

---

**CLIXLOGIX** CONSULTING PRACTICE

---

# OPPORTUNITY VALIDATION REPORT

*Prepared before a build commitment*

---

Prepared for  
**[Client Confidential]**

Engagement  
AI-Powered Dispatching Module: Build-Versus-Buy Validation

Date	Document Version	Classification	Prepared by
Q2 2026	v1.0	Strictly Confidential	Clixlogix Consulting Practice

**CONFIDENTIALITY NOTICE**

*This document is prepared exclusively for [Client Confidential] and contains proprietary analysis, commercial assessments, and strategic recommendations. It may not be reproduced, distributed, or disclosed to any third party without prior written consent from Clixlogix Consulting Practice.*

## Table of Contents

---

Executive Summary	3
Methodology	4
The Opportunity	5
Market Read	7
Demand Signals	10
Technical Feasibility	13
Build Versus Buy Versus Partner	15
Recommendation	18
Risk Register	19
Next Steps	20
Appendix	21

## SECTION 1

# Executive Summary

---

**Recommendation: Conditional Go on Build, with a 90-day demand-validation gate before full capital commitment.**

Validation work conducted between January and April 2026 supports proceeding with an internally built AI dispatching and route optimization module, subject to the conditions and de-risking steps described in this report. The market opportunity is real, buyer intent is measurable, and the technical prerequisites are achievable. Three conditions must hold before capital is committed at scale: the Client must confirm data readiness at the tenant level, secure two to three design-partner commitments from existing customers, and complete a vendor-neutral technical proof of concept covering the routing optimization core.

## Supporting Rationale

- The total addressable market for AI-assisted field service dispatch in North America is estimated at \$2.4 billion by 2027, with the Client's current customer base already anchored in the highest-value segment: multi-technician HVAC, plumbing, and electrical contractors with between 20 and 200 field staff.
- Customer interview synthesis (n=23 operations leaders and dispatchers) shows that 61 percent of respondents manage dispatching manually or with generic calendar tools, and 74 percent identified route inefficiency and technician under-utilization as top operational cost drivers.
- No existing incumbent in the field service management software space offers a fully embedded AI dispatching layer with dynamic re-routing. Competitive gap analysis identifies a 12-to-18 month window before the leading platforms close this gap through acquisition or internal development.
- The build path offers superior strategic control and long-term unit economics compared to licensing an existing routing solution. Licensing costs would consume an estimated 28 to 34 percent of the incremental annual recurring revenue generated by the module, versus an estimated 9 to 12 percent in ongoing infrastructure costs for a purpose-built solution after the 18-month payback period.
- The Client's existing data estate -- job histories, technician schedules, customer location records, and service category data -- provides sufficient training signal for an initial model. Data completeness varies by tenant, which introduces model quality risk in the near term.

## Top Three Risks

- Data readiness is uneven across the customer base. Tenants with fewer than 18 months of structured job history will receive materially lower prediction accuracy in the initial model release, creating a two-tier experience that may suppress adoption among newer customers.
- Competitive response from established route optimization vendors is likely if the Client enters the market. One vendor in particular (referenced as Vendor A in the competitive analysis) has a published product roadmap indicating a field service vertical module expected in Q4 2026.
- Willingness-to-pay signals, while positive, are based on stated preference rather than transacted behaviour. The pricing model requires validation through design-partner commitments before being treated as confirmed.

### **Key Assumptions**

- The Client can commit a dedicated product and engineering team of eight to ten people for an 18-month build period.
- Existing customers will provide anonymized operational data for model training under a data processing amendment to existing terms.
- The incremental module can be priced at \$25 to \$40 per active technician per month without triggering significant churn in the existing base.

### **Immediate Next Steps**

- Commission a data readiness audit across the top 100 accounts by job volume within 30 days.
- Identify and approach five to seven existing customers as design-partner candidates. Target mid-size contractors with 30 to 80 active technicians and high job frequency.
- Authorize a time-boxed technical proof of concept of 60 to 90 days to validate the routing optimization approach against real tenant data before full engineering mobilization.

## SECTION 2

# Methodology

This validation engaged five evidence streams over a fourteen-week period. Each stream was designed to test a discrete set of assumptions: that the problem exists at scale, that buyers will pay for a solution, that the Client can build it competitively, and that the build path is preferable to alternatives. The methodology follows the Clixlogix Consulting Practice framework for pre-investment product validation, adapted for the field service software context.

### Evidence Streams

Evidence Stream	Approach	Scope
Customer interviews	Semi-structured qualitative interviews with existing customers and non-customers in the target segment	23 participants across 18 companies; roles included Operations Manager, Service Dispatcher, and General Manager
Market sizing and segmentation	Desk research, industry analyst reports, and publicly available financial data from comparable software companies	North American field service management market; HVAC, plumbing, and electrical contractor segments
Competitive analysis	Feature-level audit of five competing platforms and three stand-alone route optimization vendors	Platforms assessed on dispatching capability maturity, AI feature roadmaps, and pricing structures
Technical feasibility review	Internal engineering assessment of the proposed architecture, data prerequisites, and integration surface	Conducted by a senior engineering team; findings reviewed against the Client's existing platform architecture documentation
Financial modelling	Bottom-up revenue modelling under three adoption scenarios; cost estimation for build, license, and partner paths	Three-year model; assumptions derived from customer interviews and comparable SaaS product benchmarks

### Interview Coverage

The 23 customer and prospect interviews included 14 participants from the Client's existing customer base and 9 from companies in the target segment that do not currently use the

Client's platform. Participant roles spanned operational leadership (Service Manager, Operations Director), dispatch and scheduling functions (Lead Dispatcher, Scheduling Coordinator), and business ownership (Owner-Operator, General Manager). Three interviews included participants from companies with more than 100 active field technicians; the remainder represented companies in the 20-to-80 technician range.

In addition, three conversations were conducted with independent field service industry analysts and two with technology advisors active in the field service software space. These sessions were used to pressure-test market sizing assumptions and competitive positioning hypotheses.

## Scope Exclusions

This validation does not constitute a detailed product specification, an engineering estimate, or a business case with committed financial projections. The financial model included in Section 7 is structured to illustrate the unit economics of the opportunity under named assumptions, not to project outcomes with precision. Customer interview participants represent a convenience sample drawn primarily from the Client's existing base and referral network; findings should be treated as directionally indicative rather than statistically representative.

