 · CONTINUED

Pilot Scope

Automate a single high-volume, rules-based back-office process, such as invoice-to-payment matching, including document extraction, validation, and exception routing. The pilot runs in parallel with the manual process on a defined transaction subset to measure accuracy, cycle time, and effort saved.

Success metric: At least a 40% reduction in manual handling time per transaction on the pilot process, with accuracy matching or exceeding the manual baseline.

SECTION 3 OF 6 · USE CASE 2 OF 3 · IMPLEMENTATION PLAN

2. Back-Office Workflow Automation

Owner: Director, Finance Shared Services Operations

TYPE **Automation** | EFFORT **High** | PATH **Partner-led** | TIME BACK **150 to 300 hours a week**

How to Get It Done

Plan on about 16 weeks, most of it spent mapping the chosen process, integrating finance and document systems, and tuning exception handling before a pilot on one high-volume workflow.

1. Weeks 1–3: choose the pilot process, map current steps and exceptions, and set baseline cycle-time and cost metrics.
2. Weeks 3–8: secure system access, clean source data, and standardize document formats for the pilot process.
3. Weeks 6–13: build the automated workflow, document extraction, validation rules, and exception routing, integrated with source systems.
4. Weeks 13–16: run alongside the manual process, tune exceptions, measure results, and prepare the go/no-go readout.

PREREQUISITES

- System access approvals for the pilot finance and procurement processes
- Clean, consistent source data for the chosen process
- Financial controls and segregation-of-duties design signed off
- Specialist integration capacity secured before the build phase

Delivery Options

Self-serve
Your team could try off-the-shelf automation tools on a narrow process, but they would not cover multi-system integration and custom exception handling.

THG-led
Outside design, prioritization, and change support would help with process selection and adoption, but not the integration build itself.

Partner-led RECOMMENDED
The work requires specialist integration across several enterprise systems and a custom workflow build, beyond what your team can do with off-the-shelf tools alone.

How to Measure Success

KPI	How to capture today's baseline
Manual handling time per transaction on the pilot process	Time a sample of pilot-process transactions end to end before launch.
Accuracy compared with the manual baseline	Record current error and exception rates on the pilot process before the parallel run.

SECTION 3 OF 6 · USE CASE 3 OF 3

3. Automated Warranty Claims Adjudication

Horizon: H1 · Solution Area: Automate · Months 1–5

The Challenge Today

Your teams process a very large volume of warranty claims from dealers and service channels. Adjudication means manually reviewing claim details, repair histories, and coverage rules, which is slow, labor-intensive, and prone to inconsistent decisions across regions. Your warranty and field service cost line is not separately disclosed in the Form 10-K.

Industry warranty costs reached about $58 billion globally in 2024 and exceed 4% of revenue for some OEMs, while Honda already applies AI to detect safety-critical warranty claims. Reducing warranty cost is a target you named, and without faster, consistent adjudication, claims cost keeps absorbing margin already pressured by tariffs.

The AI-Enabled Future State

"Imagine if routine warranty claims were reviewed, checked against coverage rules and repair history, and approved in minutes, with only the unusual or high-value claims going to a human adjudicator who already has the case summarized."

An automated adjudication capability reads each claim, validates it against coverage terms and part and repair history, and approves clear cases. It flags anomalies and complex claims for adjudicators with a summarized rationale and keeps a full decision record, so your adjudicators spend their time on the cases that need judgment.

IMPLEMENTATION PATH

A cloud-based machine learning and rules-decisioning capability with document-processing support, connected to your claims and dealer systems. The most important build consideration is consolidating claim and repair-history data for one claim category, then running the model in shadow mode with an auditable decision record before any automation.

KEY RISKS

Consolidating claim, repair-history, and coverage data across systems and regional formats into one clean dataset is the biggest blocker and is likely to extend the timeline to about 20 weeks. Starting consolidation in Week 3 for one claim category contains it. Auditability of automated claim decisions is a second consideration.

Business Value

3-Year Value Range: $19.48M–$77.9M · **ROI Range:** 2.1x – 14.2x · **5%–10% warranty cost reduction from predictive warranty and quality analytics in automotive** (Research precedent: predictive warranty and quality analytics) · **52% labor cost reduction from warranty claims automation in automotive** (Research precedent: warranty claims automation) · **About $58 billion in global industry warranty costs in 2024, above 4% of revenue for some OEMs** (McKinsey, 2026)

Continued on the next page

SECTION 3 OF 6 · USE CASE 3 OF 3 · CONTINUED

Pilot Scope

Run an automated adjudication model in shadow mode on one high-volume, routine warranty claim category. The model recommends approve, deny, or refer decisions that are compared with human adjudicator outcomes, with no automatic payments, to measure accuracy, consistency, and potential processing time saved.

Success metric: At least 85% agreement between model recommendations and senior adjudicator decisions on the pilot category, with a projected reduction of at least 40% in manual review time.

SECTION 3 OF 6 · USE CASE 3 OF 3 · IMPLEMENTATION PLAN

3. Automated Warranty Claims Adjudication

Owner: Director, Warranty Operations

TYPE **Agent** | EFFORT **High** | PATH **Partner-led** | TIME BACK **80 to 200 hours a week**

How to Get It Done

Plan on about 20 weeks, most of it spent consolidating claim and part-history data and encoding coverage rules before a shadow-mode pilot on one claim category.

1. Weeks 1–3: select the pilot claim category, document adjudication rules, and measure baseline processing time, cost, and consistency.
2. Weeks 3–10: consolidate and standardize claims, repair-history, and coverage data for the pilot category across regional formats.
3. Weeks 8–16: encode coverage rules, train the decision-support model, build integrations, and set audit logging and referral thresholds.
4. Weeks 16–20: run the model alongside human adjudicators in shadow mode, tune thresholds, and prepare the go/no-go readout.

PREREQUISITES

- Claim, repair-history, and coverage data accessible for one claim category
- Senior adjudicators available to document rules and review model decisions
- Audit logging requirements for automated decisions agreed
- Regional data formats mapped and standardized

Delivery Options

Self-serve
Off-the-shelf tools could summarize claims for adjudicators, but they would not deliver a trained decision model integrated with your claims and dealer systems.

THG-led
Outside design, prioritization, and change support would sharpen scoping and adjudicator adoption, but the model and integrations still need specialist build capability.

Partner-led RECOMMENDED
The work needs a trained decision model and integration of claims, repair-history, rules, and dealer systems, which calls for specialist build and integration capability.

How to Measure Success

KPI	How to capture today's baseline
Agreement between model recommendations and senior adjudicator decisions	Have senior adjudicators decide a sample of pilot-category claims before launch to set the comparison benchmark.
Manual review time per claim in the pilot category	Measure current average adjudication time and cost per claim for the pilot category before the shadow run.

