

Building Your Own AI Growth Engine

A research-backed plan for accelerating the ProPharma Distribution sales department with AI you build in-house

Prepared for Mark Snyder, Founder & Chairman, ProPharma Distribution

June 2026

Every factual claim in this report carries a bracketed citation to the Sources section. Where the evidence is a vendor's own case study, it is labeled as such. Where a number is our extrapolation rather than a published finding, it is explicitly labeled "our estimate." This report is business research, not legal advice; the legal section should be validated with telecom counsel before any AI-voice program launches.

Contents

Executive Summary

Your instinct is correct, and the data is on your side: the constraint on ProPharma's growth is not the number of reps you employ but the number of hours those reps actually spend selling — and what they do with those hours. Salesforce's global surveys found reps spending only 28% of their week selling in 2022, recovering to about 40% by late 2025 largely through AI-driven automation of the rest^[20, 22]. HubSpot's 2024 survey put it at roughly two hours of real selling per day^[23], and the one study that measured reps with tracking devices rather than surveys found 23% of the week spent selling^[24]. McKinsey estimates that more than 30% of sales tasks are automatable with current technology, and that early automators see 10–15% efficiency gains^[25]. In other words: there is close to a second sales force hiding inside your existing payroll, and it is reachable with software your team can build on systems you already own.

On your legal question — the belief that “AI calling is illegal” is understandable but inaccurate, and the distinction matters commercially. The FCC's February 2024 ruling did not ban AI calls; it classified AI-generated voices as “artificial or prerecorded voice” under the TCPA, which makes them *consent-regulated*, exactly like prerecorded messages have been for decades^[2, 3]. Everything your reps do today by hand remains untouched. AI whispering guidance into a live rep's ear (what Balto already does) is untouched. AI answering your inbound phone line is outside the rule entirely. And outbound AI calls become lawful the moment the customer has given prior express written consent — a checkbox you can collect during onboarding and portal signup. Section 1 maps this precisely.

The strategic conclusion that follows from both facts: **do not try to replace your sellers with AI — multiply them.** The one well-documented attempt at fully-autonomous AI selling at scale (the AI-SDR startup 11x) reportedly lost 70–80% of its customers^[28], and Gartner projects that while AI agents will outnumber human sellers tenfold by 2028, fewer than 40% of sellers will say agents improved their productivity — there is a “value ceiling” when AI is bolted on instead of aimed at the real bottlenecks^[27]. The pattern that does work, in McKinsey's phrase, is “humans and machines working together”^[25]: AI does the research, logging, drafting, verifying, and watching; humans do the conversations that close pharma accounts.

This report ranks six initiatives your team can build internally on Balto, RingCentral, HubSpot, Acumatica, BigCommerce, BambooHR, and your A2A agent platform — no third-party AI SaaS touching your data. The ranking weighs expected return, strength of evidence, and build cost together:

#	Initiative	Primary payoff	Build complexity
1	Playbook harmonization engine (Balto + RingCentral + HubSpot)	Mine what your best reps actually do, push it to everyone through Balto in real time; ends “everyone pulling in different directions”	Low–Medium — Balto is already live

2	Rep co-pilot: admin-offload agents	Auto CRM logging, call summaries, follow-up drafts, pre-call briefs — converts admin hours into selling hours; the single biggest capacity lever	Medium
3	Dormant-account radar and win-back (Acumatica)	Detect accounts going quiet before they are gone; lapsed customers convert at a multiple of cold prospects. Smallest build on this list — likely your first dollar of AI-driven revenue	Low
4	Compliance fast-lane onboarding	Automate licensure/DEA/credit verification so a clean account opens in hours, not weeks; speed-to-first-order becomes a competitive weapon and a consent-collection point	Medium
5	Nurse-friendly async channel: SMS reorder + logged-in portal	Reach the buyer your phones cannot — negotiated pricing stays behind login (BigCommerce price lists), so no public-pricing problem	Medium–High
6	Consent-based AI voice layer	Inbound AI order line, after-hours coverage, and customer-requested AI callbacks — built last, on the consent rails created by #4 and #5	High

These compound. The transcripts mined in #1 feed the briefs in #2. The consent checkboxes collected in #4 and #5 are what make #6 legal. The dormant-account alerts in #3 give your newly freed-up reps (#2) the highest-probability list to call. Built in this order, each initiative makes the next one cheaper.

1. The Legal Reality of AI Calling

The short version: AI voice calling in the United States is not banned. It is regulated by consent — and several of the most valuable uses for ProPharma require no new consent at all. This section is precise because the penalties for getting it wrong are real, and because knowing exactly where the lines are will let you build things your competitors assume are illegal.

1.1 What the FCC actually ruled in February 2024

On February 8, 2024, the FCC issued a unanimous Declaratory Ruling (FCC 24-17) confirming that voices generated or cloned by AI count as “artificial or prerecorded voice” under the Telephone Consumer Protection Act^[1, 2]. The FCC’s own press release used the word “illegal,” which is where the popular impression comes from — but law firms analyzing the ruling were quick to point out that the headline overstates it. As Mayer Brown put it, such calls are “legal but subject to a range of restrictions”^[3]. Specifically, an AI-voice call is treated exactly like a prerecorded-message call has been treated since 1991: it requires *prior express consent* for informational calls to cell phones, *prior express written consent* for marketing calls, plus identification of the caller at the start of the call and an opt-out mechanism^[2, 4]. Nothing in the ruling touches calls where a live human is the one talking.

1.2 What is restricted, and what is clearly permitted

The table below is the operational map. The legal trigger under TCPA § 227(b) is whether an artificial or prerecorded voice is *delivered to the called party* on an *outbound* call — not whether AI was involved anywhere in the process^[4, 3].

