



Meet Beacon: Your Shop's AI Assistant, Built Into WebMB

What Beacon is, how it works, why your data is safe — and when it arrives.

July 2026 • Topics: AI, Operations, Security, Architecture | Read time: 6 min

IN THIS ISSUE

- 01 The Big Picture:** The shift happening across the ERP industry, and where Beacon fits.
- 02 Owners & Operators:** What Beacon does for your shop. Plain English, no tech jargon.
- 03 Technical Operations:** How it works, the security model, and confirmed launch capabilities.

01 THE BIG PICTURE

The Way You Interact With Your ERP Is Changing

In May 2026, McKinsey published a paper titled "[The End of ERP As We Know It: Five Ways AI Is Disrupting ERP](#)". Their core finding was not about technology—it was about interaction. AI is not just changing how enterprise software is built. It is changing how you use it.

The prediction is specific: instead of navigating menus and running reports, users will ask questions and get answers. The software becomes invisible. The outcome stays.

Every major ERP vendor read that paper. Every one of them is now announcing an AI assistant.

The difference with Beacon is that we did not read the paper and then start building. Beacon was already under development when the paper came out. And it was built specifically for motor repair ERP—not adapted from a generic enterprise template.

"We were not quiet about AI because we were behind. We were quiet because we were working."

— Troy D. Locke, Co-Founder | CEO, Spring Point · EASA 2026

Where Beacon Is Going

The January 2027 launch is the starting point, not the destination.

The first release is focused on one thing: making your data accessible without a report or a phone call. Ask a question, get an answer, backed by your live shop data.

What comes after that is proactive intelligence—Beacon surfacing things you did not know to ask. Flagging at-risk jobs before you walk in Monday morning. Identifying patterns that would never surface in a standard report. Deeper analysis: quoted versus actual job performance, technician productivity, margin by customer, receivables aging, financial trends.

And further out, with your explicit authorization: guided actions. Flagging a job for priority review. Drafting a customer follow-up based on job status. When Beacon takes any action on your behalf, it will go through the same authorization model as everything else in WebMB—same permissions, same access controls, same division scoping. No bypass. Every irreversible action requires your confirmation.

We will communicate each phase clearly before it ships—not as a vague roadmap, but as specific capabilities with honest timelines.

02 OWNERS & OPERATORS

What Beacon Does for Your Shop

Beacon is a chat assistant built directly into the WebMB screen you already use. You type a question. Beacon reads your actual shop data and gives you a direct answer. You do not navigate to it. You do not log into a separate tool. You do not need to learn anything new. If you can type a question, you can use Beacon.

Operational questions Beacon can answer at launch:

QUESTION	WHAT BEACON DOES
"What jobs are stalled this week?"	Scans open jobs for activity gaps and surfaces the ones that have stopped moving, before a customer calls to ask.
"Which jobs are at risk of missing their deadline?"	Flags jobs where the trajectory suggests they will not ship on time, so you can get ahead of it.
"How is the shop performing financially this month?"	A financial summary in a sentence, not a spreadsheet.
"What is the breakdown of activity across open jobs?"	Job activity patterns across the shop floor, on demand.
"What are our open receivables right now?"	AR aging detail on demand, without pulling a report.
"What new jobs came in this week?"	Intake summary across the shop.

WebMB how-to questions Beacon can answer at launch:


- "How do I post a transaction?"
- "What does the WIP report calculate?"
- "Walk me through the end-of-month close process."
- "What does the QM hold status mean?"

These are answered from our curated WebMB knowledge base—documentation, workflows, field definitions, and resolved support cases, all searchable by Beacon in real time.

The launch question set covers the most common daily needs for shop owners, operations managers, and controllers. The question inventory grows with each release—deeper analysis, profitability trends, technician performance, and more follow in subsequent phases.

Watch Beacon in Action

Watch the Beacon Demo: A quick look at Beacon answering real shop questions—live data, no scripts.



▶ springpt.com/wp-content/uploads/2026/07/2026-07-10-Beacon-QuickPeek-Amy-v4.mp4

What Beacon Does Not Do — At Launch

The initial release is read-only by design. Beacon answers questions—it does not post transactions, close jobs, or modify records. This is a deliberate choice: build the foundation on trust before expanding capability.