SECTION 4 OF 6

Business Case — Financial Summary

The following table summarizes the SAVA™ v2 (Solution-Area Value Anchoring) financial model for Ford Motor Company's shortlisted use case portfolio. SAVA anchors each use case's value and investment range to its Solution Area, deployment horizon, and scoring, rather than relying on a single generic ROI assumption — full methodology details are in Appendix A. All values represent 3-year Notional-level estimates, built from published industry benchmark ranges and calibrated to Ford Motor Company's specific operational profile.

Financial ranges (Low/High) reflect conservative and aggressive adoption scenarios respectively. Investment ranges include discovery, development, integration, and change management costs.

ID	Use Case	Horizon	Value Low	Value High	Investment Range	Payback	ROI Low	ROI High
UC-013	**Engineering and Software Development Copilot**	**H1**	**$73.03M**	**$121.72M**	**$7.00M - $9.80M**	**0.26 yr**	**7.5x**	**17.4x**
UC-006	**Back-Office Workflow Automation**	**H1**	**$7.52M**	**$30.09M**	**$3.00M - $5.20M**	**0.65 yr**	**1.4x**	**10.0x**
UC-007	**Automated Warranty Claims Adjudication**	**H1**	**$19.48M**	**$77.90M**	**$5.50M - $9.50M**	**0.46 yr**	**2.1x**	**14.2x**
UC-012	**Finance and Shared Services Productivity Copilot**	**H1**	**$28.21M**	**$47.01M**	**$3.60M - $6.40M**	**0.4 yr**	**4.4x**	**13.1x**
UC-015	Recruiting and Onboarding Assistant	**H2**	$19.39M	$32.32M	$9.00M - $15.00M	**1.39 yr**	**1.3x**	**3.6x**
UC-002	Early Defect and Warranty Cost Prediction	**H2**	$16.74M	$33.47M	$14.70M - $22.00M	**2.19 yr**	**0.8x**	**2.3x**
UC-001	Predictive Maintenance Expansion Across Plants	**H2**	$16.74M	$33.47M	$17.00M - $30.00M	**2.81 yr**	**0.6x**	**2.0x**
	PORTFOLIO TOTAL		**$181.11M**	**$375.98M**	**$59.8M – $97.9M**	**—**	**~8.1x avg**	

Methodology: SAVA™ v2 (Solution-Area Value Anchoring with Horizon-Based Ceilings). ROI = 3-Year Value ÷ Total Investment. H1 use cases (blue) are recommended for immediate action. All values are Notional [Est.] — directional for board framing and prioritization.

SECTION 5 OF 6

12-Month AI Transformation Roadmap

Ford Motor Company's AI roadmap is organized across Three Horizons, each with a clear objective, a defined use case portfolio, an investment envelope, and a Gate decision that governs progression to the next phase. The Three Horizons structure is not a rigid schedule — it is a sequencing strategy that manages dependency risk and ensures each phase learns from the last.

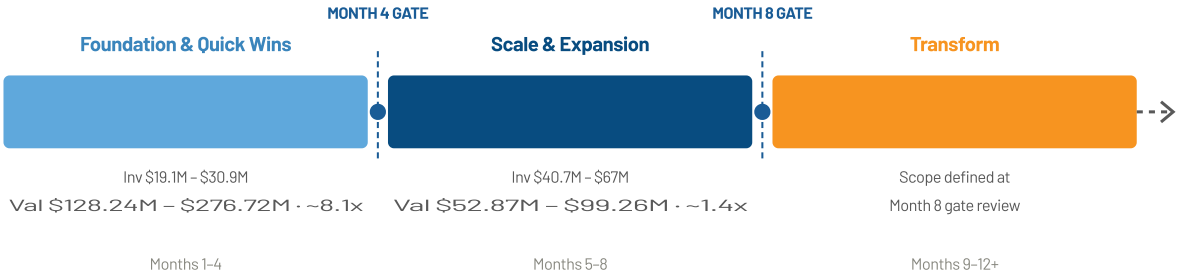

H1 | Foundation & Quick Wins | Months 1–4

In the first four months you bank your near-term, ROI-defensible efficiency wins. The Engineering and Software Development Copilot and the Finance and Shared Services Productivity Copilot move into pilot first, and neither depends on another use case. Back-Office Workflow Automation reports in Month 4, while Automated Warranty Claims Adjudication consolidates claim data and reports shadow-mode results in Month 5. Security review for source code and access approvals to financial systems are the likeliest delays. H1 carries $128.24M–$276.72M of 3-year value against $19.1M–$30.9M of investment at 80% confidence.

Investment Envelope: $19.1M – $30.9M **Value Range:** $128.24M – $276.72M **Avg ROI: ~8.1x**

▶ Month 4 Gate Decision

Do the Engineering and Software Development Copilot, Finance and Shared Services Productivity Copilot, and Back-Office Workflow Automation pilots show their success metrics, with warranty claims data consolidation on track? If yes, H2 funding and starts are released. If not, you hold H2 and address the blocker first.

SECTION 5 OF 6 · CONTINUED

H2 | Scale & Expansion | Months 5–8

Months 5 through 8 are where proof becomes platform. With H1 results in hand, you start the Recruiting and Onboarding Assistant, Early Defect and Warranty Cost Prediction, and Predictive Maintenance Expansion Across Plants. None reports a pilot result by Month 8; the work is data and integration across systems, regions, and plants. H2 carries $52.87M–$99.26M of value against $40.7M–$67M of investment at 55% confidence, a 1.4x mid-case return. Treat this as your downside risk scenario: low-case returns are 1.3x, 0.8x, and 0.6x respectively, driven by harmonized warranty data, privacy and fairness review, and multi-country plant data pipelines. Anchor board commitment on the conservative scenario, because these figures are inferred from your profile rather than a live discovery engagement.

Investment Envelope: $40.7M – $67M **Value Range:** $52.87M – $99.26M **Avg ROI:** **~1.4x**

Month 8 Gate Decision

Has H1 delivered the value forecast, and are warranty and plant data pipelines progressing well enough to justify continuing H2 investment? No H2 pilot reports by Month 8, so this gate rests on H1 results and pipeline progress. It is also the point to revisit deferred use cases.

SECTION 5 OF 6 · CONTINUED

H3 | Transform | Months 9–12+

No use cases fall into H3 for this engagement. The seven highest-ranked use cases landed in H1 and H2, and no transformation candidate met all thresholds for the reserved seat, so your first twelve months center on execution rather than a transformation bet. That is a sequencing result, not a ceiling on ambition. Transformation candidates become eligible once H2 data foundations mature.

Investment Envelope: Not applicable — no H3 use cases in this shortlist **Value Range:** Not applicable

SECTION 5 OF 6 · CONTINUED

Program Risk Register

The following risks are program-level concerns that apply across the AI transformation portfolio. Each risk is drawn from discovery data, SAVA financial modeling, and use case implementation analysis. Mitigations are actionable and owner-attributed — not generic disclaimers.