Scenario	Status	Why
Human rep manually dials and speaks (your model today)	Lawful	No artificial voice is delivered; § 227(b) is not triggered. Calls to business landlines are also outside the Do-Not-Call rules, which protect residential subscribers [4, 5]
AI real-time guidance/whisper during a live human call (Balto)	Lawful	The called party hears only a human. No FCC ruling addresses this expressly — it follows from the statute’s text, a view shared by practitioner commentary [3, 4]. Call-recording consent laws still apply, as they already do for your Balto deployment
AI answers your inbound phone line	Lawful	TCPA § 227(b) restricts calls that are made or initiated, not technologies that answer customer-initiated calls [2, 3]. California and Utah disclosure rules are easy to satisfy (Section 1.4)

Outbound AI-voice call to a customer who gave prior express written consent (requested callback, opted-in reorder reminder)	Lawful	Consent is the path the FCC built. Standard E-SIGN-compliant written consent suffices; the stricter “one-to-one” consent rule was vacated by the Eleventh Circuit in January 2025 and formally repealed by the FCC in September 2025 [9, 10]
Outbound AI-voice call to a clinic’s business landline, no consent	Gray zone	§ 227(b)(1)(B) covers residential lines, so business landlines are largely outside the federal prerecorded-voice rule [4] — but state mini-TCPAs can reach these calls, and the brand risk is yours. We recommend treating this as consent-required anyway
Outbound AI-voice or autodialed call/text to any cell phone without consent — including a nurse’s or pharmacist’s personal cell	Prohibited	This is the actual “robocall ban.” The statute protects the number type, not the relationship — there is no B2B exemption for cell phones [4, 5]. Courts have also presumed mixed-use (personal + business) cell phones are protected [6]

Two B2B nuances worth internalizing. First, the federal Do-Not-Call registry and most telemarketing-solicitation rules protect *residential* subscribers — calls by your reps to the business lines of clinics, surgery centers, and pharmacies are exempt^[5]. Second, the moment a number is a cell phone, § 227(b)’s artificial-voice and autodialer restrictions apply regardless of whether the owner is a business contact, and the Ninth Circuit has presumed that mixed-use cell phones count as residential for Do-Not-Call purposes^[6]. The safe operating rule for any automated outbound: **treat every 10-digit mobile number as protected unless you hold written consent for it.**

1.3 The consent path: how ProPharma makes AI calling legal

Because you onboard every customer through a credit application and compliance paperwork anyway, you have something most outbound shops do not: a natural, high-completion moment to collect consent. The mechanics:

- **Collect prior express written consent at onboarding and portal signup.** A clear, unbundled disclosure with an electronic signature or checkbox — e.g., “We may call and text you at the numbers provided, including with automated and AI-generated voice and text, about orders, pricing, and account matters. Consent is not a condition of purchase.” E-SIGN-compliant electronic consent is valid, and after the Eleventh Circuit’s ruling and the FCC’s September 2025 repeal, no one-to-one formality applies federally^[9, 10].

- **Honor revocation fast and everywhere.** Since April 11, 2025, customers may revoke consent “in any reasonable manner” and you must honor it within 10 business days^[7]. The harsher “revoke-all” provision (one opt-out kills consent across all message types) has been delayed to January 31, 2027 and may be modified before then — but building a single consent ledger in HubSpot now makes the question moot^[8].
- **Disclose the AI.** An FCC proposal from August 2024 would require consent forms to mention AI-generated calls and require in-call disclosure; it has not been finalized as of June 2026^[11], and the FCC’s current posture is deregulatory — its October 2025 rulemaking proposes deleting several TCPA rules^[12]. California already requires AI-voice disclosure on automated calls (AB 2905, effective January 2025)^[16], and Utah requires disclosure if asked^[17]. Practical rule: have every AI call open with “This is ProPharma’s automated assistant” — it satisfies the strictest current state law, future-proofs against the federal proposal, and customers prefer it anyway.
- **Screen state by state.** At least 15 states now run “mini-TCPA” regimes^[15]. Florida’s FTSA requires written consent for automated sales calls with a private right of action at \$500–\$1,500 per call and a 3-call-per-24-hours cap^[13]; Texas SB 140 (effective September 2025) extended its telemarketing act to texts with treble-damages exposure^[14]. Your A2A platform should carry a per-state rules table that gates any automated outbound by the destination number’s state.

The stakes if ignored are concrete: federal statutory damages run \$500 per call or text, trebled to \$1,500 for willful violations, with no cap on aggregate class exposure^[4]. In the one canonical AI-voice enforcement case — the fake-Biden New Hampshire primary robocalls — the transmitting carrier paid a \$1 million FCC penalty and the consultant behind the calls was fined \$6 million^[18, 19]. Those were cold, deceptive, non-consensual political robocalls to consumers — the exact opposite of a disclosed, consented, customer-requested callback. The lesson is not “don’t use AI voice”; it is “never call anyone with it who hasn’t asked you to.”

Bottom line: your entire current operation, plus Balto, plus inbound AI, plus consented outbound AI is lawful today. The initiative ranking below is built so that consent collection (initiatives 4 and 5) precedes any AI voice (initiative 6).

2. The Math of Your Sales Department

Two published patterns explain almost everything about why a 100-person, outbound-phone-driven distributor grows slower than its market allows: reps don't sell most of the day, and reps don't sell the same way as each other. Every initiative in Section 3 attacks one or both.

2.1 The selling-time problem

The evidence base is unusually consistent across methodologies. Salesforce's 5th-edition State of Sales (7,775 sales professionals, 2022) found reps spending just 28% of their week selling, with the rest consumed by "deal management and data entry"^[20]; its 6th edition (2024) phrased the same finding as 70% of time on non-selling tasks^[21]. HubSpot's 2024 survey of 1,400+ sales professionals produced the figure you quoted in your brief: about two hours per day of actual selling, with roughly an hour a day lost to pure administration — and, notably, reps using AI reported saving about two hours a day^[23]. The most rigorous source — Pace Productivity, which measured reps with logging devices rather than asking them — found 10.7 hours per week (23%) on selling, against an almost identical 10.5 hours on administration^[24]. And the newest data point cuts in your favor: Salesforce's 7th edition (February 2026) shows the average seller now at 40% selling time, with the improvement attributed largely to AI offload of the rest^[22] — evidence that the gap is actually closable, not just lamentable.

McKinsey's work says how much of the gap is buildable-away: more than 30% of sales activities can be automated with existing technology; early adopters report 10–15% efficiency improvements and sales uplift potential up to 10%; automated lead management alone can raise selling time 15–20%^[25]. Its generative-AI research adds that companies investing in AI for sales and marketing report revenue uplifts of 3–15%^[26].

