Palmetto Interactive

FROM ASSISTED WORK TO CONTROLLED ACTION

# Agentic AI Systems

AI agents that use your business knowledge and tools to finish real work. With explicit controls, and a named person accountable for every consequential call.

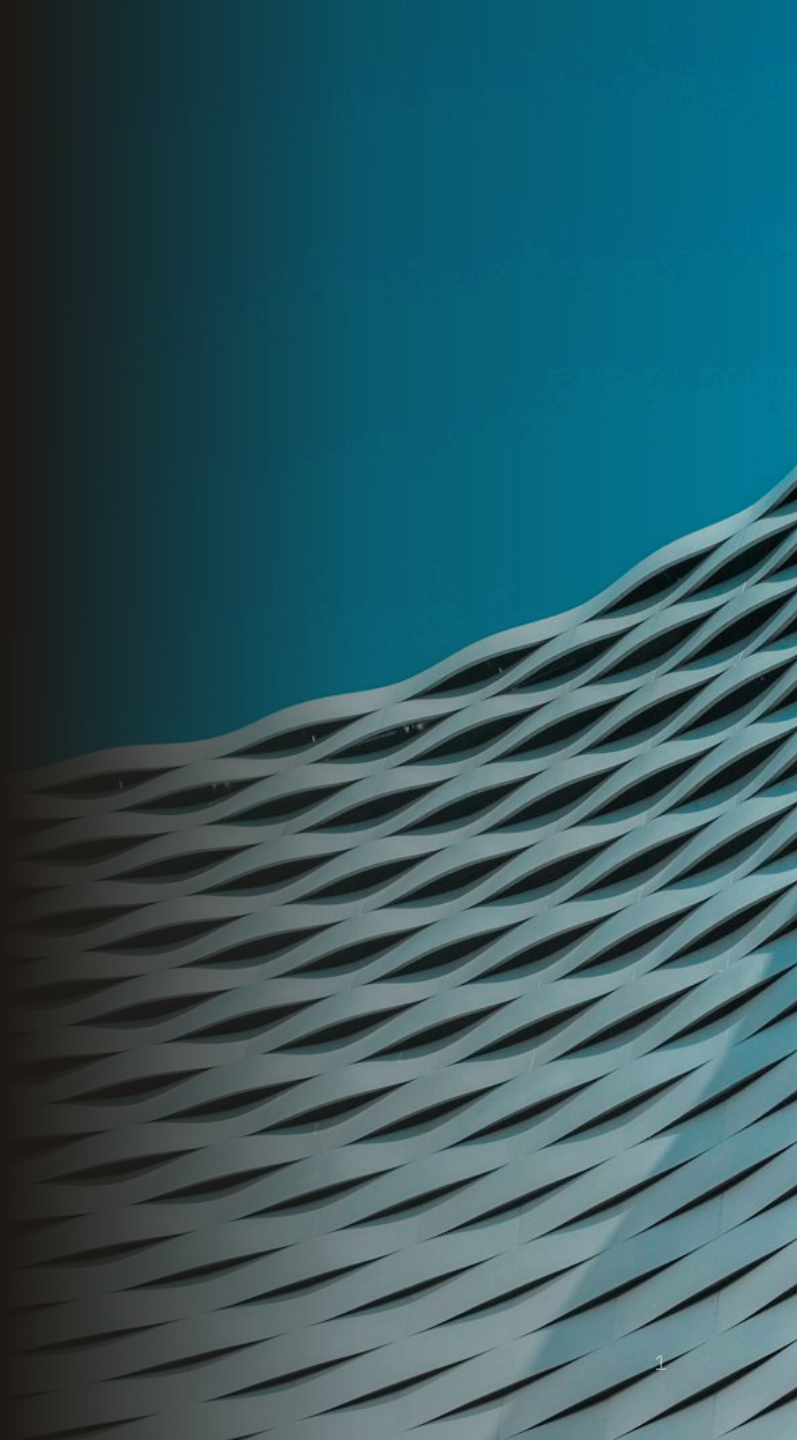

ENGAGEMENT OVERVIEW · AUGUST 2026

Confidential · Palmetto Interactive Services, LLC

WHY THIS CONVERSATION IS HAPPENING

# You’re probably here for one of these

01 A workflow depends on people reading, reconciling, deciding, and updating several systems by hand

02 An existing copilot answers questions but cannot reliably complete the surrounding work

03 Cycle time or backlog is growing while experienced people spend their days on repeatable coordination

04 Leadership wants an evidence-based business case before committing to a larger AI platform

If none of these sounds like you, this probably isn’t the right engagement. We’d rather say that early than scope work nobody needs.

WHAT YOU GET AND WHEN

# Three rungs. Every one is a real stopping point

## 01

Principal only

DIAGNOSTIC · 3 WEEKS

### Agentic Workflow Diagnostic

| | |
|---|---|
| Wk 1 | Workflow & baseline mapping |
| Wk 2 | Automation boundary & value hypothesis |
| Wk 3 | Decision brief & sequencing |

ENDS WITH

You decide go or no-go on a pilot from the decision brief.

## 02

Principal-led team

PILOT · 10 WEEKS

### Production-Shaped Agent Pilot

| | |
|---|---|
| Wk 1–2 | Architecture & tool contracts |
| Wk 3–5 | Build & evaluation harness |
| Wk 6–8 | Integrated pilot with controls |
| Wk 9–10 | Evidence review & decision |

ENDS WITH

A working pilot and real evaluation evidence.