Risk	Category	Horizons	Description	Mitigation
Warranty data residency and consolidation delays	Data Governance	H1, H2	Claim, repair-history, and coverage data sit in separate systems and regional formats, and cross-border residency rules apply across your 6+ countries. This is the long pole for warranty adjudication and defect prediction.	*Director, Warranty Operations confirms data lineage and regional residency constraints for the pilot claim category before Week 3, and sets harmonized failure coding before defect prediction starts.*
Cautious change readiness limits pilot adoption	Change Management	H1, H2	You rate change readiness as Cautious and report difficulty retaining skilled staff. Copilots and automation deliver the 5% productivity target only if pilot users actually adopt them.	*Director, Engineering Productivity captures baseline metrics and runs enablement sessions for pilot developers in Weeks 5–6, then reports adoption and satisfaction at the Month 4 gate.*
Multi-jurisdiction AI regulation across 6+ countries	Regulatory	H1, H2	The EU AI Act product-embedded deadline is now August 2028, China holds manufacturers liable for automated-driving violations, and automated claim decisions need auditability. Recruiting use adds employee privacy exposure.	*Before H2 funding is released at the Month 8 gate, map EU AI Act and China requirements to each pilot and confirm audit logging exists for automated claim decisions.*

SECTION 5 OF 6 · CONTINUED

Program Risk Register (continued)

Risk	Category	Horizons	Description	Mitigation
Source code IP exposure to hosted AI	Technical	H1	Exposing proprietary vehicle and software source code to a hosted AI service requires security and intellectual-property review that may run longer than technical setup, gating the first pilot.	*Chief Information Security Officer completes the review in Weeks 1–3 and limits the pilot to non-restricted repositories until approvals close.*
Plant data variation slows predictive maintenance scale-up	Technical	H2	Plants differ in equipment, sensors, historians, and maintenance systems across countries, so reliable data pipelines are the critical path and the 56-week timeline is hard to compress.	*VP, Manufacturing Operations and Maintenance selects the three pilot plants and confirms data residency and network security requirements in Weeks 1–8, before pipeline build begins.*

Risk assessment performed in SPARKle simulation mode. All risks are directional — severity and ownership would be confirmed through further validation, such as a Day 2 Technical Deep Dive. Feasibility scores are estimated [Est.].

SECTION 6 OF 6

First 90 Days + Closing Recommendation

The following five actions are specific, sequenced, and attributable. Each has a named owner role and a week-level deadline. Executing these five actions in the order listed is the critical path to a successful H1 launch.

#	Action	Owner	Timeline
1	Complete security, IP, and data-protection review and set repository scope for engineering copilot pilot teams.	Information Security Reviewer	**Weeks 1–3**
2	Select high-volume request types, consolidate finance policy content, and define access rules and review workflow for the finance copilot.	Director, Finance Shared Services	**Weeks 1–4**
3	Choose the back-office pilot process, map current steps and exceptions, and set baseline cycle-time and cost metrics.	Director, Finance Shared Services Operations	**Weeks 1–3**
4	Select the pilot warranty claim category, document adjudication rules, and measure baseline processing time, cost, and consistency.	Director, Warranty Operations	**Weeks 1–3**
5	Consolidate and standardize claim, repair-history, and coverage data for the pilot warranty claim category.	Director, Warranty Operations	**Weeks 3–10**

Closing Recommendation

Mike — based on this validated roadmap, The Holistyx Group recommends proceeding with the H1 Quick Win Package:

Engineering and Software Development Copilot + Back-Office Workflow Automation + Automated Warranty Claims Adjudication + Finance and Shared Services Productivity Copilot

— with Consolidate warranty claim, repair-history, and coverage data into one dataset, in parallel with the pilots.

Investment: $19.1M – $30.9M **3-Year Value:** $128.24M – $276.72M **Avg ROI:** **8.1x average across these 4 Quick Win use cases.**

To discuss next steps and receive your Statement of Work, contact The Holistyx Group at info@theholistyxgroup.com

About The Holistyx Group

The Holistyx Group is a boutique AI transformation consultancy specializing in helping enterprise organizations move from AI enthusiasm to funded, implementation-ready execution. Our proprietary SPARK™ methodology combines structured discovery, financial modeling, and sequenced roadmapping to deliver board-ready AI investment strategies. We simplify the complexity of change so leadership can act with confidence.

theholistyxgroup.com | SPARK™ SCOPE™ VAF™ SAVA™

APPENDIX A

SAVA™ v2 Financial Model

Portfolio Financial Summary

All figures represent 3-year Notional-level estimates. ROI Conservative = Low Value ÷ High Investment. ROI Optimistic = High Value ÷ Low Investment.

ID	Use Case	Hz	Val Low	Val Mid	Val High	Inv Range	ROI Low	ROI High	Payback
H1 — Foundation & Quick Wins — Months 1–4									
UC-013	Engineering and Software Development Copilot	H1	$73.03M	$97.38M	$121.72M	$7.00M - $9.80M	7.5x	17.4x	3.1 mo
UC-006	Back-Office Workflow Automation	H1	$7.52M	$18.80M	$30.09M	$3.00M - $5.20M	1.4x	10.0x	7.8 mo
UC-007	Automated Warranty Claims Adjudication	H1	$19.48M	$48.69M	$77.90M	$5.50M - $9.50M	2.1x	14.2x	5.5 mo
UC-012	Finance and Shared Services Productivity Copilot	H1	$28.21M	$37.61M	$47.01M	$3.60M - $6.40M	4.4x	13.1x	4.8 mo
H2 — Scale & Expansion — Months 5–8									
UC-015	Recruiting and Onboarding Assistant	H2	$19.39M	$25.86M	$32.32M	$9.00M - $15.00M	1.3x	3.6x	16.7 mo
UC-002	Early Defect and Warranty Cost Prediction	H2	$16.74M	$25.11M	$33.47M	$14.70M - $22.00M	0.8x	2.3x	26.3 mo
UC-001	Predictive Maintenance Expansion Across Plants	H2	$16.74M	$25.11M	$33.47M	$17.00M - $30.00M	0.6x	2.0x	33.7 mo
PORTFOLIO TOTAL | All Horizons			**$181.11M**	**$278.56M**	**$375.98M**	$59.8M - $97.9M	1.9x	6.3x	—

SAVA™ v2 (Solution-Area Value Anchoring). All values are Notional-level [Est.] — directional for prioritization and board framing. Investment-grade projections require further validation beyond this report's scope.