2.2 The harmonization problem

"All pulling in different directions" is not a ProPharma quirk — it is the statistically normal state of a sales floor, and it is expensive. Cresta's analysis of 100,000+ recorded sales conversations found top performers outproducing bottom-quartile peers by 2.05x on average (up to 3.5x in some organizations), and estimated that lifting everyone to top-performer behavior would grow revenue roughly 25%^[35] (vendor research, but consistent with the peer-reviewed finding that individual performance follows a power-law, with stars roughly 4x average productivity^[36]). The behaviors that separate them are observable and teachable: in Cresta's data, top performers ask 2.1x more discovery questions and overcome more than 80% of objections versus 40% for the bottom quartile; reps who adhere to a standardized playbook generate 26% more revenue per conversation^[35]. Gong's analysis of millions of B2B calls adds the famous talk-to-listen finding — winners hover near 43:57, and win rates fall when reps talk more than ~65% of the call^[38]. Critically for a heterogeneous team, Harvard Business Review's analysis of sales coaching found the payoff concentrates in the *middle* 60% of performers — up to 19% improvement — rather than the stars or the laggards^[37]. Harmonization is therefore not about

cloning your best rep; it is about moving your large middle toward what your best rep demonstrably does.

3. The Build List, Ranked

Each initiative below follows the same structure: what it is, how it maps to systems you already own, why the evidence says it works, expected impact (with evidence and our estimates separated), and build complexity. “Build complexity” assumes what you told us: an in-house team that already ships custom AI automation on your A2A platform, with live API access to Balto (org 10959), RingCentral, HubSpot, BambooHR, Acumatica, and BigCommerce.

Initiative 1 — The Playbook Harmonization Engine

What it is

A closed loop that turns your calls into your playbook, and your playbook back into your calls. Stage one is mining: an A2A pipeline pulls every call transcript and outcome — Balto provides the transcripts and in-call behaviors, RingCentral the call metadata (duration, time of day, attempts, dispositions), HubSpot the downstream truth (did the deal advance, what was the order value), and BambooHR the rep roster and tenure so you can separate “what top reps do” from “what senior reps do.” The pipeline answers concrete questions: What do your top reps say in the first 30 seconds? How and when do they handle the licensure objection, the pricing question, the “we already have a supplier” brush-off? What is their talk ratio, question count, follow-up cadence? Stage two is broadcasting: the winning patterns become Balto dynamic prompts, checklists, and playbook updates pushed through the Balto API — so the guidance every rep hears mid-call is no longer a generic script but a continuously updated distillation of what is currently winning on your own floor. Stage three is measurement: Balto adherence scoring plus HubSpot outcomes tell you whether each playbook change actually moved conversion, which means you can A/B test the playbook itself.

Why it works — the evidence

This initiative stacks three independently evidenced effects. First, real-time guidance: Balto’s strongest published result is an A/B test at National General (~750 reps) where reps using Balto closed 16% more sales in 14 weeks (vendor-published, but experimental in design — the strongest methodology in this literature)^[31]. Second, coaching discipline: CSO Insights’ 2016 study — the actual primary source behind the “coaching lifts win rates ~28%” claim that circulates in sales blogs — found organizations with a *dynamic* coaching process showing win rates 27.9% above the study average, and quota attainment 10.2% higher^[30] (correlational survey research, so treat as an association, not a guarantee). Third, playbook standardization: Cresta’s conversation-mining research found script-adherent reps generating 26% more revenue per conversation with 21% better conversion^[35], and Balto documents a customer winning 25% more competitor objections from a single dynamic-prompt change deployed in 10 minutes^[32]. There is also a ramp-time dividend for new hires: Balto’s pharmacy-benefits customer EmpiRx cut new-hire ramp up to 50%^[33], and the only independent-methodology number in this space — Forrester’s Total Economic Impact model for conversation intelligence — also lands at 50% faster onboarding^[34]. A note on honesty: the “15–28% win-rate lift” and

“22–50% ramp compression” ranges in your prior research are composites assembled from these separate sources, not findings of any single study. The anchors above are the real, checkable numbers.

Expected impact

Evidence says: mid-teens conversion lift is the best-documented direct effect of real-time guidance (16%, A/B-tested)^[31]; the coaching and standardization literature brackets a 20–28% association at the high end^[30, 35]. **Our estimate:** because your team is unusually heterogeneous by your own description, ProPharma sits closer to the situation Cresta measured (2.05x top-to-bottom spread) than the average; harmonization initiatives pay most where dispersion is widest. If outbound-acquired revenue responds even at the conservative end of the evidence range, the effect on a ~\$90M base is a seven-figure annual swing. We deliberately do not put a dollar point-estimate on it — you will be able to measure the real number within one quarter, because the same pipeline that pushes the playbook also scores it.

Build complexity: Low–Medium

This is the highest-leverage-per-engineering-hour item on the list because Balto is already deployed and already has API access live. The new work is a transcript/outcome warehouse, the mining analytics, and the discipline of a weekly playbook-update cycle. No new customer-facing surface, no compliance exposure, no consent requirements. The main risk is organizational, not technical: someone senior must own the weekly “what’s winning” review, or the loop decays into a dashboard nobody reads. Given that the same warehouse feeds Initiatives 2 and 3, this is also the natural foundation build.

Initiative 2 — The Rep Co-Pilot: Admin-Offload Agents

What it is

A set of A2A agents that take over the non-selling work that currently eats most of each rep’s day. Four agents, in priority order. **The scribe:** after every RingCentral call, an agent consumes the Balto transcript and writes the HubSpot record automatically — contact updates, call notes, disposition, next step, follow-up task — with the rep approving rather than typing. **The drafter:** follow-up emails and quote summaries generated from the actual conversation (“per our call, here are the three items and your pricing”), pulling each account’s negotiated prices from Acumatica so reps stop re-keying quotes. **The briefer:** before each call on the day’s list, a one-screen brief: order history and margin from Acumatica, open deals and last-touch from HubSpot, what happened on the last three calls from the transcript warehouse, licensure/credit status, and a suggested opener from the Initiative-1 playbook. **The list-builder:** each morning, a ranked call list assembled from HubSpot pipeline stage, Acumatica reorder timing, and Initiative-3 dormancy alerts — so the two-hours-of-selling problem is attacked from both sides: more selling hours, aimed at better targets.