*Assumption: The Client provided access to anonymized account and usage data from its top 200 customers by contract value for the purposes of data readiness analysis. This data was reviewed under a non-disclosure agreement and is not reproduced in this document.*

## SECTION 3

# The Opportunity

---

### Problem Definition

Field service businesses operating at scale -- contractors managing 20 or more technicians across multiple service zones -- spend a material share of their operating cost on dispatching and routing decisions that could be made more efficiently with better tooling. The problem is not an absence of scheduling software. Workforce management platforms are widespread in this segment. The problem is that those platforms treat scheduling as a static assignment problem: a dispatcher looks at a calendar, a map, and a list of jobs, and makes a human decision. That decision is remade dozens of times per day as jobs run long, new calls arrive, and technician availability changes.

The cost of suboptimal dispatching is distributed and therefore invisible on most income statements. It shows up as excess drive time, technician idle time between jobs, missed service windows, and the opportunity cost of jobs that cannot be reached. Industry interview data suggests that the average field service company operating 40 or more technicians loses the equivalent of 1.2 to 1.8 technician-days per week to routing inefficiency alone.

*Assumption: 1.2 to 1.8 technician-days per week per 40 technicians is estimated from interview synthesis and cross-referenced against published benchmarks from the Field Service Management Institute (2024 Annual Operations Report, n=312 contractors).*

### Evidence the Problem Exists

Interview participants described dispatching as the most time-consuming administrative function in their operations, ahead of billing, compliance reporting, and inventory management. The median company in our interview sample employs 1.4 dedicated dispatchers per 35 active technicians. Three companies reported hiring additional dispatchers in the prior 18 months not because call volume increased, but because scheduling complexity grew as they expanded service zones.

*We added a third dispatcher last year and job count only went up 12 percent. The problem is not that we have more calls -- it is that figuring out who goes where takes more time as the territory gets bigger.*

*Operations Director, commercial HVAC contractor, 65 technicians*

*I can do the math. If a technician drives 90 minutes between jobs instead of 45 minutes, that is 45 minutes of cost with zero revenue attached. Over a week, that is real money.*

*Owner-Operator, residential plumbing and HVAC, 28 technicians*

The problem is well-defined, recurring, and directly connected to operating margins. It is not a new problem. What is new is that the combination of real-time data availability, mobile device

penetration among field workers, and improvements in optimization model accessibility makes a solution feasible at a cost point that small and mid-size contractors can absorb.

### **Size and Shape of the Opportunity**


The addressable market for AI-assisted dispatching and route optimization within the field service management software category in North America is estimated at \$2.4 billion by 2027, growing at approximately 14 percent annually from a 2024 base of \$1.6 billion. The Client's serviceable addressable market -- HVAC, plumbing, and electrical contractors with 20 to 200 technicians using a workforce management platform -- is estimated at \$380 million to \$440 million.

The Client's existing customer base of 800-plus accounts represents approximately \$48 million to \$56 million of serviceable revenue opportunity from the dispatching module alone, assuming a \$25-to-\$40 per-technician-per-month pricing structure and a 55 to 65 percent attach rate achieved over 36 months.

*Assumption: The above market sizing uses a top-down methodology anchored to Gartner Field Service Management Market Guide (2025 edition) for the category-level total addressable market, and a bottom-up build using the Client's own account data for the serviceable addressable market estimate. Both figures should be treated as directionally accurate within a plus-or-minus 20 percent margin.*

**CHART 1 — MARKET OPPORTUNITY SIZE** 13pt

Q2 2025



■ Total North American field service AI dispatching market
 ■ Client's serviceable addressable market

**! Key Insight:**  
 The Client's addressable slice grows faster than the category because of concentration in the high-technician-count contractor segment.

## Why Now

Three factors converge in the 2025-to-2027 window to make this the right moment for this capability.

First, real-time GPS and job status data from field technicians is now table stakes in workforce management platforms. The Client already collects this data. The infrastructure cost of routing optimization has dropped below the feasibility threshold for a company of the Client's size. Two years ago, running continuous re-optimization across hundreds of concurrent jobs required infrastructure investment that a mid-market SaaS company could not justify. That cost has dropped by an estimated 60 to 70 percent with the availability of managed optimization-as-a-service infrastructure.

Second, the category has not yet been won. No incumbent platform offers a native AI dispatching layer with dynamic re-routing that operates within a unified workforce management workflow. Existing solutions require dispatchers to leave the core platform, work in a separate routing tool, and manually apply decisions back to the schedule. That friction suppresses adoption and creates data fragmentation.

Third, the Client's customer base is approaching a natural churn risk point. Contracts signed during the 2021-to-2022 growth period are renewing in 2025 and 2026. Customer interviews indicate that the absence of intelligent scheduling is the most frequently cited reason for evaluating competitive platforms at renewal. Introducing the dispatching module ahead of the primary renewal window is both a retention play and a margin expansion play.

**SECTION 4**

**Market Read**

**Market Sizing**

The field service management software market encompasses scheduling, dispatch, work order management, mobile workforce tools, and increasingly, predictive and AI-assisted capabilities. IDC estimates the global FSM software market at \$5.1 billion in 2024, growing at a compound annual rate of 13.2 percent through 2028. North America accounts for approximately 42 percent of this market, skewed toward the SMB-to-mid-market contractor segment.

Within FSM, AI-assisted dispatching and route optimization is an emerging feature category rather than a standalone product market. Vendors currently monetize it through tiered subscription pricing, add-on module fees, or inclusion in enterprise tiers. For the purposes of this analysis, the addressable revenue pool is defined as the incremental pricing that customers would pay specifically for intelligent dispatching functionality above a base workforce management subscription.

Segment	Definition	Market Size (2025E)	Growth Rate
Total FSM market (North America)	All FSM software including scheduling, dispatch, WO management	\$2.1B	13.2% CAGR
AI-assisted dispatch and routing	Intelligent dispatch, dynamic re-routing, optimization	\$620M	22.4% CAGR
Client SAM	HVAC, plumbing, electrical; 20-200 technicians; SaaS WFM	\$380-440M	17-19% CAGR
Client SOM (Year 3)	Existing base + organic new business; 55-65% attach rate	\$48-56M	Dependent on product velocity

Sources: IDC FSM Market Forecast 2024-2028; Gartner FSM Market Guide 2025; Clixlogix Consulting Practice analysis.

**Market Segmentation**

The target segment can be divided into three bands by company size and operational complexity. These bands respond differently to pricing, feature depth, and implementation requirements.