APPENDIX A · CONTINUED

Payback Period Analysis

Payback = Investment Mid ÷ Annual Value · Annual Value = 3-Year Value Mid ÷ 3 · Investment Mid = (Low + High) ÷ 2

ID	Use Case	Hz	Conf.	Val Mid	Annual Val	Inv Range	Payback (mo)	Payback (yr)
H1 — Foundation & Quick Wins — Months 1–4								
UC-013	Engineering and Software Development Copilot	H1	80%	$97.38M	$32.46M	$7.00M - $9.80M	3.1 mo	0.26 yr
UC-006	Back-Office Workflow Automation	H1	80%	$18.80M	$6.27M	$3.00M - $5.20M	7.8 mo	0.65 yr
UC-007	Automated Warranty Claims Adjudication	H1	80%	$48.69M	$16.23M	$5.50M - $9.50M	5.5 mo	0.46 yr
UC-012	Finance and Shared Services Productivity Copilot	H1	80%	$37.61M	$12.54M	$3.60M - $6.40M	4.8 mo	0.4 yr
H2 — Scale & Expansion — Months 5–8								
UC-015	Recruiting and Onboarding Assistant	H2	55%	$25.86M	$8.62M	$9.00M - $15.00M	16.7 mo	1.39 yr
UC-002	Early Defect and Warranty Cost Prediction	H2	55%	$25.11M	$8.37M	$14.70M - $22.00M	26.3 mo	2.19 yr
UC-001	Predictive Maintenance Expansion Across Plants	H2	55%	$25.11M	$8.37M	$17.00M - $30.00M	33.7 mo	2.81 yr

APPENDIX A · CONTINUED

A. Horizon Confidence Ceilings

Horizon	Confidence	Description
H1 — Foundation & Quick Wins	**80%**	Conservative — near-term, mature data
H2 — Scale & Expansion	**55%**	Moderate — requires H1 data infrastructure
H3 — Transform	**35%**	Directional — dependent on H1 and H2 outcomes

C. ROI Calculation Methodology

Metric	Formula	Description
ROI Conservative	Low Value ÷ High Investment	Worst-case value against worst-case cost
ROI Mid	Mid Value ÷ Mid Investment	Expected value against expected cost
ROI Optimistic	High Value ÷ Low Investment	Best-case value against best-case cost
Payback Period	Investment Mid ÷ (Value Mid ÷ 3)	Months to recover investment at mid-scenario annual value
Value Basis	3-Year Undiscounted Cumulative	No discount rate applied — Notional level

SAVA™ v2 is a proprietary methodology of The Holistyx Group. All values are Notional-level estimates — directional for board framing, not capital commitment. Investment-grade projections require further validation beyond this report's scope.

B. Industry Benchmark Ranges

Solution Area	Low %	High %	Source
Predict	**5%**	**10%**	McKinsey, When Warranty Costs Rival R&D Spend: Remaking Vehicle Quality with AI, 2026 — warranty cost reduction
Automate	**4%**	**16%**	McKinsey, Artificial Intelligence: Automotive's New Value-Creating Engine, 2018 — cost reduction in procurement, manufacturing and supply chain
Employees	**15%**	**25%**	McKinsey, AI Is Already Rewiring the Aftermarket and Services, 2026 — technician labor and nonproductive time reduction
Customers	**10%**	**13%**	McKinsey, How to Win Tomorrow's Car Buyers, 2019, and McKinsey, Artificial Intelligence: Automotive's New Value-Creating Engine, 2018 — marketing and sales cost reduction
Magic Wand	**5%**	**9%**	IBM Institute for Business Value, AI in the Automotive Industry, 2025 — share of total OEM revenue attributable to AI

Solution Area definitions appear on page 30 (Input 1 — Solution Area Anchor).

APPENDIX A · CONTINUED

Horizon Investment & Value Plan

Three Horizons Summary — Investment envelopes and value ranges from SAVA™ v2 model

Hz	Horizon	Timeframe	Conf.	Value Range	Investment Envelope	Avg ROI
H1	Foundation & Quick Wins	Months 1–4	80%	$128.24M – $276.72M	$19.1M – $30.9M	~8.1x
H2	Scale & Expansion	Months 5–8	55%	$52.87M – $99.26M	$40.7M – $67M	~1.4x
H3	Transform	Months 9–12+	35%	Not applicable — no H3 use cases in this shortlist	Not applicable — no H3 use cases in this shortlist	~N/A
PORTFOLIO TOTAL				$181.11M – $375.98M	$59.8M – $97.9M	1.9x – 6.3x

APPENDIX A · CONTINUED

About the SAVA™ Model

All financial projections in this report were produced using THG's proprietary SAVA™ v2 (Solution-Area Value Anchoring) financial modeling methodology. SAVA is purpose-built for AI transformation engagements where traditional ROI models fail — specifically, where the value of AI is diffuse across multiple operational functions and cannot be reduced to a single cost-line calculation.

SAVA produces defensible, board-ready financial ranges by anchoring estimates to three independent inputs that combine to determine the final value projection for each use case.

The Three Inputs

Input 1 — Solution Area Anchor

Every AI use case in this portfolio is assigned to one of five Solution Areas: Predict, Automate, Employees, Customers, or Magic Wand. Each Solution Area has a published benchmark range derived from current industry-specific AI adoption research calibrated to Ford Motor Company's actual sector — see the Industry Benchmark Ranges table in this appendix for the specific source cited per Solution Area.

Solution Area	What It Covers
Predict	Forecasting, anomaly detection, and predictive analytics — anticipating outcomes before they happen (demand, risk, churn, failure).
Automate	Workflow and process automation — removing manual, repetitive work from existing operational processes.
Employees	Internal productivity tools — AI that helps staff work faster or better (search, drafting, internal knowledge, decision support).
Customers	Customer-facing AI — service, personalization, and next-best-experience applications aimed at external users.
Magic Wand	Transformative or novel use cases that don't fit the other four categories — new capabilities or business models AI makes possible for the first time.

APPENDIX A · CONTINUED

Input 2 — Horizon Ceiling

Each use case is assigned to a deployment horizon — H1, H2, or H3 — based on data readiness, technical complexity, and organizational readiness. The horizon assignment applies a confidence ceiling to the value projection, preventing the model from projecting full benchmark value for use cases that carry meaningful execution risk or data dependency.

Input 3 — Workshop Scores

Each use case is scored on two axes during the SPARK engagement: Strategic Impact (1–10) and Implementation Feasibility (1–10). In SPARKle mode, these scores are estimated by THG from the client intake profile and sector benchmarks rather than generated in a live workshop — all estimated scores are tagged [Est.] and would be confirmed through further validation, such as a Day 2 Technical Deep Dive.

Composite Score = (Impact × 0.40) + (Feasibility × 0.25) + (Strategic Alignment × 0.20) + (Time to Value × 0.15). This is SPARK's default weighting, applied unless Ford Motor Company specified custom priorities during discovery. This score determines each use case's ranking in the Full Scored Portfolio (Appendix D); raw Impact and Feasibility scores alone determine placement in the Prioritization Matrix.

Strategic Alignment (1–10) reflects how well a use case fits executive visibility, regulatory deadlines, existing technology investment, or priorities Ford Motor Company stated during discovery — it is scored independently of Impact and Feasibility. **Time to Value** (1–10) is one of the four weighted sub-dimensions inside the Feasibility Score itself, then carried into the Composite Score a second time as its own factor, since speed to a working pilot weighs heavily in how Ford Motor Company would realistically prioritize this work.