Why it works — the evidence

The size of the prize is the selling-time evidence in Section 2.1: 28–40% of time spent selling across Salesforce editions^[20, 22], two hours a day in HubSpot's survey^[23], 23% in the device-measured study — which also found reps spending as much time on administration as on selling^[24]. McKinsey's automation research maps the same territory from the supply side: over 30% of sales tasks automatable now, 10–15% efficiency gains for early adopters^[25]. The directional 2026 evidence is that this is exactly the mechanism behind the industry's selling-time recovery: Salesforce attributes the climb from 28% to 40% selling time largely to AI handling the rest, and reports sellers expecting AI agents to cut research time ~34% and email drafting ~36%^[22]. HubSpot's reps using AI self-reported saving ~2 hours/day^[23]. One important caution, and the reason this initiative is designed as approval-based assist rather than autonomy: Gartner finds fewer than 40% of sellers report productivity improvement from AI agents and warns of a “value ceiling”^[27], and projects over 40% of agentic-AI projects canceled by 2027^[29]. The failures concentrate where AI replaces the human conversation; the successes concentrate where it removes the typing around it.

Expected impact

Evidence says: AI offload is associated with roughly 2 saved hours per rep per day in survey data^[23] and 10–15% measured efficiency gains in McKinsey's adopter research^[25]. **Our estimate:** if your reps are near the industry's ~2 selling hours/day and the co-pilot returns even half of the published savings, you move toward 3–4 selling hours per rep per day — functionally a 50–100% increase in selling capacity with zero added headcount. That is the arithmetic behind calling this the biggest single lever; it is an extrapolation from survey-based savings figures, not a published case study of a pharma distributor, and you should instrument it (RingCentral talk-time per rep per day is the natural KPI) from week one.

Build complexity: Medium

Four integrations your team already has tokens for, plus LLM summarization/drafting your team already knows how to ship. The scribe and drafter are weeks of work each; the briefer depends on the Initiative-1 warehouse; the list-builder is mostly ranking logic. Two design rules from the evidence: keep a human approval step on everything customer-facing (the drafter proposes, the rep sends), and roll out rep-by-rep with talk-time measurement so adoption is provable — the Gartner failure mode is tools nobody uses^[27].

Initiative 3 — Dormant-Account Radar and AI-Assisted Win-Back

What it is

An agent that watches Acumatica order history and catches customers going quiet before anyone notices. For every active account it learns the normal rhythm — median days between orders, typical basket, seasonal pattern — then flags three conditions: overdue (no order at 1.5–2x the account's own interval), shrinking (order values or line-item breadth trending down, a

wallet-share leak to another distributor), and gone-dark (no orders in a defined window). Each flag generates a win-back brief pushed to HubSpot as a task for the owning rep: what the account bought, margin contribution, what changed, any service issues mined from call transcripts and notes, and a drafted re-engagement opener. The rep makes the call — a human-initiated call, with zero TCPA implications.

Why it works — the evidence

Win-back is one of the few sales motions with peer-reviewed support. Kumar, Bhagwat and Zhang's Journal of Marketing study of eight years of customer data found that winning back lost customers is reliably profitable, that a firm's data on its own defectors (purchase history, complaint records) makes them far more predictable targets than cold prospects, and that the most effective offers address the reason the customer left — a service fix beats a discount^[66]. The widely-quoted conversion heuristic — 60–70% probability of selling to existing customers, 20–40% to lapsed customers, 5–20% to new prospects — comes from Griffin & Lowenstein's Customer Winback citing a Marketing Metrics study; the underlying methodology was never published, so treat it as an industry heuristic rather than research, but note that even its lower bounds put lapsed customers at a multiple of cold-call conversion^[67]. Harvard Business Review's summary of the cost side: acquiring a new customer runs 5–25x the cost of retaining one^[68]. In distribution specifically, the evidence is vendor-published but directionally consistent: Proton.ai estimates distributors lose 4.3–15.1% of annual revenue to churn^[69], and Zilliant's case studies report an MRO distributor growing accounts by \$60M+ after deploying churn and wallet-share alerts on ERP transaction data, and an electronics distributor lifting customer revenue ~20%^[70].

Expected impact

Evidence says: churn silently costs distributors mid-single-digit to low-double-digit percentages of revenue (vendor estimate)^[69]; recovered accounts convert at multiples of cold outreach (heuristic + peer-reviewed support for profitability)^[67, 66]. **Our estimate:** on a ~\$90M base, even the bottom of the vendor-estimated churn range implies millions of dollars a year quietly walking out through accounts nobody flagged. A radar that recovers a fraction of that is very likely the fastest payback on this list — not because the ceiling is highest, but because the build is smallest and the target list is made of people who already know you, already passed compliance onboarding, and already have negotiated pricing on file. Time-to-first-recovered-order is plausibly weeks from go-live. This is the pilot we would build first to prove the A2A sales-agent pattern internally.

Build complexity: Low

One data source (Acumatica), straightforward per-account statistics (no ML required to start — an account's own order-interval history is the model), one HubSpot task integration, one brief-drafting prompt. The transcript-mining enrichment ("why did they leave") can arrive later via the

Initiative-1 warehouse. The main design decision is threshold tuning to keep flag volume credible — a radar that cries wolf gets ignored.

Initiative 4 — The Compliance Fast-Lane: Onboarding as a Conversion Weapon

What it is

Reframe the licensure/DEA/credit gauntlet from a back-office cost into the sales department's sharpest conversion lever. The build: a digital onboarding flow (hosted on your BigCommerce/portal surface, feeding Acumatica and HubSpot) where a new customer enters their state license number, DEA registration (if ordering controlled substances), and credit application once — and A2A verification agents do the rest in parallel. DEA registration validates against the DEA registrant database, which registrants like ProPharma can access as a downloadable dataset and validation tool expressly intended for in-house verification systems^[58]. State licensure checks run against state board lookups and NABP's licensure data services^[59], with NCPDP dataQ as programmatic enrichment for pharmacy customers^[60] and the industry's emerging Open Credentialing Initiative ATP credentials as a forward path^[61]. The credit application digitizes into Acumatica with auto-decisioning for clean, low-limit accounts and human review above thresholds. Exceptions route to your compliance staff; everything else completes while the prospect is still excited. The KPI is one number: **time from “yes” on the phone to first order shipped**. And one more thing rides along at zero marginal cost: the onboarding flow is where the TCPA-grade communications consent from Section 1.3 gets collected — the legal foundation for Initiatives 5 and 6.