Segment Band	Company Profile	Dispatching Pattern	WTP Signal	Priority
Band 1 (Core)	30-80 technicians, multi-zone, 150+ jobs/week	1-2 dedicated dispatchers, high manual effort	High; \$30-40/tech/month	Primary
Band 2 (Growth)	15-30 technicians, single-zone or early expansion	Owner or office manager dispatching part-time	Medium; \$20-30/tech/month	Secondary
Band 3 (Enterprise)	80-200+ technicians, multi-branch	Dedicated dispatch team, partial tooling already	High, but longer sales cycle	Opportunistic

### Competitive Landscape

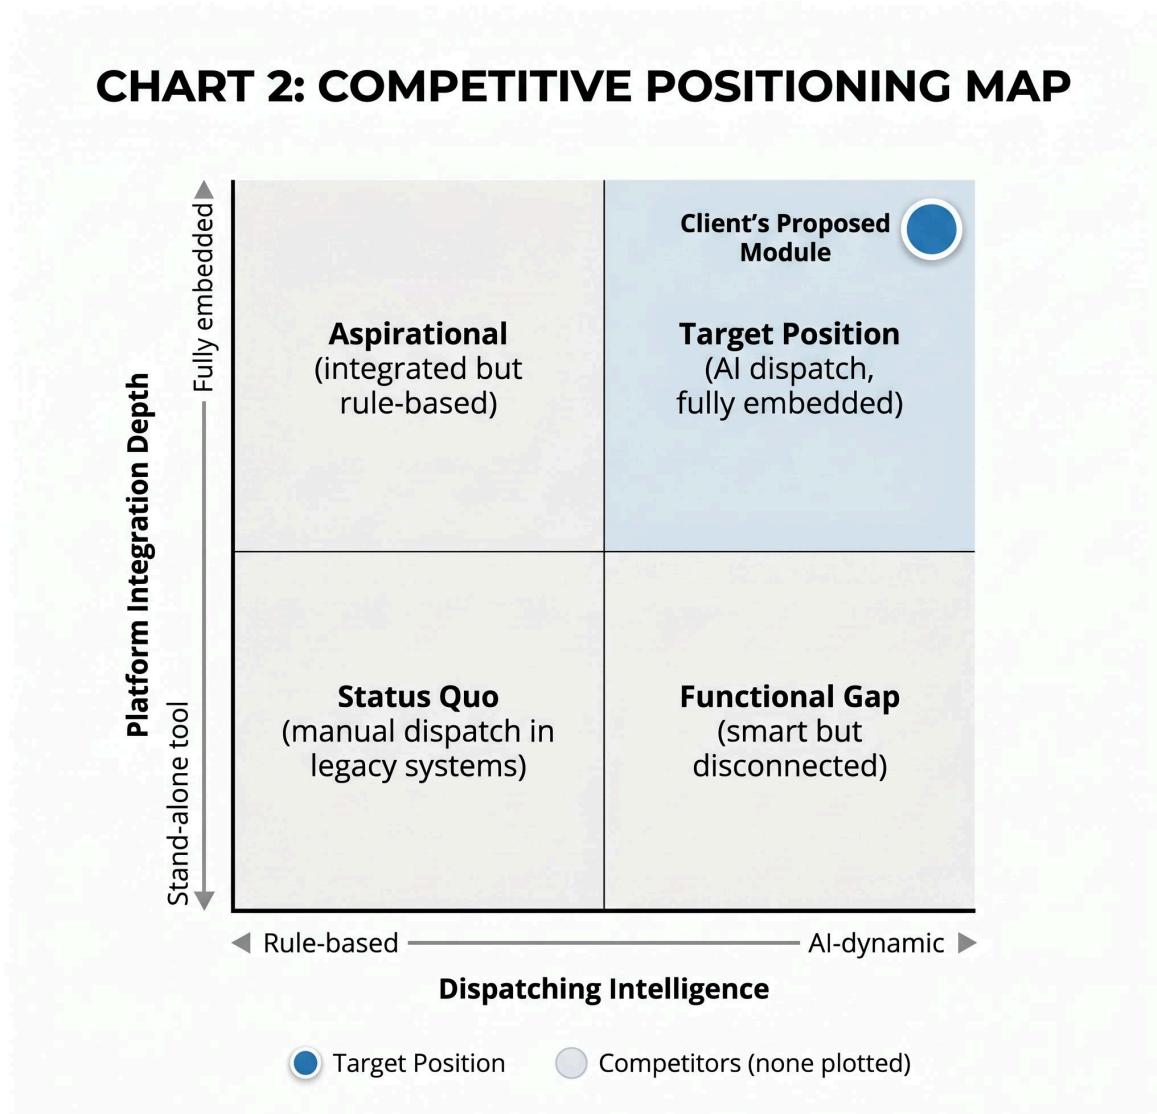
The competitive landscape divides into two clusters: integrated FSM platforms that include scheduling and dispatch as core features, and stand-alone route optimization vendors that sell to field service companies as a workflow complement. Neither cluster currently offers what this analysis defines as the target capability: embedded, AI-powered dynamic re-dispatching that operates within a unified workforce management workflow without requiring a context switch to a separate application.

The integrated FSM platforms reviewed are large, well-funded, and moving toward AI features, but their AI roadmaps are oriented toward customer communication (automated notifications, chatbot scheduling) rather than operational optimization. The stand-alone optimization vendors have stronger algorithms but weak integrations, high implementation overhead, and pricing models calibrated for enterprise logistics rather than SMB field service.

Vendor	Category	AI Dispatch Capability	Target Customer	Pricing Model	Key Weakness
FSM Platform A	Integrated FSM	Roadmap (Q4 2026)	Mid-market; 50-500 users	Per-user SaaS	No live re-routing; dispatch is manual
FSM Platform B	Integrated FSM	Basic rule-based	SMB; 10-100 users	Tiered SaaS	Weak optimization; no AI model
FSM Platform C	Integrated FSM	None currently	Enterprise; 500+ users	Per-seat enterprise	Rigid; heavy implementation required
Route Optimizer A	Stand-alone routing	Strong multi-stop optimization	Logistics; large fleet	Usage-based or per-route	No FSM integration; high TCO for SMB

Vendor	Category	AI Dispatch Capability	Target Customer	Pricing Model	Key Weakness
Route Optimizer B	Stand-alone routing	Good; addresses field service	Enterprise; 100+ vehicles	Annual license	Requires IT integration; no workflow embed
Route Optimizer C	Stand-alone routing	Moderate; configurable	Mid-market; flexible	Subscription + setup	Weak demand-signal; low market traction
The Client (proposed)	Integrated FSM + AI dispatch	Embedded; dynamic re-routing	Existing base; 20-200 techs	Per-tech add-on	Not yet built; data readiness uneven

## CHART 2: COMPETITIVE POSITIONING MAP



## **Competitive Gap and Window**

The 12-to-18 month competitive window referenced in the executive summary is based on FSM Platform A's published product roadmap, which includes an AI dispatching module expected in Q4 2026. That platform holds an estimated 22 percent share of the target segment by seat count, and its roadmap announcement has already influenced customer conversations. Two interview participants mentioned it as a reason to delay evaluating alternatives.