The composite formula is weighted rather than a simple product of Impact and Feasibility because a straight multiplication can let one axis dominate — for example, a high-Impact idea with modest Feasibility can outrank a well-rounded, executive-sponsored use case purely on scale. Weighting Impact most heavily while still crediting Strategic Alignment and Time to Value keeps the ranking anchored to what a client can realistically fund and deliver in the near term, not just what scores highest on paper.

APPENDIX A · CONTINUED

How the Value Range Is Calculated

Low scenario — Conservative adoption rate × Solution Area benchmark low end × Horizon ceiling applied.

Mid scenario — Midpoint adoption rate × Solution Area benchmark midpoint × Horizon ceiling applied.

High scenario — Full adoption rate × Solution Area benchmark high end × Horizon ceiling applied.

All values represent undiscounted 3-year cumulative value. No discount rate is applied at this stage, consistent with Notional-level estimation practice.

How ROI Is Calculated

- **Conservative ROI** = 3-Year Value Low ÷ Total Investment High
- **Optimistic ROI** = 3-Year Value High ÷ Total Investment Low
- **Mid ROI** = 3-Year Value Mid ÷ Total Investment Mid

Payback period is calculated at the Mid scenario: Total Investment Mid ÷ Annual Value Mid (3-Year Value Mid ÷ 3).

Confidence and Limitations

These projections are Notional-level estimates — directional for prioritization and board framing, not capital commitment figures.

Benchmark Sources

Source	Application in This Model
Industry-specific AI benchmark research	Solution Area value benchmark ranges — see the Industry Benchmark Ranges table earlier in this appendix for the specific study cited per Solution Area
THG SAVA™ v2 Rate Card	Investment range calculations
THG Engagement History	Adoption curve and payback period assumptions

APPENDIX B

About SPARK™: The AI Transformation Methodology

What Is SPARK™?

SPARK™ is The Holistyx Group's proprietary AI transformation methodology — a structured, repeatable process for helping enterprise organizations move from AI enthusiasm to funded, implementation-ready execution. SPARK was designed to solve the boardroom funding problem: most AI initiatives fail not in production, but because they lack a structured process for translating ambition into a defensible, sequenced investment portfolio that leadership will actually fund. SPARK was developed and refined over hundreds of live client consulting engagements, and SPARKle applies that same proven methodology through an AI-assisted delivery mode built for rapid turnaround.

The Five Phases

Phase 1 — Discovery

Structured discovery builds a complete picture of the client's strategic context, AI maturity, data readiness, and opportunity landscape. Discovery produces a Pre-Flight Briefing — an internal intelligence document that synthesizes market research, competitive positioning, and client-specific context into workshop facilitation intelligence.

Phase 2 — Use Case Generation

Using the Imagine If™ framework, SPARK generates a comprehensive portfolio of AI use case candidates across five Solution Areas: Predict, Automate, Employees, Customers, and Magic Wand. A well-run SPARK use case generation session produces 15–25 candidates.

Phase 3 — Prioritization

SPARK applies a dual-axis Impact/Feasibility scoring model to rank every use case candidate, producing a four-quadrant Prioritization Matrix that separates Quick Wins from Strategic Investments and Experimental ideas.

Phase 4 — Financial Modeling

SPARK applies the SAVA™ v2 financial model to quantify 3-year value potential, investment range, and ROI for every shortlisted use case, producing figures that can withstand CFO scrutiny.

Phase 5 — Roadmap and Executive Readout

SPARK synthesizes the prioritized portfolio and financial model into a Three Horizons AI Roadmap with defined investment envelopes, Go/No-Go gates, and a First 90 Days action plan.

SPARK Engagements at a Glance

Format	Delivery	Output	Best For
SPARKle Free	24 hours	Industry AI Snapshot — sector benchmarks and top use case hypotheses	First exposure to AI opportunity landscape
SPARKle Express	24 hours	Full portfolio, financial model, and 12-month roadmap	Leadership alignment and board framing
SPARKle Pro	48 hours	Everything in Express plus use case cards, QA review, and implementation guidance	SOW scoping and H1 investment authorization
SPARK Live	2 workshop days	Confirmed portfolio with live-scored feasibility, validated financial model, and facilitated stakeholder alignment	Investment-grade capital commitment decisions

Ready to Go Deeper?

If you would like to progress from this SPARKle report to a confirmed investment plan, the next step is a **Day 2 Technical Deep Dive** — a structured working session with your engineering, data, and operations leads that validates feasibility assumptions, confirms data readiness, and produces an Investment-grade financial model suitable for capital budget approval.

Upon confirmation, The Holistyx Group will deliver a Statement of Work covering the Day 2 scope within 48 hours.

To discuss next steps, contact The Holistyx Group at info@theholistyxgroup.com

theholistyxgroup.com | SPARK™ SCOPE™ VAF™ SAVA™ SPARKle™

SPARK™ is a proprietary methodology of The Holistyx Group. Simplifying the complexity of change.

APPENDIX C

Use Case Dossiers — Shortlisted Portfolio

Compact reference dossiers for all 7 shortlisted use cases. Financial figures from SAVA™ v2. All financial fields are estimated (Est.) and would be confirmed through further validation, such as a Day 2 Technical Deep Dive.

UC-013 · EMPLOYEES · H1

Engineering and Software Development Copilot

Quick Win

Score: 8.2/10

CHALLENGE

Ford Motor Company's engineering and software development teams spend significant time on repetitive coding, code review, documentation, and test creation across large, complex codebases, and AI tooling gains are not captured consistently.

PILOT: ENGINEERING COPILOT PILOT — SELECTED SOFTWARE ENGINEERING TEAMS

Deploy the copilot to selected software engineering teams on non-restricted repositories for code suggestions, test generation, documentation drafting, and code-change summaries. Measure productivity against a pre-pilot baseline and capture developer feedback.

Success: Measured reduction in time-to-complete for routine engineering tasks (such as pull-request cycle time) against baseline, with a developer satisfaction score of at least 70% favorable.

Timeline: About 8 weeks

Owner: Director, Engineering Productivity

FINANCIAL SUMMARY

3-Yr Value Low	**$73.03M**
3-Yr Value Mid	**$97.38M**
3-Yr Value High	**$121.72M**
Investment Range	**$7M – $9.8M**
ROI Range	**7.5x - 17.4x**
Payback Period	**3.1 months**

HOW IT GETS DONE

TYPE **Assistant** | EFFORT **Low** | PATH **Self-serve**

TIME BACK **3,000 to 6,000 hours a week**

Plan on about 8 weeks, most of it spent on security review, repository scoping, and developer onboarding before a pilot with selected engineering teams.

1. Weeks 1-3: Security and policy review
2. Weeks 3-5: Configuration and access setup
3. Weeks 5-6: Pilot launch and training
4. Weeks 6-8: Pilot run and measurement

APPENDIX C · CONTINUED

UC-006 · AUTOMATE · H1

Back-Office Workflow Automation

Quick Win

Score: 7.625/10

CHALLENGE

Ford Motor Company's back-office functions such as finance operations, procurement, and administrative processing rely on repetitive, rules-based tasks with manual data entry, reconciliation, and routing across systems, slowing cycle times globally.