Why it works — the evidence

The compliance obligations themselves are non-negotiable and now fully enforced: under DSCSA you may only transact with authorized trading partners and must verify their licensure/registration, with the wholesale-distributor enforcement exemption expired as of August 2025^[54, 55]; DEA expects documented know-your-customer diligence and suspicious-order monitoring for controlled substances^[56, 57]. So the question is not whether to do this work — only whether it takes weeks of document-chasing or hours of automated verification. On the value of speed, the classic evidence is Harvard Business Review's lead-response study of 1.25 million leads: firms that contacted a prospect within an hour were about seven times as likely to qualify them as firms waiting even an hour longer, and over sixty times as likely as those waiting a day^[62] — buyer enthusiasm decays fast, and every day of onboarding friction is decay time. McKinsey's study of corporate-client onboarding (banking, the closest well-studied analog to compliance-heavy B2B onboarding) found cycles up to 100 days and identified compression as a direct revenue and differentiation lever^[63]. In adjacent healthcare credentialing, automation vendors claim manual cycles of ~120 days collapsing to hours with automated primary-source verification^[64], and credit-application automation vendors claim roughly 2x faster credit decisions^[65] — both vendor claims, labeled as such. We could find no published benchmark for pharma-

wholesaler account-opening time specifically; that absence is itself useful, because it means nobody in your market is competing on a published number yet.

Expected impact

Evidence says: response speed compounds into conversion at dramatic multiples in the one large-sample study on the subject^[62]; onboarding compression is an established revenue lever in compliance-heavy B2B^[63]. **Our estimate:** the win is threefold and partly unquantifiable in advance: (a) fewer deals dying between verbal yes and first order — your reps' hardest-won conversions are currently exposed to days or weeks of risk; (b) rep hours returned — every hour a rep spends chasing a DEA cert PDF is an hour subtracted from the two they sell; (c) a marketable promise no competitor makes ("verified and ordering same-day"). We would expect the measured time-to-first-order to fall by at least half for clean accounts; that is our engineering judgment from the verification sources available, not a published statistic.

Build complexity: Medium

The hard part is not the AI — it is the data plumbing. The DEA dataset is a file feed, not a modern API; state boards vary from clean lookups to scrape-or-call; NABP's programmatic access is oriented to member boards, so your team should validate commercial access early^[59]. Document-extraction agents (reading uploaded licenses and certs) are well within your team's demonstrated capability. Phase it: DEA + your top-5 states by pipeline first, then the long tail. Keep humans on every rejection and every controlled-substance edge case — the automation should make your compliance people faster, never absent.

Initiative 5 — The Nurse Channel: SMS Reorder + Logged-In Portal

What it is

A way to sell to the buyer your phones structurally cannot reach. Nurses and clinical buyers at clinics and surgery centers cannot take calls mid-shift — but they can glance at a text between cases and tap a link at 7 p.m. The build has two halves that share one constraint: **negotiated pricing never appears in public**. The portal half: a logged-in BigCommerce storefront where each account's customer group maps to its Acumatica-negotiated price list — prices exist only behind authentication, which is exactly how the large distributors already solve the same problem (McKesson's SupplyManager surfaces contracted items per customer; Henry Schein runs negotiated programs through its e-commerce)^[52, 53]. Reorder-from-history is the killer feature: the agent pre-builds carts from Acumatica purchase history so a nurse's reorder is three taps. The SMS half: an opted-in (Section 1.3 consent, collected at onboarding) transactional texting agent on a 10DLC-registered number — order confirmations, ship notices, and a reorder nudge timed to each account's consumption rhythm ("You're likely running low on X — reply YES or tap to review your cart"). Inbound texts ("need 2 more boxes of the usual") get matched against order history and confirmed with a portal link. Reps stay attached to every account — the channel routes activity into HubSpot, and anything unusual becomes a rep follow-up, not a silent self-serve loss.

Why it works — the evidence (with the folklore removed)

First, honesty about the SMS statistics in your prior research: the famous “98% open rate” is folklore. SMS has no open-tracking mechanism — the figure traces to a 2010 white paper whose own author firm has since walked it back to roughly 55% of messages read^[40]; the “98%/45%” pairing entered circulation through a 2016 Gartner article that itself attributed the numbers to unnamed “various sources”^[39]. The defensible modern evidence is still strong, just smaller: 89% of consumers have opted in to texts from at least one business, 72% read business texts within five minutes, 82% reply within fifteen^[41]; 74% check text notifications within five minutes^[42]; and on measured platform data, SMS click-through runs roughly 9–15% versus ~2% for email^[43] (consumer-skewed data — no credible B2B-specific SMS benchmark exists; we flag that gap rather than paper over it).

Second, the persona case. There is no published study of nurses’ receptivity to supplier texting — we checked — so we argue it from three adjacent, well-evidenced facts: nurses are interrupted constantly during shifts (up to ~12–14 interruptions per hour in emergency and high-acuity settings), which is why synchronous sales calls fail^[46]; access to clinical buyers for sales reps has been collapsing for a decade (rep-accessible physicians fell from 77% to 44% in ZS’s tracking)^[47]; and B2B buyers at large increasingly prefer asynchronous, self-directed buying — Gartner’s 2025 survey puts rep-free preference at 67%^[48], McKinsey finds a stable third of buyers preferring digital self-serve at every stage, with 39% willing to place \$500K+ orders through self-serve or remote channels^[50, 51]. Two counterweights keep the design honest: Gartner also finds rep-free buyers experience more purchase regret^[48] and predicts a swing back toward human-prioritized buying by 2030^[49] — which is why this channel is built to keep your reps attached to every account rather than replace them.

Third, compliance is solvable but real: marketing texts to cell numbers require prior express written consent even when the recipient is a business contact^[44]; the sending number must be 10DLC-registered with carriers^[45]; STOP/HELP keywords and 10-business-day revocation handling are mandatory^[7]. Your onboarding flow (Initiative 4) collects the consent; your A2A consent ledger enforces it.

Expected impact

Evidence says: a third of B2B buyers already prefer self-serve at every purchasing stage^[50]; texting reaches people within minutes who cannot take calls at all^[41]; the large distributors’ portals demonstrate the model works at scale in this exact industry^[52]. **Our estimate:** this is the only initiative that grows the addressable market rather than the efficiency of serving it — clinics and ASCs whose buyers never answered a rep’s call become winnable, and existing accounts’ reorder revenue becomes more defensible (a three-tap reorder is a switching cost). Revenue impact is genuinely unforecastable from published evidence — no healthcare-supply SMS-reorder case study exists — so we recommend treating phase one (transactional SMS + reorder portal for 20 friendly accounts) as the experiment that prices the channel.