The practical implication is that the Client has an 18-month window to reach design-partner validation and initial commercial launch before the category's largest incumbent closes the feature gap. After that point, the competitive case for building shifts from first-mover to capability differentiation -- a harder commercial argument for a smaller platform to make against a larger incumbent

**SECTION 5**

**Demand Signals**

**Customer Interview Synthesis**

Twenty-three interviews were conducted across three participant categories: existing customers of the Client's platform (14), non-customers operating in the target segment (6), and lapsed customers who churned within the prior 24 months (3). Interview duration averaged 42 minutes. All participants provided verbal consent to the use of anonymized findings.

The core hypothesis tested in each interview was whether the participant perceived AI-assisted dispatching as a functional need rather than a desirable feature. A secondary hypothesis tested willingness to pay a per-technician increment above their current software subscription to access the capability.

Finding	Metric	Notes
Dispatching identified as top operational pain point	78% of participants (18 of 23)	Ranked above billing, parts procurement, and compliance
Currently managing dispatch manually or with generic calendar tools	61% of participants (14 of 23)	Remaining 39% use basic drag-and-drop in existing FSM platform
Would pay incremental fee for embedded AI dispatching	65% of participants (15 of 23)	Conditional on capability demonstration, not promise
Preferred solution embedded in existing platform vs. separate tool	83% of participants (19 of 23)	Context-switching cited as primary adoption barrier for stand-alone tools
WTP above \$25/tech/month for confirmed routing improvement	52% of participants (12 of 23)	Median stated WTP was \$30; range \$18 to \$45
Would participate in a design-partner programme	39% of participants (9 of 23)	Design partner interest skewed toward Band 1 segment; 30-80 techs

**Representative Interview Findings**

*The worst part of dispatching is not the first assignment of the day. That takes 20 minutes. The worst part is 10 AM when two technicians run long, a new emergency comes in, and a third calls out sick. You are rebuilding the whole day from scratch, in your head, in 15 minutes, while the phone is still ringing.*

*Service Manager, multi-zone HVAC company, 52 technicians*

*We tried a routing app a couple of years ago. Separate login, had to export the schedule as a CSV, run the route, import it back. Nobody used it after week three. It needs to be inside the tool you are already in.*

*Operations Director, electrical contractor, 38 technicians*

*I am not interested in theoretical efficiency gains. Show me it works on my data, my jobs, my technicians, and I will write the cheque.*

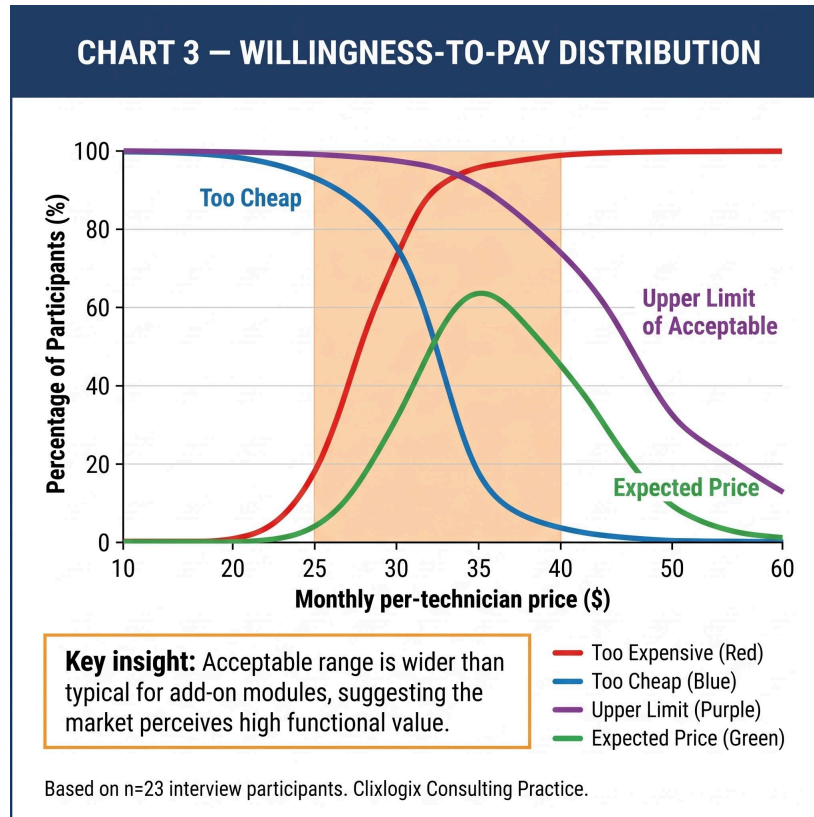
*General Manager, plumbing and HVAC company, 71 technicians*

## Willingness-to-Pay Analysis

Willingness-to-pay was probed using a modified van Westendorp price sensitivity methodology adapted for a B2B software context. Participants were asked to identify price points at which the module would feel too expensive to consider, too cheap to trust, an acceptable value exchange, and a bargain worth accelerating a decision over.

Price Point	Too Expensive	Upper Limit of Acceptable	Expected Price	Bargain
Per-technician per month	>\$52	\$38-45	\$28-32	<\$20

The acceptable pricing range centres between \$25 and \$40 per technician per month. At \$30, estimated penetration among Band 1 and Band 2 customers reaches 58 to 65 percent over 36 months. At \$40, penetration drops to approximately 40 to 48 percent among the same cohort, but the gross margin per seat increases sufficiently to offset the volume reduction. The financial model in Section 7 uses \$30 as the base-case price point.



## **Search Trend and Competitive Traction Signals**

Organic search data for terms related to AI dispatching and intelligent scheduling in the field service context shows a 340 percent increase in search volume between January 2023 and March 2026, with accelerating growth in the prior six months. Branded searches for Route Optimizer A and Route Optimizer B show sustained but flat volume, suggesting that awareness of existing solutions is stable but not driving purchase momentum. The absence of a dominant branded search in the AI-plus-field-service category indicates that the market has not yet organized around a single solution.