PILOT: BACK-OFFICE AUTOMATION PILOT — ONE HIGH-VOLUME PROCESS

Automate one high-volume, rules-based process, such as invoice-to-payment matching, including document extraction, validation, and exception routing. Run in parallel with the manual process on a transaction subset to measure accuracy, cycle time, and effort saved.

Success: At least a 40% reduction in manual handling time per transaction on the pilot process, with accuracy matching or exceeding the manual baseline.

Timeline: About 16 weeks

Owner: Director, Finance Shared Services Operations

FINANCIAL SUMMARY

3-Yr Value Low	**$7.52M**
3-Yr Value Mid	**$18.8M**
3-Yr Value High	**$30.09M**
Investment Range	**$3M – $5.2M**
ROI Range	**1.4x - 10.0x**
Payback Period	**7.8 months**

HOW IT GETS DONE

TYPE **Automation** EFFORT **High**

PATH **Partner-led** TIME BACK **150 to 300 hours a week**

Plan on about 16 weeks, most of it spent mapping the chosen process, integrating finance and document systems, and tuning exception handling before a pilot on one high-volume workflow.

1. Weeks 1-3: Process selection and mapping
2. Weeks 3-8: Data and system access preparation
3. Weeks 6-13: Workflow build and integration
4. Weeks 13-16: Parallel run and pilot

APPENDIX C · CONTINUED

UC-007 · AUTOMATE · H1

Automated Warranty Claims Adjudication

Quick Win

Score: 7.6/10

CHALLENGE

Ford Motor Company processes a very large volume of dealer and service warranty claims. Adjudication relies on manual review of claim details, repair histories, and coverage rules, which is slow, labor-intensive, and inconsistent across regions.

PILOT: WARRANTY CLAIMS ADJUDICATION PILOT — SINGLE CLAIM CATEGORY (SHADOW MODE)

Run an automated adjudication model in shadow mode on one high-volume, routine claim category. The model recommends approve, deny, or refer decisions compared with human outcomes, with no automatic payments.

Success: At least 85% agreement between model recommendations and senior adjudicator decisions on the pilot category, with a projected reduction of at least 40% in manual review time.

Timeline: About 20 weeks

Owner: Director, Warranty Operations

FINANCIAL SUMMARY

3-Yr Value Low	**$19.48M**
3-Yr Value Mid	**$48.69M**
3-Yr Value High	**$77.9M**
Investment Range	**$5.5M – $9.5M**
ROI Range	**2.1x - 14.2x**
Payback Period	**5.5 months**

HOW IT GETS DONE

TYPE **Agent** | EFFORT **High** | PATH **Partner-led**

TIME BACK **80 to 200 hours a week**

Plan on about 20 weeks, most of it spent consolidating claim and part-history data and encoding coverage rules before a shadow-mode pilot on one claim category.

1. Weeks 1-3: Scoping and baseline
2. Weeks 3-10: Data consolidation
3. Weeks 8-16: Model and rules build
4. Weeks 16-20: Shadow-mode pilot

APPENDIX C · CONTINUED

UC-012 · EMPLOYEES · H1

Finance and Shared Services Productivity Copilot

Quick Win

Score: 7.525/10

CHALLENGE

Ford Motor Company's finance and shared-services teams spend significant time reconciling transactions, answering policy questions, drafting routine communications, and assembling reports across fragmented tools. At 100,000+ employees, small inefficiencies become large capacity drains.

PILOT: SHARED-SERVICES ASSISTANT PILOT — TWO REQUEST TYPES

Pilot the assistant on two high-volume shared-services request types, such as accounts payable inquiries and routine reconciliation summaries. It drafts responses and summaries for staff review, with no autonomous posting or approvals.

Success: Average handling time per pilot request type falls at least 25% versus baseline while maintaining reviewer-approved accuracy above an agreed threshold.

Timeline: About 10 weeks

Owner: Director, Finance Shared Services

FINANCIAL SUMMARY

3-Yr Value Low	**$28.21M**
3-Yr Value Mid	**$37.61M**
3-Yr Value High	**$47.01M**
Investment Range	**$3.6M – $6.4M**
ROI Range	**4.4x - 13.1x**
Payback Period	**4.8 months**

HOW IT GETS DONE

TYPE **Assistant** | EFFORT **Medium** | PATH **THG-led**

TIME BACK **150 to 300 hours a week**

Plan on about 10 weeks to a running pilot, most of it spent preparing policy content and access controls before a review-based pilot on two request types.

1. Weeks 1-4: Process selection and content preparation
2. Weeks 4-7: Assistant configuration and connection
3. Weeks 7-10: Pilot launch and tuning

APPENDIX C · CONTINUED

UC-015 · EMPLOYEES · H2

Recruiting and Onboarding Assistant

Experimental

Score: 7.375/10

CHALLENGE

Ford Motor Company hires and onboards at a scale where recruiter screening, scheduling, candidate communications, and new-hire paperwork consume large amounts of coordinator time, risking lost candidates and delayed new-hire productivity.

PILOT: RECRUITER SCHEDULING AND NEW-HIRE Q&A PILOT

Pilot interview scheduling, candidate communication drafts, and a new-hire question assistant for one set of requisitions and one onboarding cohort. No candidate ranking or automated hiring decisions; recruiters review all communications.

Success: Recruiter coordination time per hire falls at least 20%, with new-hire onboarding satisfaction improving versus the prior cohort baseline.

Timeline: About 20 weeks

Owner: Director, Talent Acquisition Operations

FINANCIAL SUMMARY

3-Yr Value Low	**$19.39M**
3-Yr Value Mid	**$25.86M**
3-Yr Value High	**$32.32M**
Investment Range	**$9M – $15M**
ROI Range	**1.3x - 3.6x**
Payback Period	**16.7 months**

HOW IT GETS DONE

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

TIME BACK **100 to 200 hours a week**

Plan on about 20 weeks to a running pilot, most of it spent integrating talent and HR systems and completing privacy and fairness review before a limited recruiter pilot.

1. Weeks 1-6: Privacy and fairness review
2. Weeks 3-12: System integration
3. Weeks 10-16: Assistant configuration
4. Weeks 16-20: Pilot launch

APPENDIX C · CONTINUED

UC-002 · PREDICT · H2

Early Defect and Warranty Cost Prediction

Quick Win

Score: 7.25/10

CHALLENGE

Ford Motor Company learns about emerging defects largely after vehicles reach customers, when warranty claims accumulate. Warranty, field service, and quality signals sit in separate systems, making early pattern detection hard.

PILOT: EARLY WARRANTY SIGNAL SHADOW PILOT — LIMITED COMPONENTS

Run predictive defect-cluster and warranty cost alerts in shadow mode on limited components in one vehicle program, comparing alerts against actual claim development. Alerts inform analysts but trigger no automatic actions.