Build complexity: Medium–High

BigCommerce customer-group price lists synced from Acumatica are well-trodden e-commerce engineering, but pricing-sync correctness is unforgiving — a wrong price shown once costs trust that took years to build. The SMS agent itself is modest (RingCentral provides the messaging rails); 10DLC registration takes days-to-weeks of lead time. Phase: portal reorder for existing accounts → transactional SMS → conversational reorder → (only with marketing consent) proactive replenishment nudges.

Initiative 6 — The Consent-Based AI Voice Layer

What it is

The initiative your competitors think is illegal — built strictly inside the lines from Section 1. Three applications, in order of risk. **Inbound AI agent:** an AI receptionist on your RingCentral line for after-hours and overflow — order status, reorder capture against Acumatica history, callback scheduling, warm routing to on-call staff. Outside TCPA entirely (it answers; it never initiates)^[3]. **Requested AI callbacks:** when a customer asks for a callback (via portal, SMS, or the inbound agent), an AI voice agent returns the call at the chosen time, handles the routine portion (order confirmation, delivery scheduling), and hands anything substantive to a rep. The customer's request is the consent for that responsive call; the onboarding consent (Initiative 4) covers it in writing. **Opted-in reorder-reminder calls:** for accounts that prefer a call to a text — some clinic back-offices do — the same replenishment logic as Initiative 5 delivered by AI voice, only to numbers with prior express written consent on file. Every AI call opens with identification (“This is ProPharma’s automated assistant”), satisfying California’s AB 2905 and the FCC’s proposed-but-unfinalized disclosure rule^[16, 11], offers a human transfer, and logs to the consent ledger. What this initiative never does: cold outbound AI calls. Not because all of them would be illegal — business landlines are a federal gray zone (Section 1.2) — but because the legal margin is thin, state laws vary, and the brand cost of being “the robocall distributor” in a relationship industry outweighs the reach.

Why it works — the evidence

The legal architecture is Section 1’s: inbound AI outside § 227(b)^[3], consented outbound squarely lawful under the FCC’s own framework^[2], written-consent mechanics simplified by the 2025 repeal of one-to-one consent^[10]. The commercial logic is coverage arithmetic: your sales day ends at 5 p.m.; clinic ordering needs do not, and the nurse persona’s free moments cluster outside business hours (Section 3, Initiative 5 evidence). An after-hours order captured by an AI agent is revenue that currently goes to whichever competitor answers first the next morning — or to a McKesson portal that never sleeps^[52]. We are deliberately sparing with vendor success statistics here: published “AI phone agent” results in 2026 are mostly unaudited marketing, and the honest evidence base for voice-AI revenue impact in B2B distribution is thin. This initiative is justified by coverage and consent economics, not by a borrowed case study.

Expected impact

Evidence says: the legal path is clear and documented; the channel-coverage gap is real (Section 1; Initiative 5 evidence). **Our estimate:** impact is concentrated in (a) after-hours and weekend order capture, (b) callback responsiveness — the HBR speed evidence^[62] suggests answering a callback request within minutes instead of next-day materially raises conversion, and (c) rep leverage, since every routine call an AI handles is rep time returned. We rank it last not because the ceiling is low but because its prerequisites (consent rails, order-history agents, brand trust in automation) are exactly what Initiatives 1–5 build, and because voice AI done badly is the most publicly visible way to fail. Gartner's agentic-AI cancellation forecast^[29] argues for scoping the pilot to one use case — the after-hours inbound line — and measuring captured orders before expanding.

Build complexity: High

Real-time voice AI (latency, interruption handling, graceful failure to human voicemail) is the hardest engineering on this list, even with RingCentral's telephony APIs and your A2A platform hosting the logic. The consent ledger, state-rules gate, and disclosure scripting must ship before the first outbound AI call — they are part of the build, not an afterthought. Have telecom counsel review the consent language and call flows before launch; the cost of that review is trivial against \$500–\$1,500-per-call statutory exposure^[4].

4. Sequencing: How the Six Compound

First 90 days. Build the transcript/outcome warehouse (the foundation of Initiative 1) and ship the dormant-account radar (Initiative 3). The radar is the confidence play: smallest build, fastest revenue, and it proves the A2A agent-to-HubSpot-task pattern that everything else reuses. In parallel, run the first playbook-mining pass and push the first evidence-based prompt updates into Balto — the Balto case literature shows prompt changes can be deployed and measured in days^[32].

Months 3–9. Roll out the rep co-pilot (Initiative 2) one agent at a time — scribe first, since auto-logging is the most universally resented admin task — measuring RingCentral talk-time per rep as the headline KPI. Build the compliance fast-lane (Initiative 4) starting with DEA validation plus your top pipeline states, and start collecting communications consent in the onboarding flow from day one, even though nothing uses it yet: consent collected now is optionality banked for Initiatives 5 and 6.

Months 9–18. Launch the portal-plus-SMS channel (Initiative 5) with a friendly-account pilot, then the after-hours inbound AI line (Initiative 6, first use case only). By this point the playbook loop has had several quarters of iterations, the co-pilot has expanded selling hours, and both new channels inherit consent, order-history agents, and a sales floor accustomed to working with AI rather than around it.

A final framing for the board: none of this is speculative platform-building. Each step is a measurable intervention on a known bottleneck — selling hours, conversation quality, onboarding latency, channel coverage, churn detection — using data and systems already paid for. The published evidence brackets the upside; your own instrumentation will price it precisely within a quarter or two of each launch.

5. What the Evidence Does Not Say (Read Before Quoting This Report)

Deep research is partly the art of finding out which famous numbers are real. The following claims from the prior research — and from sales-tech marketing generally — should not be repeated as fact:

- **“SMS has a 98% open rate.”** Unmeasurable (SMS has no open tracking); traces to a 2010 white paper its own authors later walked back to ~55% read^[40]. Use the modern survey figures instead: 72% of business texts read within five minutes^[41].
- **“45% SMS response rate.”** Appears only as a secondhand “various sources report” figure in a 2016 Gartner article^[39]; no primary methodology exists.
- **“AI coaching lifts win rates 15–28% and cuts ramp 22–50%.”** Composite ranges, not findings of any study. The real anchors: CSO Insights’ 27.9% win-rate association with

dynamic coaching^[30], Balto's 16% A/B-tested sales lift^[31], and ~50% ramp reductions in Forrester's model and Balto's pharmacy-benefits case^[34, 33]. We found no source at all for a "22%" ramp figure.