## 03

Principal-led team

SCALE · ~4 MONTHS

### Agent Rollout & Operating Handoff

| | |
|---|---|
| Wk 1–3 | Hardening & control expansion |
| Wk 4–8 | Integration & adoption rollout |
| Wk 9–13 | Runbook & knowledge transfer |
| Wk 14–16 | Measurement & backlog handoff |

ENDS WITH

Production capability handed over with runbooks and named owners.

Every engagement ends with a decision you actually make. Timings are planning ranges. The right path depends on scope, access, risk, and what the evidence shows.

DIAGNOSTIC · 3 WEEKS

# Agentic Workflow Diagnostic

WEEK 1

## Workflow & baseline mapping

Map the current workflow, the systems it touches, and human touch time to establish the baseline.

WEEK 2

## Automation boundary & value hypothesis

Define the prioritized automation boundary and the value hypothesis measured against that baseline.

WEEK 3

## Decision brief & sequencing

Deliver the decision brief, target state, and sequenced recommendations for a pilot go/no-go.

GATE

## Decision point

You decide go or no-go on a pilot from the decision brief. This diagnostic stands on its own. Nothing here assumes a pilot follows.

WHAT YOU RECEIVE

- Workflow map and current-state baseline
- Value hypothesis with explicit assumptions
- Prioritized automation boundary
- Decision brief, target state, and sequenced recommendations

WHO DELIVERS IT

Principal only

3 weeks, ending in the decision point above.

PILOT · 10 WEEKS

# Production-Shaped Agent Pilot

WEEKS 1–2

## Architecture & tool contracts

Design the agent architecture, tool contracts, permission model, and the points where a human decides.

WEEKS 3–5

## Build & evaluation harness

Build against representative data, and construct the test set and evaluation scorecard alongside it.

WEEKS 6–8

## Integrated pilot with controls

Integrate with the agreed systems. Wire up least-privilege access, provenance, and escalation, then run it with real users.

WEEKS 9–10

## Evidence review & decision

Score against release criteria and deliver telemetry, operating lessons, and a go-forward recommendation.

GATE

## Decision point

A working pilot and real evaluation evidence. You decide: scale it, iterate, or stop.

WHAT YOU RECEIVE

- Agent architecture, tool contracts, and permission model
- Documented human decision points
- Representative test set, evaluation scorecard, and release criteria
- A production-shaped agent workflow integrated with agreed systems

WHO DELIVERS IT

Principal-led. We add an integration engineer when the pilot spans several systems, and an evaluation specialist when the scorecard needs real infrastructure.

10 weeks, ending in the decision point above.

SCALE · ~4 MONTHS

# Agent Rollout & Operating Handoff

WEEKS 1–3

## Hardening & control expansion

Harden approval gates and traceability, and build the cost and latency monitoring that scale requires.

PRINCIPAL + MLOPS ENGINEER

WEEKS 4–8

## Integration & adoption rollout

Integrate across multiple systems and teams, running the adoption rollout in parallel rather than after.

INTEGRATION ENGINEERS + ADOPTION LEAD

WEEKS 9–13

## Runbook & knowledge transfer

Build the operating runbook and transfer ownership to the internal team that will run it.

PRINCIPAL + INTEGRATION ENGINEER

WEEKS 14–16

## Measurement & backlog handoff

Validate the measures against the original baseline and hand over the prioritized improvement backlog.

PRINCIPAL

GATE

## Decision point

Production capability handed over with runbooks and named owners. We step back to advisory, or we exit cleanly.

WHAT YOU RECEIVE

- Production capability in the operating environment
- Telemetry, operating runbook, and rollout plan
- Prioritized improvement backlog
- Ownership handoff to a named internal team

WHO DELIVERS IT

Palmetto's principal owns architecture, controls, and your executive readouts. Alongside the principal: integration engineers, an evaluation engineer, and an adoption lead. We size the team to the systems and teams in scope.

~4 months, ending in the decision point above.

WHAT WE NEED FROM YOU

# This doesn’t work without your people

- An executive sponsor and an accountable workflow owner
- Domain experts who can define good decisions, edge cases, and unacceptable outcomes

The most common reason this work stalls isn’t technical. It’s that nobody with authority was free to settle a definition or approve an access request.

HOW SUCCESS IS MEASURED

- Cycle time and human touch time against baseline
- Throughput, backlog, and service-level performance
- Adoption, unit cost, and workflow-specific quality

WHEN THIS ISN’T THE RIGHT FIT

You want the agent making the call instead of a person. Or nobody can get us real data, real access, and someone to validate the output.

ENABLING CAPABILITY

# Digital Product Modernization & Integration

Not every valuable technology problem needs AI. Sometimes a legacy application, a manual handoff, or a missing integration is the real constraint behind the AI outcome you want. We name that work during the diagnostic and price it before the pilot fee is fixed, then deliver it inside the engagement at the standard implementation rate. You won't find it discovered halfway through and quietly absorbed. And if it turns out to be most of the work, we'll say so and quote it as a modernization engagement instead.

## Bring us the hardest question you've got.

If a bounded diagnostic isn't the right next step, we'll tell you in the first conversation. Not the first invoice.

THE ENGAGEMENT PATH

**Diagnostic**

Typically 2–6 weeks

Decision brief, baseline, target state, and sequenced recommendations

**Production-shaped pilot**

Typically 6–12 weeks

Working pilot, evaluation evidence, operating lessons, and go-forward decision

**Scale & rollout**

Often 3–6 months

Production capability, rollout evidence, runbooks, and ownership handoff

matt@palmettointeractive.com
(843) 732-1460
palmettointeractive.com