Success: Models flag emerging defect clusters at least four weeks earlier than current detection on back-tested and live cases, at an agreed precision threshold.

Timeline: About 30 weeks

Owner: Director, Global Warranty and Quality Analytics

FINANCIAL SUMMARY

3-Yr Value Low	**$16.74M**
3-Yr Value Mid	**$25.11M**
3-Yr Value High	**$33.47M**
Investment Range	**$14.7M – $22M**
ROI Range	**0.8x - 2.3x**
Payback Period	**26.3 months**

HOW IT GETS DONE

TYPE **Assistant** EFFORT **High** PATH **Partner-led**

Plan on about 30 weeks to a running pilot, most of it spent harmonizing warranty and quality data across systems and regions before a shadow-mode pilot on a limited component set.

1. Weeks 1-5: Data scoping and residency review
2. Weeks 4-16: Data consolidation and cleansing
3. Weeks 14-24: Model development and validation
4. Weeks 24-30: Shadow-mode pilot

APPENDIX C · CONTINUED

UC-001 · PREDICT · H2

Predictive Maintenance Expansion Across Plants

Strategic Investment

Score: 7.25/10

CHALLENGE

Ford Motor Company has proven predictive maintenance in selected settings, but plants differ in equipment, sensors, and data practices. Unplanned downtime persists where models are not deployed, and plant-by-plant extension is slow and inconsistent.

PILOT: THREE-PLANT PREDICTIVE MAINTENANCE EXPANSION PILOT

Extend predictive maintenance to critical equipment classes at three new plants across two countries using a standardized pipeline and model pattern. Alerts run alongside existing practice; planners decide actions. No automated work orders.

Success: Unplanned downtime on pilot equipment classes falls at least 15% versus baseline, with alert precision above an agreed threshold.

Timeline: About 56 weeks

Owner: VP, Manufacturing Operations and Maintenance

FINANCIAL SUMMARY

3-Yr Value Low	**$16.74M**
3-Yr Value Mid	**$25.11M**
3-Yr Value High	**$33.47M**
Investment Range	**$17M – $30M**
ROI Range	**0.6x - 2.0x**
Payback Period	**33.7 months**

HOW IT GETS DONE

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

Plan on about 56 weeks to a running multi-plant pilot, most of it spent building plant data pipelines and handling multi-country integration before a pilot at three plants.

1. Weeks 1-8: Plant selection and data assessment
2. Weeks 6-16: Standard pattern design
3. Weeks 12-40: Plant data pipeline build
4. Weeks 32-50: Model adaptation and validation
5. Weeks 50-56: Multi-plant pilot

APPENDIX D · PRIORITIZATION MAP

Impact vs. Feasibility — Full Portfolio

The 20 highest-ranked of the 25 use cases evaluated, plotted by Impact score (Y axis) and Feasibility score (X axis). Horizon zones divide the portfolio into sequencing bands. Shortlisted use cases are shown in solid color; evaluated-only use cases appear lighter. Feasibility scores are estimated [Est.] in SPARKle simulation mode.

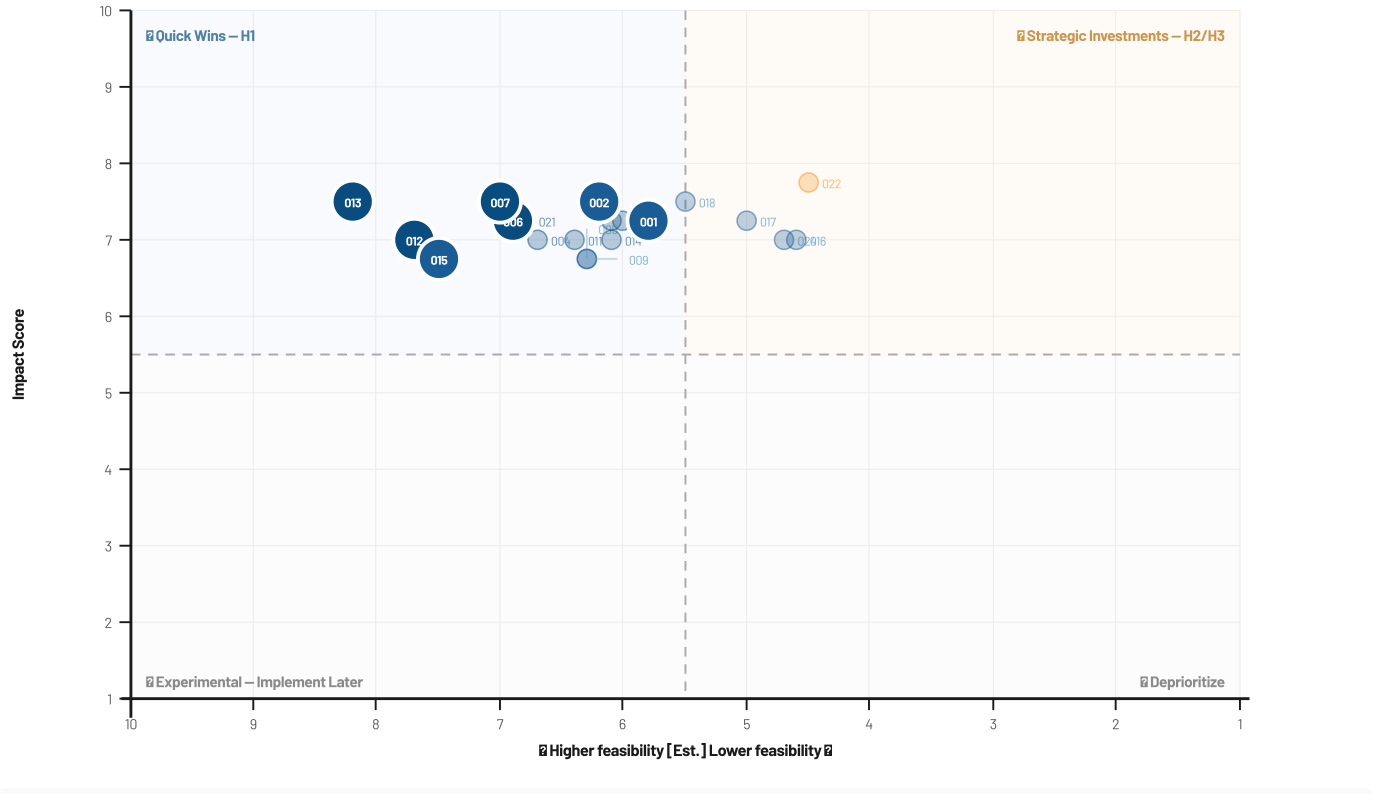

Shortlisted use cases (solid circles, 7 total) labeled by ID. Evaluated-only use cases (faded dots) shown for portfolio context — full scored portfolio in Appendix D. X axis reversed: feasibility reads high (left) to low (right). Use cases sharing identical scores are offset for legibility. Feasibility scores are estimated [Est.] in SPARKle simulation mode.