- **"60–70% of sales to existing customers close, vs. 5–20% for prospects."** A real published figure — but from a 2001 book citing a consultancy study whose methodology was never released^[67]. Useful heuristic; not research.
- **"Nurses prefer texting from vendors."** No study exists. Our argument is an inference from interruption research^[46], collapsing rep access^[47], and general B2B self-serve preference^[48] — strong circumstantially, unproven directly. The friendly-account pilot in Initiative 5 is how you turn it into your own data.
- **Pharma-wholesaler onboarding cycle times.** No published benchmark exists for account-opening speed at McKesson, Cardinal, or ABC. The 100-day banking figure^[63] and ~120-day credentialing figure^[64] are analogies, not industry data.
- **Vendor case studies throughout** (Balto, Cresta, Gong, Zilliant, Proton.ai, Qualified) are labeled where used. They are directionally informative and methodologically self-interested. The Balto National General A/B test is the strongest of them because of its experimental design^[31].

Sources

- [1] FCC news release, “FCC Makes AI-Generated Voices in Robocalls Illegal” (Feb. 8, 2024). <https://www.fcc.gov/document/fcc-makes-ai-generated-voices-robocalls-illegal>
- [2] FCC Declaratory Ruling, CG Docket No. 23-362, FCC 24-17 (Feb. 8, 2024). <https://docs.fcc.gov/public/attachments/FCC-24-17A1.pdf>
- [3] Mayer Brown, “FCC Declares Authority and Intent to Regulate AI-Generated Calls under the TCPA” (Feb. 2024). <https://www.mayerbrown.com/en/insights/publications/2024/02/fcc-declares-authority-and-intent-to-regulate-ai-generated-calls-under-the-tcpa>
- [4] 47 U.S.C. § 227 (Telephone Consumer Protection Act), Cornell Legal Information Institute. <https://www.law.cornell.edu/uscode/text/47/227>
- [5] DNC.com, “Are B2B calls exempt from TCPA regulations?”. <https://www.dnc.com/faq/are-b2b-calls-exempt-tcpa-regulations>
- [6] Mac Murray & Shuster, “Ninth Circuit’s Do Not Call Decision Effects B2B Callers” (re Chennette v. Porch.com). <https://mslawgroup.com/ninth-circuits-do-not-call-decision-effects-b2b-callers/>
- [7] Klein Moynihan Turco, “FCC’s TCPA Consent Revocation Rule Effective April 11, 2025”. <https://kleinmoynihan.com/fccs-tcpa-consent-revocation-rule-effective-april-11-2025/>
- [8] Burr & Forman, “FCC Delays Effective Date of TCPA Revoke-All Rule Until January 31, 2027” (Jan. 2026). <https://www.burr.com/telephone-consumer-protection-act/the-fcc-delays-effective-date-of-tcpa-revoke-all-rule-until-january-31-2027>
- [9] Wiley, “11th Circuit Vacates FCC’s One-to-One TCPA Consent Rule” (Insurance Marketing Coalition v. FCC, Jan. 2025). <https://www.wiley.law/alert-UPDATE-11th-Circuit-Vacates-FCCs-One-to-One-TCPA-Consent-Rule>
- [10] Goodwin, “The FCC Issues Final Rule Formally Eliminating the One-to-One Consent Requirement” (Sept. 2025). <https://www.goodwinlaw.com/en/insights/blogs/2025/09/the-fcc-issues-final-rule-formally-eliminating-the-one-to-one-consent-requirement>
- [11] FCC, “FCC Proposes First AI-Generated Robocall & Robotext Rules” (NPRM, Aug. 7, 2024; not finalized as of June 2026). <https://www.fcc.gov/document/fcc-proposes-first-ai-generated-robocall-robotext-rules-0>
- [12] Kelley Drye, “FCC Includes Proposals to Roll Back Several TCPA Rules as Part of Call Branding Initiative” (Oct. 2025). <https://www.kelleydrye.com/viewpoints/blogs/ad-law-access/look-whos-calling-fcc-includes-proposals-to-roll-back-several-tcpa-rules-as-part-of-call-branding-initiative>
- [13] Florida Telephone Solicitation Act, Fla. Stat. § 501.059. <https://www.flsenate.gov/laws/statutes/2021/501.059>
- [14] Paul Hastings, “Marketing Texts in Texas: SB 140 Broadens State Telemarketing Regulations” (2025). <https://www.paulhastings.com/insights/ph-privacy/marketing-texts-in-texas-sb-140-broadens-state-telemarketing-regulations>
- [15] Kaufman Dolowich, “State Mini TCPA Laws Growing — Texas Latest to Update Its Telemarketing Rules” (Aug. 2025). <https://www.kaufmandolowich.com/news-resources/>
- [16] California AB 2905 (2024) — AI-voice disclosure for automated calls, Cal. Pub. Util. Code § 2874. https://leginfo.ca.gov/faces/billNavClient.xhtml?bill_id=202320240AB2905
- [17] Davis Wright Tremaine, “Utah Enacts Consumer Disclosure Requirements for Businesses Using Bots and Generative AI” (Apr. 2024).

<https://www.dwt.com/blogs/artificial-intelligence-law-advisor/2024/04/utah-enacts-ai-and-bot-business-disclosure-law>