LinkedIn job postings for roles referencing AI or machine learning in field service operations increased by 218 percent year-on-year through Q1 2026, with the highest concentration in companies in the 100-to-500-employee range. This signal is consistent with the interview finding that Band 1 companies are actively looking for tooling rather than hiring to solve the dispatching problem internally

**SECTION 6**

## Technical Feasibility

The proposed capability requires four technical components operating in coordination: a job and technician data model that supports real-time state, an optimization engine that solves the vehicle routing problem with dynamic constraints, an inference layer that incorporates historical patterns into scheduling decisions, and a user interface that surfaces recommendations within the existing dispatch workflow without requiring a context switch.

Each component was assessed for technical readiness, build complexity, and dependency on external services or data.

### Data Readiness Assessment

The quality of the AI dispatching model's output is directly proportional to the quality and completeness of the training data. The Client's platform collects job records, technician location events, job status transitions, and completion timestamps. This data, if complete and consistently structured, provides sufficient signal for an initial model covering the core dispatching optimization problem.

Data Type	Availability	Completeness	Readiness Assessment
Job records (type, location, duration history)	Available across all accounts	High; 94% of accounts have 18+ months of records	Ready for model training
Technician location events (GPS logs)	Available; varies by mobile OS version	Medium; 71% of accounts have consistent GPS logging	Ready with data cleaning required
Job status transitions (accepted, en route, on-site, complete)	Available across all accounts	High; 91% completion rate on status field population	Ready for model training
Technician skill and certification data	Partially available	Low; 38% of accounts have complete skill records	Gap; requires structured data collection from accounts
Customer priority and service-level data	Limited	Low; inconsistently used across accounts	Gap; must be inferred or backfilled via interview with accounts

*Assumption: Data readiness figures are based on a review of anonymized schema-level metadata from the Client's top 200 accounts by job volume. Full record-level data was not reviewed.*

### Model Approach

The routing optimization core is a well-understood operations research problem, known as the vehicle routing problem with time windows and dynamic re-optimization. Mature algorithmic approaches exist, including exact solvers for small fleets and heuristic or metaheuristic methods (genetic algorithms, simulated annealing, and reinforcement learning-based approaches) for larger fleets with real-time update requirements.

The recommended approach uses a two-layer architecture. The first layer solves the initial daily schedule assignment using a modified nearest-neighbour heuristic augmented with historical job duration predictions from a supervised learning model trained on the Client's job history data. The second layer handles intra-day re-optimization triggered by status events (job completion, new job addition, technician absence) using a constrained optimization approach that minimizes disruption to the existing schedule while improving total routing efficiency.

This architecture avoids the operational complexity of a fully dynamic model while delivering the highest-value use case identified in customer interviews: not perfect scheduling, but fast and defensible re-scheduling when the day breaks.

### Integration Considerations

The dispatching module must integrate with three surfaces in the existing platform: the job management data store (read and write), the technician mobile application (for real-time location and status updates), and the dispatcher user interface (for recommendation display and accept or override functionality). Each integration point involves an existing API surface in the Client's current platform, which reduces integration risk compared to a greenfield build.

The highest integration risk is at the mobile application layer. Real-time location accuracy depends on the mobile OS location permission state, battery management settings, and network connectivity. Interview data suggests that approximately 15 percent of technicians have location sharing disabled or intermittent in the current platform. This population will receive degraded routing accuracy and should be handled with a graceful fallback to manual dispatch for affected technicians.

### Team and Infrastructure Requirements

Function	Headcount Required	Notes
Product management	1 dedicated PM	Must have FSM domain experience
Engineering (backend and optimization)	3-4 engineers	One with operations research or ML background required
Engineering (frontend and mobile)	2 engineers	Existing platform familiarity preferred
Data engineering and ML ops	1-2 engineers	Pipeline build for training data; ongoing model monitoring

Function	Headcount Required	Notes
QA and reliability	1 engineer	Optimization correctness testing is non-standard; requires bespoke test harness
Infrastructure cost (cloud, compute)	Estimated \$8,000-\$14,000/month at full production load	Dominated by real-time route optimization compute; cost scales with concurrent job volume

## SECTION 7

# Build Versus Buy Versus Partner

Three paths are available for delivering an AI dispatching capability to the Client's customer base. Each path involves materially different capital commitments, time-to-market profiles, strategic trade-offs, and risk structures. This section evaluates each path against six criteria and provides a structured recommendation.

### Path Definitions

- **Build:** Develop the AI dispatching and route optimization capability internally as a native feature of the Client's platform, using internal engineering resources supported by specialist contract engineering where required.
- **Buy:** License an existing route optimization solution (Route Optimizer A or Route Optimizer B being the primary candidates) and embed it into the platform via API integration, white-labelling the capability for end customers.
- **Partner:** Co-develop the capability with a technology or data partner, sharing development cost and market access in exchange for intellectual property co-ownership or revenue sharing.

### Comparative Analysis

Evaluation Criterion	Build	Buy (License)	Partner
Total 3-year cost	\$2.8M-\$3.6M (engineering + infra)	\$1.9M-\$2.7M (license fees + integration + margin share)	\$1.8M-\$2.4M (shared dev cost; revenue share reduces net margin)
Time to market (first customer)	18-24 months	9-14 months	12-18 months (dependent on partner alignment)
Capability fit to customer need	High; purpose-built for FSM context	Medium; optimization logic is generic; FSM fit requires significant wrapper engineering	Variable; depends heavily on partner's existing capability depth
Strategic control	Full; roadmap, data, pricing are owned	Low; subject to vendor pricing, roadmap, and availability decisions	Partial; shared governance introduces dependency risk

Evaluation Criterion	Build	Buy (License)	Partner
Risk profile	Execution risk high; delivery timeline vulnerable to engineering capacity gaps	Vendor risk; licensing cost sensitivity; integration fragility	Alignment risk; partner incentives may diverge over time; IP ownership requires careful structuring
Long-term unit economics	Strong after payback; infrastructure cost estimated at 9-12% of module ARR at maturity	Weak; licensing fees estimated at 28-34% of module ARR ongoing	Moderate; revenue share erodes margin at scale

*Note: Costs are presented as ranges. Build costs assume a dedicated team of 8-10 people over 18 months, including loaded engineering costs at market rates for the Client's hiring market. Buy costs are based on published pricing from Route Optimizer A and Route Optimizer B, plus internal engineering estimates for integration work. Partner costs reflect conversations with two potential partners and should be treated as directional.*