SPARK™ Three Horizons — Delivery Sequencing

H1 — FOUNDATION & QUICK WINS

Months 1–4

High feasibility, mature data, deployable now with minimal dependencies. Builds AI credibility and data infrastructure for later horizons.

H2 — SCALE & EXPANSION

Months 5–8

Moderate feasibility. Builds on H1 data infrastructure and process changes. Broader integration and change management required.

H3 — TRANSFORM

Months 9–12+

Highest impact, lowest near-term feasibility. Platform-level bets that require H1 and H2 foundations to be fully operational first.

APPENDIX D

Full Scored AI Use Case Portfolio

The 20 highest-ranked of the 25 use cases evaluated during the SPARK™ Imagine If™ brainstorm session, ranked by composite score. Shortlisted use cases are highlighted.

Scoring methodology: Composite Score = (Impact × 0.40) + (Feasibility × 0.25) + (Strategic Alignment [Align] × 0.20) + (Time to Value [TTV] × 0.15) — SPARK's default weighting, applied unless the client specified custom priorities during discovery. See Appendix A, Input 3, for the full rationale. Feasibility scores are estimated (Est.) in SPARKle mode and would be confirmed through further validation, such as a Day 2 Technical Deep Dive. Use cases sorted by composite score descending.

Abbreviations: Hz = Horizon | Feasib. = Feasibility | Align = Strategic Alignment | TTV = Time to Value

ID	Use Case	Area	Hz	Quadrant	Impact	Feasib.	Align	TTV	Score	Status
UC-013	**Engineering and Software Development Copilot**	Employees	**H1**	Quick Win	**7.5**	**8.2**	**9**	**9**	**8.2**	**Shortlisted**
UC-006	**Back-Office Workflow Automation**	Automate	**H1**	Quick Win	**7.25**	**6.9**	**9**	**8**	**7.625**	**Shortlisted**
UC-007	**Automated Warranty Claims Adjudication**	Automate	**H1**	Quick Win	**7.5**	**7**	**9**	**7**	**7.6**	**Shortlisted**
UC-012	**Finance and Shared Services Productivity Copilot**	Employees	**H1**	Quick Win	**7**	**7.7**	**8**	**8**	**7.525**	**Shortlisted**
UC-015	**Recruiting and Onboarding Assistant**	Employees	**H2**	Experimental	**6.75**	**7.5**	**8**	**8**	**7.375**	**Shortlisted**
UC-002	**Early Defect and Warranty Cost Prediction**	Predict	**H2**	Quick Win	**7.5**	**6.2**	**9**	**6**	**7.25**	**Shortlisted**
UC-001	**Predictive Maintenance Expansion Across Plants**	Predict	**H2**	Strategic Investment	**7.25**	**5.8**	**10**	**6**	**7.25**	**Shortlisted**
UC-011	Unified Enterprise Knowledge and Data Search	Employees	H1	Quick Win	7	6.4	9	7	7.25	Evaluated
	Cleared both thresholds but ranked 8th on composite; viable Quick Win for next cycle or H1 expansion once shortlisted work is underway.									

APPENDIX D · CONTINUED

Abbreviations: Hz = Horizon | Feasib. = Feasibility | Align = Strategic Alignment | TTV = Time to Value

ID	Use Case	Area	Hz	Quadrant	Impact	Feasib.	Align	TTV	Score	Status
UC-004	Tariff and Input Cost Impact Forecaster	Predict	H1	Quick Win	7	6.7	8	7	7.125	Evaluated
	Cleared both thresholds but ranked 9th on composite; viable cost-efficiency opportunity for next-cycle addition once shortlisted H1 items are in motion.									
UC-003	Supply Chain Demand and Disruption Forecasting	Predict	H1	Quick Win	7.25	6	9	6	7.1	Evaluated
	Cleared both thresholds (feasibility at the 6.0 floor) but ranked 10th; natural candidate for the next prioritization cycle.									
UC-021	Ford Credit Customer Servicing Assistant	Customers	H1	Quick Win	7.25	6.8	7	7	7.05	Evaluated
	Cleared both thresholds but ranked 11th on composite; strong Ford Credit servicing opportunity once current H1 capacity frees up.									
UC-005	Critical Skills Attrition and Gap Predictor	Predict	H2	Experimental	6.75	6.3	8	6	6.775	Evaluated
	Feasibility clears the floor but impact (6.75) is narrower than shortlisted seven; value is real but modest, centered on talent-risk visibility.									
UC-019	Dealer Ordering and Inventory Advisor	Customers	H1	Quick Win	7.25	6.1	7	6	6.725	Evaluated
	Cleared both thresholds but ranked 13th on composite; viable dealer-facing efficiency use case for later-cycle consideration.									
UC-014	Field Service Technician Copilot	Employees	H1	Quick Win	7	6.1	7	6	6.625	Evaluated
	Cleared both thresholds but ranked 14th on composite; legitimate field-service productivity opportunity for future-cycle inclusion.									
UC-018	Proactive Vehicle Health Alerts	Customers	H2	Strategic Investment	7.5	5.5	7	5	6.525	Evaluated
	Strong impact (7.5) but feasibility (5.5) limited by real-time telemetry integration; sequencing after UC-001 would improve readiness.									
UC-009	Global Regulatory Change Monitor	Automate	H2	Experimental	6.75	6.3	6	7	6.525	Evaluated
	Feasibility clears the floor but impact (6.75) is narrower; value is more compliance-adjacent than core cost or productivity impact this cycle.									

APPENDIX D · CONTINUED

Abbreviations: Hz = Horizon | Feasib. = Feasibility | Align = Strategic Alignment | TTV = Time to Value

ID	Use Case	Area	Hz	Quadrant	Impact	Feasib.	Align	TTV	Score	Status
UC-017	Agentic Sales and Service Assistant	Customers	H2	Strategic Investment	7.25	5	8	5	6.5	Evaluated
	Impact solid (7.25) but feasibility (5.0) constrained by agentic orchestration complexity and real-time cross-system integration, limiting near-term readiness.									
UC-022	Closed-Loop Product and Process Intelligence Hub	Predict	H3	Strategic Investment	7.75	4.5	8	4	6.425	Evaluated
	High impact (7.75) but feasibility (4.5) limited; needs the enterprise data foundation UC-024 would establish first.									
UC-024	Enterprise AI-Ready Data Product Layer	Magic Wand	H2	Strategic Investment	7	4.7	9	4	6.375	Evaluated
	Impact solid (7.0) but feasibility (4.7) low; foundational data-governance work needed before downstream agentic and intelligence-hub use cases mature.									
UC-016	Enterprise AI Agent Workforce Operating Model	Magic Wand	H2	Strategic Investment	7	4.6	9	4	6.35	Evaluated
	Impact solid (7.0) but feasibility (4.6) limited by missing agent-governance model and data product layer (UC-024) that must exist first.									