Future releases will add guided actions—for example, flagging a job for priority review or drafting a customer follow-up. When that happens, every action will go through the same authorization model as WebMB: your session, your menu-level permissions, your division scoping. Beacon will not have an elevated access path, and it will not take any irreversible action without your explicit confirmation.

Beacon also does not replace WebMB screens or act as a persistent dashboard. Answers live in the conversation, not on a page you pin and return to.

If Beacon does not know the answer—because the question falls outside its knowledge base or beyond its current tool set—it says so and routes you to support: [888-382-7835](tel:888-382-7835) or support@springpt.com.

Your Data Is Safe

When people hear "AI has access to my financial data," a fair question follows: how far does that access go?

"Think of Beacon as a knowledgeable colleague who can look things up but only in the same locked filing cabinet you have access to, using lookups our engineering team has reviewed and approved. There is no back door."

Beacon cannot do anything you are not already allowed to do in WebMB. If your account cannot see a division's financial data, Beacon cannot show it to you either. The same permission controls that govern your WebMB access govern Beacon.

Beacon also does not invent its own questions to ask your database. The lookups it performs are written by our engineers, not generated by AI. The AI decides which approved lookup to run and what safe inputs to pass—it cannot reach beyond that defined set.

When Will Beacon Be Available?

Beacon is currently in development. We are targeting a rollout beginning September 2026, with full availability by January 2027 aligned with the broader WebMB modernization milestone announced in Issue 01.

We will begin releasing capabilities ahead of the full launch so customers can start using Beacon before the year is out. If you want to be among the first, email support@springpt.com or contact your account manager now—we will reach out as early access opens.

03 TECHNICAL OPERATIONS

How Beacon Works

Beacon is a RAG (Retrieval-Augmented Generation) architecture with parameterized data access, permission-aware tooling, and security as the first design constraint. Here is an accurate, non-marketing summary of the request flow.

- 1 User submits a question** via the WebMB chat drawer. No new login or context switch required.
- 2 Knowledge base search:** Beacon runs a vector similarity search against a curated set of Markdown articles — WebMB features, workflows, field definitions, and support case history stored as embeddings in pgvector on Aurora Postgres.
- 3 Tool selection and execution:** For data questions, Beacon selects one pre-approved tool from a fixed set and executes it with parameterized queries against the customer SQL Server database. The AI selects the tool and supplies safe input parameters. It does not write SQL.
- 4 Response synthesis:** AWS Bedrock (Claude Sonnet) synthesizes a response from KB results and tool output. Every conversation turn is persisted to Postgres. Sessions survive pod restarts without losing context.

The Security Model — What Is Actually True

This is not a list of security marketing claims. It is a description of specific design decisions made before the first line of production code was written.

DESIGN DECISION	WHAT IT MEANS
The AI never generates SQL	Every data query Beacon executes was written by a Spring Point engineer, reviewed, and deployed as a fixed parameterized query. The model decides which query to run and what values to pass—nothing else. This eliminates the class of risk where an AI generates a malicious or out-of-scope database query.
Authorization enforced at the tool layer	Every tool call goes through the same menu-level authorization used across WebMB. If a user account does not have access to a module, the tool call is rejected—Beacon cannot route around it.
Division-level data scoping respected	If your account is scoped to specific divisions in WebMB, that scoping applies to every Beacon data query. There is no elevated access path.
System behavior locked at deployment	The system prompt that governs Beacon's behavior is a version-controlled file deployed with the API. It cannot be modified at runtime and is not stored in a database where it could be altered without a code deployment.
AWS IAM authentication only	There are no API keys in the system that could be leaked or rotated incorrectly. Access to the underlying model is governed by the same IAM role that governs the rest of the service's AWS access.

What Beacon Will Be Able to Do at Launch

The following capabilities are built and in active testing. This is the confirmed launch set—not a wishlist.

All capabilities targeting rollout Sept 2026–Jan 2027.