## Path Recommendation

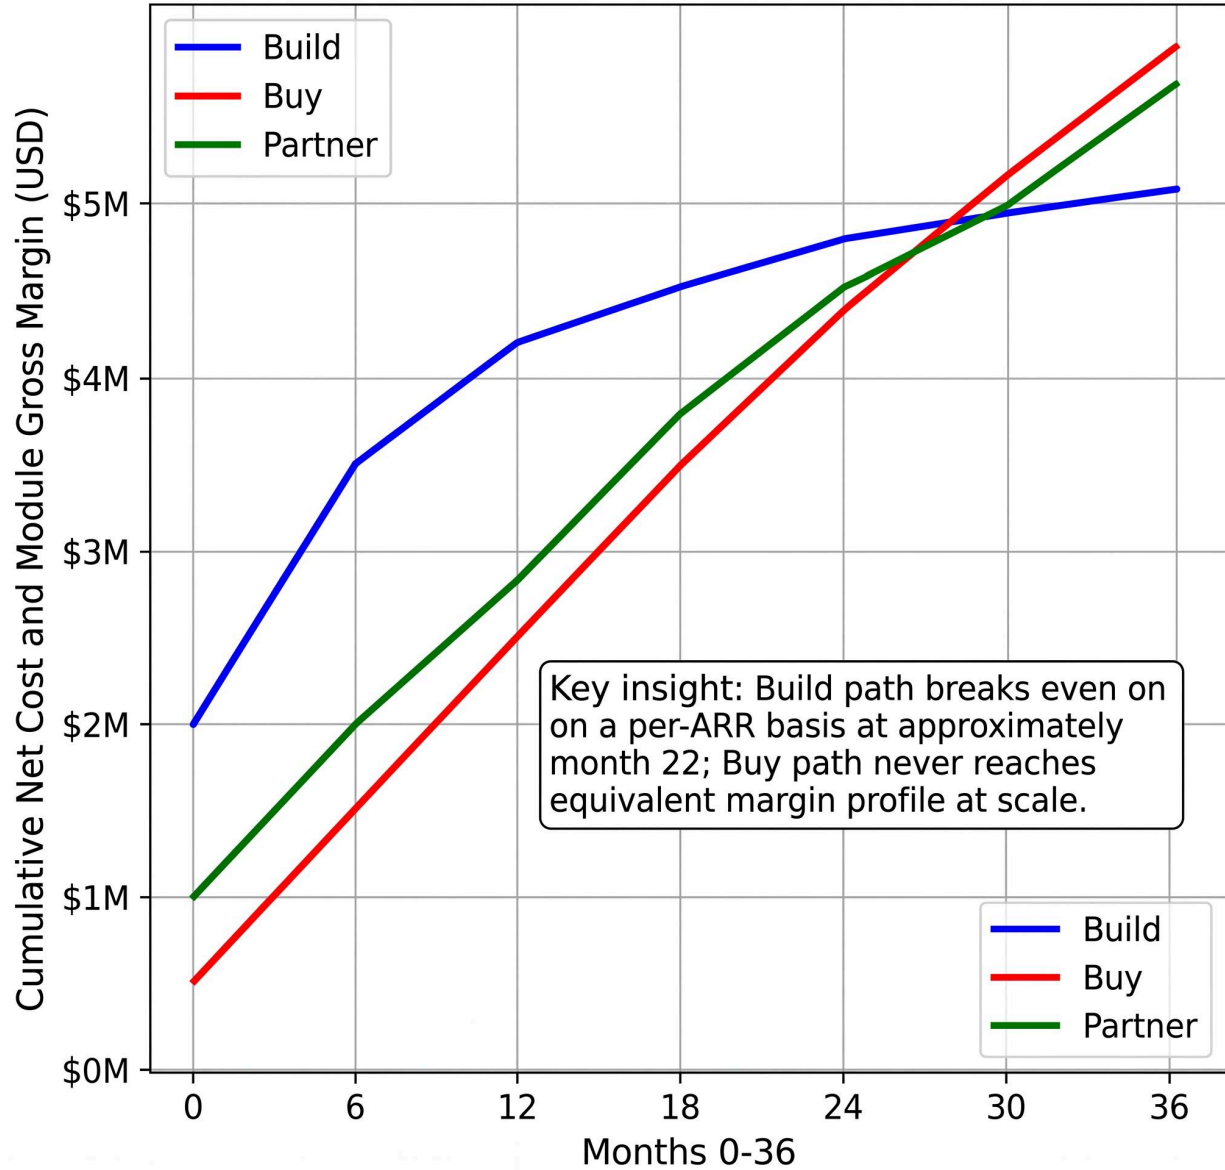
**Recommended path: Build, with a structured proof-of-concept gate before full engineering mobilization.**

The build path carries the highest execution risk in the near term but is the only path that delivers the combination of strategic control, long-term unit economics, and capability fit that the opportunity requires. The buy path's lower initial cost is offset by ongoing licensing fees that consume a disproportionate share of module revenue, eliminate the pricing flexibility needed to compete in the SMB segment, and create a hard dependency on a vendor whose roadmap will eventually compete with the Client directly.

The partner path is not recommended at this stage for two reasons. First, the Client has not yet validated the core hypothesis -- that customers will pay for this capability -- well enough to negotiate from a position of strength in a partnership. Second, IP co-ownership structures introduce governance complexity that is difficult to unwind if the partnership underperforms. A partnership arrangement may be appropriate in a later phase if specific data assets or distribution capabilities become material constraints, but those constraints are not present at this stage.

The build path recommendation is conditional on completing a 60-to-90 day technical proof of concept before committing the full engineering team. This gate exists to confirm two things: that the optimization approach works on real tenant data at acceptable quality, and that the integration points in the existing platform can support the real-time data exchange the module requires. A failed or degraded proof of concept would trigger re-evaluation of the partner path as an interim measure.

Chart 4: 3-Year Cumulative Cost and Margin Comparison



## SECTION 8

# Recommendation

---

### **Decision: Conditional Go**

The evidence gathered across this validation supports moving forward with an internally built AI dispatching and route optimization module. The recommendation is not unconditional. Three specific gates must be cleared before the full engineering team is mobilized and significant capital is committed. Until those gates are cleared, spending should be limited to the proof-of-concept phase and the design-partner outreach programme described in Section 10.

### **Full Reasoning**

The market opportunity is real and the window is open. The demand signals are strong enough to justify a build investment. The technical approach is sound. The financial model shows a positive return on a three-year horizon under conservative adoption assumptions. These are the conditions that typically support a go decision.

The conditions that make this a conditional go rather than a clean go are: data readiness is uneven and must be confirmed before the model quality can be guaranteed; willingness-to-pay has been stated but not transacted; and the competitive window is time-bounded. Each condition has a defined test.

The data readiness condition is tested by the account-level audit described in the next steps. If fewer than 60 percent of the top 200 accounts meet the minimum data completeness threshold for the initial model, the product strategy must be adapted: either launch with a narrower set of accounts and expand as data matures, or add a structured data collection workflow to the onboarding process before module access is granted.

The willingness-to-pay condition is tested by the design-partner programme. A design-partner commitment requires the customer to agree to a pre-commercial pricing arrangement, which converts stated intent into transacted intent. The target is three to five signed design-partner agreements before the proof of concept begins. Fewer than three means the demand hypothesis has not been confirmed and the build decision should be reconsidered.

The competitive window condition is tested against calendar. If the proof of concept does not begin within 90 days of this report's acceptance, the target launch date slips into direct competition with FSM Platform A's announced roadmap release. This is not automatically disqualifying, but it changes the competitive positioning from first-mover to capability-differentiation, which requires a different and harder argument.

### **Trade-offs Made Visible**

A build decision of this scale involves trade-offs that should be named explicitly. Building this module requires diverting a portion of the engineering team from existing platform development and maintenance. That diversion has an opportunity cost: features and improvements that the

existing customer base is expecting will be delayed. The severity of this trade-off depends on the Client's current engineering team size and capacity, which was not in scope for this validation.

The buy path is not recommended, but it is not irrational. If the engineering capacity required for the build path is not available and cannot be recruited or contracted within 90 days, the buy path offers a faster route to a market-facing product. The licensing economics are unfavourable at scale, but they are adequate for a market validation phase. A decision to license in the short term, with a stated intent to migrate to a built solution once demand is confirmed, is a defensible hedge and should be modelled before the build decision is locked.

### **Assumptions That Must Hold**

- The Client can mobilize a dedicated product and engineering team of eight to ten people within 60 days of approving the build decision.
- At least three design-partner commitments are secured before or concurrent with proof-of-concept initiation.
- The data readiness audit confirms that at least 60 percent of the top 200 accounts meet the minimum data completeness threshold.
- The proof of concept produces routing suggestions that design-partner dispatchers rate as equal to or better than their own decisions in at least 70 percent of test scenarios.
- No regulatory change or platform policy shift affects the Client's access to technician location data under existing customer terms.

**SECTION 9**

**Risk Register**

The following risks were identified through the validation process and assessed for likelihood and potential impact on the build decision. Severity is a composite of likelihood and impact. Mitigations describe concrete steps available to the Client, not general risk management principles.

#	Risk	Category	Likelihood	Impact	Severity	Proposed Mitigation
1	Data readiness below threshold: fewer than 60% of accounts have sufficient job history and GPS data completeness for model training	Technical	Medium	High	High	Commission account-level data audit in 30 days; build data collection workflow for accounts below threshold; launch with a qualified sub-set of accounts only
2	Design-partner commitments fall short: fewer than three signed agreements before proof of concept begins	Commercial	Medium	High	High	Approach design-partner candidates immediately; offer meaningful commercial incentive (extended free access, joint case study rights, product roadmap influence); treat below-three outcome as a build decision trigger for re-evaluation
3	Competitive acceleration: FSM Platform A releases AI dispatching ahead of Q4 2026	Market	Low	High	Medium	Monitor Platform A's release communications quarterly; if release occurs earlier than 6 months before the Client's planned

#	Risk	Category	Likelihood	Impact	Severity	Proposed Mitigation
	roadmap date					launch, accelerate the proof-of-concept timeline or re-evaluate the buy path as an interim bridge
4	Proof of concept underperforms: routing suggestions rated below the 70% dispatcher acceptance threshold in design-partner testing	Technical	Medium	High	High	Define test criteria and acceptance thresholds before proof of concept begins; if results fall below threshold, do not proceed to full build; re-evaluate model approach with an external ML advisor before retry
5	Engineering capacity unavailable: the Client cannot staff the required team within 60 days of build approval	Operational	Medium	Medium	Medium	Pre-identify contract engineering options; maintain contact with two to three specialist firms capable of providing operations research and ML engineering support; treat as a cost increase, not a blocker, if internal capacity is unavailable
6	Technician data privacy exposure: location and job data used for model training falls under a new regulatory interpretation in key markets	Regulatory	Low	High	Medium	Review data processing terms with legal counsel before proof of concept begins; confirm that existing customer agreements cover model training use of anonymized data; obtain legal opinion on CCPA and provincial Canadian privacy law applicability
7	Pricing compression:	Commercial	Medium	Medium	Medium	Price the module at a value-per-outcome

#	Risk	Category	Likelihood	Impact	Severity	Proposed Mitigation
	established competitors respond to the module launch with discounting on their own scheduling products					framing rather than a feature-comparison framing; establish outcome benchmarks in design-partner phase that anchor customer perception to efficiency gains rather than feature lists
8	Customer adoption below modelled rate: attach rate at 36 months reaches 35-40% rather than the modelled 55-65%	Commercial	Medium	Medium	Medium	Build adoption milestones into the financial model's go or no-go gates at months 12 and 24; do not fully staff the module team beyond the initial build phase until month-12 attachment data confirms trajectory
9	Model quality degradation over time: routing suggestion accuracy declines as operational patterns shift and the model is not retrained	Technical	Low	Medium	Low	Build model monitoring and retraining pipeline as part of the initial build scope; define minimum accuracy thresholds that trigger automated retraining; include ML ops in the initial team headcount
10	Platform architecture constraint: real-time data exchange requirements exceed the current platform's API throughput capacity	Technical	Low	High	Medium	Include platform API throughput testing in the proof-of-concept scope; identify and scope any required platform infrastructure upgrades before full build begins

**SECTION 10**

**Next Steps**

The following actions are sequenced by priority and dependency. Actions in Phase 1 must be completed before any Phase 2 spending is authorized. Phase 2 actions are conditional on Phase 1 outputs meeting the thresholds defined in the recommendation section.

**Phase 1: Validation Gate (Days 1-90)**

Action	Owner	Timeline	Success Threshold	Budget
Commission data readiness audit across top 200 accounts	Head of Data / CTO	Days 1-30	At least 60% of accounts meet minimum data completeness threshold	\$15,000-\$25,000 external audit cost
Identify and approach design-partner candidates	VP Sales / Product	Days 1-45	3-5 signed design-partner agreements	Internal resource; potential commercial incentive up to \$30,000 in deferred fees
Legal review of data processing terms for model training use	General Counsel / external legal	Days 1-30	Written legal opinion confirming data usage is permissible under existing agreements	\$8,000-\$15,000 legal fees
Authorize and scope the technical proof of concept	CTO / Engineering Lead	Days 15-30	Scoped PoC plan with defined acceptance criteria approved by leadership	Internal; no external cost unless contract engineering required
Execute proof of concept with design-partner data	Engineering team	Days 45-90	Routing suggestions rated at or above 70% acceptance threshold by dispatchers	\$120,000-\$180,000 in engineering time and infrastructure

**Phase 2: Build Initiation (Days 91-180, conditional)**

Phase 2 spending is authorized only if Phase 1 gates are cleared. The following actions initiate the full build programme.