TOOL	WHAT IT DOES
Open Job Summary	Summarizes open jobs by status and activity
Stalled Job Detection	Identifies jobs with no recent activity
At-Risk Job Flagging	Flags jobs at risk of missing committed deadlines
Job Activity Breakdown	Activity pattern breakdown across the shop floor
New Job Intake Summary	New jobs received by week across the shop
Financial & AR Summary	Revenue, billing state, and AR aging on demand
WebMB How-To Assistant	Answers procedural questions from the curated knowledge base

Additional capabilities will be announced with explicit target quarters as development progresses.

The Knowledge Base

The knowledge base is a curated set of Markdown articles covering WebMB features, screen-by-screen workflows, field definitions, and resolved support cases. Articles are converted to vector embeddings and stored in pgvector on Aurora Postgres.

The KB is updated continuously. Adding or updating articles does not require a code deployment or an API restart—the new content is ingested and searchable on the next embedding run.

Beacon's answers are grounded in retrieved KB content and pre-approved tool output. It does not answer from general AI knowledge. When the KB has no relevant result, Beacon acknowledges the gap and routes to support. An unconfident but honest answer is better than a confident wrong one.

Beacon is the beginning of a shift in how you interact with WebMB—from navigation to conversation. It is built to be useful at launch and to get sharper as your data grows and your questions shape it.

BEACON · SPRING POINT SOLUTIONS TECHNOLOGY TEAM · JULY 2026

To request early access: support@springpt.com | [888-382-7835](tel:888-382-7835)



Your ERP, Modernized.

What Is Changing, What Is Not, and What You Can Count On in 2027

For over 25 years, Spring Point Solutions has built software that industrial repair shops actually use every day, on the shop floor, in the office, and in front of customers. This issue lays out exactly where we are headed, why it matters, and what you can expect from us throughout the transition.

IN THIS ISSUE

- 01 The Partnership:** Our shared vision, philosophy, and what this modernization means for everyone.
- 02 Owners & Operators:** Plain English. No tech jargon. What this means for your shop, your team, your day.
- 03 Technical Operations:** Architecture, AI-centric engineering, the Fabric platform, and security for those who want the full picture.

01 THE PARTNERSHIP

Stability Through Innovation

There is a question that has been sitting in the room at every customer conversation we have had over the past year: "If your platform has roots in MS Access, what does the future actually look like?"

This issue is our direct, unambiguous answer.

We are modernizing aggressively, urgently, and deliberately. But we are doing it in a way that protects what you have built on our platform. That is not marketing language. It is a technical and operational commitment, and we will show you exactly what it means in the pages that follow.

What We Believe

Spring Point was built on a single conviction: industrial repair businesses deserve software that was actually designed for them. Not adapted from generic ERP platforms. Not bolted together with spreadsheets and workarounds. Built for motor shops, pump repair facilities, and apparatus service centers from the ground up.

That conviction has not changed. What is changing is the technology underneath it.

Our Promise Has Always Been the Same

More efficient operations. More productive teams. More profitable businesses. We have heard versions of this from Jacksonville to Portland to South Bend: when Spring Point works, it changes how a shop runs. We are modernizing to make sure it keeps working better for the next 30 years.

97%+

CUSTOMER RETENTION

Speaks for itself

25+

YEARS SERVING INDUSTRY

Domain knowledge built in

Jan 2027

FULL LAUNCH TARGET

Milestone-gated rollout

The 2027 Initiative

At our WebMB Tech Workshop in March 2026, we formalized what had been building for several years: a deliberate shift to an AI-centric engineering model. This is not a rebrand. It is a change in how we build software faster, more consistently, with better test coverage and higher release confidence.

The 2027 Initiative is the product of that shift. It is a full-stack modernization of the Spring Point platform, designed to deliver three things:

Faster delivery	Of improvements to your hands—not years from request to release.
Better reliability	Observability baked in so we catch issues before you do.
A cleaner foundation	For everything that comes next—integrations, reporting, mobile, AI-assisted workflows.

We are not doing this in one "big bang" release. We know what those look like. We have seen what happens when ERP vendors force a migration. We will not do that to you.

The Approach: Incremental modernization with validation gates at every stage. Small, helpful updates starting now. Full ecosystems launch January 2027. Your data, your workflows, and your team remain intact throughout.

What This Means for Our Partnership

Spring Point has always operated as a partner-led organization. That does not change during modernization. Our support team remains your point of contact. Support does not pause. Questions get answered.

We are asking for your trust and we are earning it by being transparent. Every issue of this blog will share specific milestones, what they mean for your operations, and how we are staging releases safely. No vague roadmaps. No marketing promises without proof points.

Read What Matters to You

- Owners & Operators** → If you run the business or the shop floor, go to Section 02. Plain English, no tech jargon—just what this means for your day.
- Technical Operations** → If you want the architecture, the AI model, the stack decisions, and the engineering detail, go to Section 03.

02 OWNERS & OPERATORS

What This Means for Your Shop

Plain English. No Tech Jargon. Just What You Need to Know.

You have been running your shop on Spring Point for years. Maybe a decade or more. You know the system, your team knows the system, and it works. So when you hear words like "modernization" and "platform migration," it is reasonable to feel a little uneasy. We get it. And we want to put that unease to rest right now.

The Short Version: *We are installing a new, high-performance engine under the hood of your ERP. From your seat, it should feel smoother, faster, and easier—not different in a way that disrupts how you do business.*

What Is Actually Changing

Think of it this way: your shop has been running a reliable truck for 25 years. It has gotten you everywhere you needed to go. Now we are upgrading that truck with a modern engine, better instruments, and a cleaner cab—without making you relearn how to drive it.

Here is what you will notice:

Faster	Everything loads faster. Slow screens and loading times are one of the most common complaints we hear. Reports that used to take 30 seconds to generate will run in a fraction of that time. Screen navigation will feel snappier, especially on older hardware.
Smarter Reporting	Your data has always been there. The new architecture makes it easier to get at it—cleaner dashboards, faster queries, and reporting that does not require you to call us to pull a custom report.
Cleaner Interface, Easier Training	Every shop owner we talk to mentions the same pain point: training new employees on the system takes too long. The updated interface is designed to be more intuitive out of the box, which means less time onboarding new hires and fewer calls to your most experienced people to explain where things are.

What Is NOT Changing

This is the part we want to be absolutely clear about.

Your Data	Every customer record, job history, invoice, and equipment detail stay exactly where it is. We are not moving your data to a new system—we are upgrading the system around your data.
Your Workflows	We are not here to tell you how to fix motors. If your team has a process that works, we are not going to break it. The changes we are making are underneath the surface.
Your Support Team	Our team is still your people. Same phone number, same faces, same commitment to picking up when you call. Modernization does not mean automation replacing humans on our end.
Your Pricing	Whether you prefer the traditional MotorBase interface* or the new WebMB platform, the features you rely on remain the same. Access to both interfaces is included under your current subscription rate, with no initial upgrade fees required.

**Note: A 3–5-year sunseting period for the traditional MotorBase interface will begin after the new WebMB platform is launched, fully stabilized, and actively replacing legacy environments.*

The Timeline: What to Expect and When

Now → Late 2026	Small, incremental improvements roll out automatically through your existing update process. You will notice things getting faster and smoother. Nothing requires action on your part.
Late 2026	We will reach out directly to walk through the updated interface with your team. Hands-on walkthrough, your questions answered, your team prepared.
January 2027	Full launch of the modernized ecosystem. By this point, your team will have already experienced most of the changes in smaller doses. No surprises.

What You Should Do Right Now

Honestly? Nothing urgent. You do not need to prepare for anything, upgrade any hardware, or brief your team today.

What we do ask:

- Keep an eye out for update notifications—these will come through your normal channel and are worth installing when they arrive.
- Let us know if anything feels off. Our team wants to hear from you if something in an update does not feel right for your operation.
- Share this with your office manager or lead tech—they will appreciate knowing what is coming.

Have Questions? Email us at support@springpt.com—our team responds quickly.