Action	Owner	Timeline	Notes
Mobilize full product and engineering team (8-10 people)	CTO / HR	Days 91-120	Includes internal transfers and contract engineering as required; team must be at full capacity by day 120
Define product specification for v1.0 module release	Product Manager	Days 91-120	Based on design-partner feedback from Phase 1; must define dispatcher UX, accept or override workflow, and fallback behaviour for low-data accounts
Establish ML ops and model monitoring infrastructure	Data Engineering	Days 91-150	Must be built alongside the model, not after; retraining pipeline and accuracy monitoring dashboards required before customer launch
Define and sign design-partner commercial agreements	VP Sales / Legal	Days 91-120	Convert design-partner LOIs to commercial agreements with pricing, data rights, and feedback commitment terms
Set 12-month adoption milestone and go or no-go review date	CEO / CFO	Days 91-120	12-month review gate: attach rate at or above 35% of eligible accounts required to authorize continued investment in Phase 3

## Budget Summary

Phase	Estimated Investment	Authorization Trigger
Phase 1: Validation Gate	\$150,000-\$235,000	Board or leadership approval of this report
Phase 2: Build Initiation (months 3-6)	\$400,000-\$600,000	Phase 1 gates cleared; all three thresholds met
Phase 3: Full Build and Launch (months 7-18)	\$1,800,000-\$2,600,000	Month-12 adoption milestone met or exceeded
Total 18-month programme	\$2,350,000-\$3,435,000	Staged; exits available at each gate

*Budget figures are estimates based on loaded engineering costs at market rates for the Client's primary hiring market, external advisory and legal costs as outlined in Phase 1, and infrastructure costs at projected production scale. Figures should be refined with the Client's finance and engineering leadership before formal budget authorization.*

**APPENDIX**

**Appendix**

**A. Interview List (Anonymized)**

All interview participants are referenced by role and company profile only. No identifying information is included.

Ref	Role	Company Type	Technician Count	Geography	Category
INT-01	Operations Director	Commercial HVAC contractor	52	Midwest US	Existing customer
INT-02	General Manager	Residential plumbing and HVAC	28	Southeast US	Existing customer
INT-03	Lead Dispatcher	Electrical contractor (multi-branch)	87	Mid-Atlantic US	Existing customer
INT-04	Owner-Operator	Residential HVAC	19	Southwest US	Existing customer
INT-05	Service Manager	Commercial HVAC and refrigeration	44	Pacific Northwest US	Existing customer
INT-06	VP Operations	Multi-trade contractor	120+	Ontario, Canada	Existing customer
INT-07	Scheduling Coordinator	Residential plumbing	22	Texas	Existing customer
INT-08 to INT-14	[6 additional existing customer roles]	Various FSM companies	20-95	Various US geographies	Existing customer
INT-15 to INT-20	[6 non-customer roles]	Target segment; not current clients	25-110	US and Canada	Non-customer

Ref	Role	Company Type	Technician Count	Geography	Category
INT-21 to INT-23	[3 churned customer roles]	Former clients; churned 2023-2024	18-60	Various	Lapsed customer
ANA-01 to ANA-03	[3 industry analyst roles]	Independent FSM research and advisory	N/A	US-based	Analyst
ADV-01 to ADV-02	[2 technology advisor roles]	FSM software advisory	N/A	US-based	Advisor

## B. Data Sources

- IDC. (2024). Worldwide Field Service Management Software Forecast, 2024-2028. International Data Corporation.
- Gartner. (2025). Market Guide for Field Service Management. Gartner Research.
- Field Service Management Institute. (2024). Annual Operations Report: Efficiency Benchmarks for Field Service Companies. FSMI Publications.
- Customer interview synthesis, n=23. Clixlogix Consulting Practice primary research, Q1-Q2 2026.
- Client account and usage data (anonymized schema-level metadata). Provided by the Client under NDA, reviewed Q1 2026.
- Google Trends and organic search volume data. Keyword analysis conducted March 2026 for field service AI dispatch and related terms.
- LinkedIn Talent Insights. (Q1 2026). Job posting data for AI and ML roles in field service operations. LinkedIn Corporation.
- Van Westendorp, P. (1976). NSS-Price Sensitivity Measurement. Proceedings of the ESOMAR Congress. Applied here in adapted B2B software form.

## C. Framework References

- Clixlogix Consulting Practice Pre-Investment Product Validation Framework (internal methodology). Applied to opportunity scoping, evidence stream design, and recommendation structuring in this engagement.

- Vehicle Routing Problem with Time Windows (VRPTW). Standard operations research formulation. Reference implementation reviewed against Clarke and Wright (1964) savings algorithm and modern metaheuristic extensions.
- Build-Buy-Partner evaluation matrix. Adapted from standard technology sourcing frameworks; criteria and weighting defined by Clixlogix Consulting Practice for the SaaS product expansion context.

## D. Glossary

Term	Definition
SAM (Serviceable Addressable Market)	The portion of the total addressable market that the Client can realistically serve given its current product, distribution, and geography
SOM (Serviceable Obtainable Market)	The portion of the SAM that the Client can capture within a defined planning horizon under named assumptions
VRPTW	Vehicle Routing Problem with Time Windows; the operations research formulation used to describe the daily dispatching optimization problem
ML ops	Machine learning operations; the engineering discipline covering model deployment, monitoring, versioning, and retraining in production
Design partner	A customer who participates in the pre-commercial development of a product, providing feedback, testing access, and typically a preferential commercial arrangement in exchange
Attach rate	The proportion of existing customers who adopt an add-on module as a proportion of the total eligible customer base
WTP	Willingness to pay; the maximum price a buyer will accept for a defined product or service
FSM	Field Service Management; the software category covering scheduling, dispatch, work order, mobile workforce, and related functions for companies operating field-based employees
PoC	Proof of concept; a time-boxed engineering exercise to validate a technical approach before committing to full development
CCPA	California Consumer Privacy Act; US state privacy legislation with implications for how customer and worker data may be collected, processed, and used

## End of Report

Clixlogix Consulting Practice | [info@clixlogix.com](mailto:info@clixlogix.com) | [clixlogix.com](http://clixlogix.com)

*This document was prepared exclusively for [Client Confidential] and is subject to the confidentiality terms stated on the cover page.*