We have been in this business long enough to know that trust is earned, not declared. We are not asking you to take our word for it. We are asking you to watch what we do and hold us to it.

03 TECHNICAL OPERATIONS

Inside the Engine Room

Full-Stack Modernization & AI-Centric Engineering at Spring Point

In January 2026, Spring Point formally shifted its Technology Organization to an AI-centric operating model—a change in how we design, build, test, and ship software. This section covers the architectural decisions, tooling, release strategy, and where AI fits into the engineering lifecycle (and where it does not).

The Case for Full-Stack Modernization

Spring Point's platform carries over two decades of domain value. That value is real and it is not going anywhere. But the delivery model that built it—sequential development cycles, a monolithic codebase with significant MS Access underpinnings, limited automation in testing and deployment—was a ceiling on how fast we could move and how reliably we could ship.

The modernization effort addresses that ceiling directly. The target state is a platform that can:

- Ship incremental improvements in days, not months
- Scale database and API performance independently of UI surface area
- Surface operational telemetry before issues reach end users
- Support new integration patterns (REST, webhooks, third-party connectors) without architectural surgery

Architecture: The Stack

Frontend: React Component Architecture

The legacy UI is being replaced with a React component model—modular feature deployment, state persistence across complex ERP modules, and 60fps target responsiveness in data-heavy views. The component model also enables A/B testing of UI changes with specific customer segments before broad release, something the legacy architecture made effectively impossible.

API Layer: RESTful Service Decoupling

Business logic is migrating from embedded UI procedures to a standardized RESTful API layer. This enables mobile client support, third-party integrations, and versionable endpoints—meaning we can evolve the backend independently from the frontend on each release cycle.

Data Tier: Relational SQL + Optimized Indexing

The data tier moves to a modern relational SQL engine with query plan optimization profiled against live production workloads. Early benchmarks show significant latency reductions on the large-dataset operations that drive the reporting screens where degradation is most visible to end users.

LAYER	TECHNOLOGY	KEY BENEFIT
Frontend	React (Component-Based)	Modular deploys, 60fps responsiveness, testable UI
API	RESTful Service Layer	Decoupled business logic, versionable endpoints
Data	Relational SQL + Indexing	Reduced query latency, scalable for large datasets

The AI-Centric Engineering Model

When we say AI-centric, we mean that Gen AI tooling is embedded across the software delivery lifecycle as a productivity multiplier for our engineering team—not as a novelty, but as a systematic accelerator. Here is what that looks like in practice:

Automated Refactoring	Gen AI accelerates conversion of legacy business logic (some 15+ years old) into modern codebases. AI handles the initial conversion, engineers review against expected behavior, test coverage gates the merge. We move faster through the backlog without sacrificing correctness.
Synthetic Test Generation	AI-driven unit and integration tests extend coverage into edge cases and legacy code paths—targeting a 99.9% uptime baseline before any significant release reaches production.
Implementation Acceleration	Gen AI as a co-pilot for new feature work: boilerplate generation, API contract proposals, schema suggestions. All outputs are draft artifacts—every line that ships has been reviewed against engineering standards.

Important: We are not using your production environment as a test bed for AI-generated code. AI accelerates our internal delivery process. Engineering review, testing pipelines, and validation gates still govern what reaches your system.

The Fabric Platform

Fabric is our internal developer platform (IDP)—the infrastructure abstraction layer that lets engineers focus on product capabilities rather than deployment pipelines and environment management. It handles standardized CI/CD with validation gates, environment provisioning for feature branches, artifact management, rollback capability, and release health metrics. Fabric is not customer-facing, but it is what makes our release cadence commitments real and reversible.

Security & Observability

Security: Data-at-rest encryption as a baseline (not an add-on). TLS 1.3 in transit. Role-based access controls tightened and audited as part of the API migration.

Observability (OpenTelemetry): OpenTelemetry-based distributed tracing and structured logging across the service layer. Queryable telemetry at the span level means we identify latency regressions and error spikes before they surface as customer-reported incidents. Our ops team sees the problem first—not yours.

Release Strategy: No Big Bang

No cutover event. No parallel system. No data migration requested of customers. This is where modernization efforts most commonly fail—we are not doing any of those things. The approach:

Incremental service migration	Individual capabilities are extracted and modernized independently, with feature flags controlling customer exposure.
Controlled rollout	New capabilities launch to a subset of customers first, with explicit validation criteria before expanding.
Validation gates	Automated test coverage thresholds, performance benchmarks, and error rate ceilings gate each release stage.
Explicit rollback paths	Every deployment is designed to be reversible within a defined SLO window.

The January 2027 target is a milestone, not a cliff. By that date, the majority of the new architecture will already be serving production traffic—the launch is the completion of an incremental process, not the beginning of a migration.

99.9%

UPTIME TARGET

Pre-release test baseline

Incremental

RELEASE MODEL

No cutover events

OpenTelemetry

OBSERVABILITY

Full distributed tracing

Real-World Problems This Solves

Architecture decisions matter most when they remove friction people have quietly tolerated for years. Here are five specific issues the modernized platform directly resolves.

Printing: No More Driver Headaches

The legacy platform has no universal print driver support and no print redirection—every workstation needs the right drivers installed locally, and remote users are largely on their own. The modernized stack moves to browser-based rendering with server-side print services, eliminating driver dependency entirely. Any device, any printer, no IT intervention required.

Email: Full OAuth2 — No More SMTP Administration in MotorBase

Outbound email currently requires SMTP credentials configured and maintained directly inside MotorBase—a maintenance burden and a credential security exposure. The new platform integrates OAuth2 natively, meaning email flows through your existing Microsoft 365 or Google Workspace identity. No app passwords. No SMTP configs to manage. No calls to us when an email password expires.

VPN & Remote Access: Moving Beyond PPTP

The PPTP VPN used for remote access is a deprecated protocol with well-documented vulnerabilities—it has no place in a modern security posture. With the platform modernization, PPTP goes away. Remote and secure data access moves to a cloud-native model built around a secure S3-based interface, removing the need to maintain VPN infrastructure and significantly tightening the access security boundary.

Note: The specifics of the new remote access architecture are worth a dedicated conversation—reach out to support@springpt.com if this affects your current setup and you want to understand the transition plan before it rolls out.

Context-Sensitive Knowledge Base

The modernized platform includes an integrated, context-aware knowledge base: when a technician opens a job screen, a part entry, or a QM workflow, relevant documentation surfaces inline in the right place, at the right time. Less searching. Fewer errors. Faster onboarding for new staff.

Modern Technology Integrations: Starting with Mapping

The new RESTful API layer opens the door to integrations that were previously architectural surgery. Mapping is the first concrete example: field service routing, customer location visualization, and address validation through modern mapping APIs. Additional integration targets include real-time parts sourcing, customer portal enhancements, and third-party reporting platforms—the foundation supports all of them without custom development effort each time.

What Comes Next

Coming issues will cover: API documentation and integration patterns, a Fabric pipeline deep-dive, before/after SQL benchmarks from real production workloads, and mobile client architecture. Have a specific area you want us to cover?

Get In Touch: Reach us at support@springpt.com—we would rather write what engineers actually want to read.

SPRING POINT SOLUTIONS

Specializing in the unique requirements of the industrial community.

"Bringing order from chaos." — since 1999

support@springpt.com · Portland, ME 04101



INTRODUCING: THE SPRING POINT SOLUTIONS TECH BLOG


WHERE INDUSTRIAL EXPERTISE MEETS MODERN SOFTWARE ENGINEERING.



OWNERS & OPERATORS:
PLAIN ENGLISH. NO TECH JARGON.

Engineering the Future of Industrial Software.

Get the operational blueprint.



TECHNICAL OPERATIONS:
CODE LOGIC INSIGHT

Full-Stack Modernization and AI-Centric Engineering.

Solving real problems without disruption.

JOIN THE CONVERSATION.

We would rather write what
engineers actually want to read.

SCAN NOW TO READ ISSUE 01



YOUR FUTURE, ACCELERATED.

support@springpt.com | www.springpt.com
"Bringing order from chaos." – since 1999