- [18] FCC consent decree — Lingo Telecom, \$1,000,000 civil penalty for AI-voice robocall transmission (Aug. 2024). <https://docs.fcc.gov/public/attachments/DOC-404951A1.pdf>
- [19] FCC forfeiture order — Steve Kramer, \$6,000,000 for the New Hampshire AI “Biden” robocalls (Sept. 2024). <https://docs.fcc.gov/public/attachments/DOC-405811A1.pdf>
- [20] Salesforce, State of Sales, 5th ed. — “Sales Reps Spend Less than 30% of Their Time Actually Selling” (Dec. 2022; n=7,775). <https://www.salesforce.com/news/stories/sales-research-2023/>
- [21] Salesforce, State of Sales, 6th ed. — 70% of rep time on non-selling tasks (July 2024; n=5,500). <https://www.salesforce.com/news/stories/sales-ai-statistics-2024/>
- [22] Salesforce, State of Sales, 7th ed. (Feb. 2026; n=4,050) — sellers at 40% selling time; Agentforce SDR results. <https://www.salesforce.com/news/stories/state-of-sales-report-announcement-2026/>
- [23] HubSpot, 2024 Sales Trends Report (survey of 1,400+ sales professionals, Aug. 2023) — “reps spend about 2 hours per day selling”. <https://www.hubspot.com/hubfs/HubSpots%202024%20Sales%20Trends%20Report.pdf>
- [24] Pace Productivity, “How Sales Reps Spend Their Time” — device-measured time studies: 10.7 hrs/week (23%) selling. <https://www.paceproductivity.com/single-post/2017/02/09/how-sales-reps-spend-their-time>
- [25] McKinsey & Company, “Sales automation: The key to boosting revenue and reducing costs” (2020). <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/sales-automation-the-key-to-boosting-revenue-and-reducing-costs>
- [26] McKinsey & Company, “AI-powered marketing and sales reach new heights with generative AI” (2023). <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/ai-powered-marketing-and-sales-reach-new-heights-with-generative-ai>
- [27] Gartner press release, “By 2028 AI Agents Will Outnumber Sellers by 10x, Yet Fewer than 40% of Sellers Will Report AI Agents Improved Productivity” (Nov./Dec. 2025). <https://www.gartner.com/en/newsroom/press-releases/2025-11-18-gartner-predicts-by-2028-ai-agents-will-outnumber-sellers-by-10x-yet-fewer-than-40-percent-of-sellers-will-report-ai-agents-improved-productivity>
- [28] TechCrunch, “a16z- and Benchmark-backed 11x has been claiming customers it doesn’t have” (Mar. 24, 2025) — investigation of a fully-autonomous AI SDR vendor. <https://techcrunch.com/2025/03/24/a16z-and-benchmark-backed-11x-has-been-claiming-customers-it-doesnt-have/>
- [29] Gartner press release, “Over 40% of Agentic AI Projects Will Be Canceled by End of 2027” (June 2025). <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>
- [30] CSO Insights (Miller Heiman Group), 2016 Sales Enablement Optimization Study — dynamic coaching associated with 27.9% higher win rates, +10.2% quota attainment. <https://cdn2.hubspot.net/hubfs/4460233/SocialSellinator-June2019/PDF/2016-Sales-Enablement-Optimization-Study.pdf>
- [31] Balto case study — National General (~750 reps): A/B test, reps using Balto closed 16% more sales in 14 weeks (vendor-published). <https://www.balto.ai/case-study/how-better-discovery-questions-led-to-16-more-sales-for-national-general/>

- [32] Balto case study — one dynamic-prompt change won 25% more competitor objections, deployed in 10 minutes (vendor-published). <https://www.balto.ai/case-study/how-one-script-update-won-25-more-competitor-objections-2/>
- [33] Balto case study — EmpiRx (pharmacy benefits): new-hire ramp time cut up to 50% (vendor-published). <https://www.balto.ai/case-study/how-empirx-scales-faster-and-delivers-white-glove-service-with-balto/>
- [34] Forrester (commissioned by Gong), Total Economic Impact study — 50% reduction in rep onboarding time (modeled composite). <https://www.dock.us/revenue-archives-asset/gong-economic-impact-study>
- [35] Cresta, “Quantifying the Performance Gap” (analysis of 100,000+ recorded conversations, 2021; vendor-published). <https://cresta.com/blog/quantifying-the-performance-gap>
- [36] O’Boyle & Aguinis, “The Best and the Rest”, Personnel Psychology (2012) — individual performance follows a power-law distribution (peer-reviewed). <http://www.hermanaguinis.com/PPPsych2012.pdf>
- [37] Dixon & Adamson, “The Dirty Secret of Effective Sales Coaching”, Harvard Business Review (Jan. 2011). <https://hbr.org/2011/01/the-dirty-secret-of-effective>
- [38] Gong Labs research — talk-to-listen ratios of top performers (~43:57) across large samples of recorded B2B sales calls (vendor data science). <https://www.gong.io/blog/talk-to-listen-conversion-ratio>
- [39] Gartner, “Tap Into the Marketing Power of SMS” (Nov. 2016) — reports third-party figures of 98% open / 45% response for SMS vs. 20% / 6% for email. <https://www.gartner.com/en/marketing/insights/articles/tap-into-the-marketing-power-of-sms>
- [40] Rally Corp, “98% SMS Open Rates Are a Myth”; MobileSquared 2022 update walking back its own 2010 figure (~55% of SMS read; SMS has no open tracking). <https://www.rallycorp.com/blog/90-sms-open-rates-are-a-myth>
- [41] EZ Texting, 2026 Consumer Texting Behavior Report (n=959): 89% of consumers opted in to texts from at least one business; 72% read business texts within 5 minutes; 82% reply within 15 minutes (vendor survey, disclosed methodology). <https://www.eztexting.com/report/2026-consumer-texting-behavior-report>
- [42] SimpleTexting, Texting & SMS Marketing Statistics (2026 ed.): 74% of consumers check text notifications within 5 minutes (vendor survey). <https://simpletexting.com/blog/texting-and-sms-marketing-statistics/>
- [43] Klaviyo platform benchmarks — SMS campaign click-through ~8.9–14.5% vs. email ~2% (measured platform data, mostly B2C). <https://www.klaviyo.com/products/sms-marketing/benchmarks>
- [44] Shumaker Williams, “B2B SMS Marketing: Navigating CAN-SPAM and TCPA Regulations” — no blanket B2B exemption for texts to cell numbers. <https://www.shumakerwilliams.com/business-to-business-sms-marketing-navigating-can-spam-and-tcpa-regulations/>
- [45] Twilio, A2P 10DLC compliance documentation — carrier registration requirements for business SMS. <https://www.twilio.com/docs/messaging/compliance/a2p-10dlc>
- [46] Peer-reviewed integrative review of nurse interruptions — 0.4–13.9 interruptions/hour; emergency-department nurses ~12/hour (PMC). <https://pmc.ncbi.nlm.nih.gov/articles/PMC4376046/>
- [47] ZS Associates AccessMonitor (via Fierce Pharma) — rep-accessible physicians fell from ~77% (2008) to 44% (2016); majority now restrict sales-rep access. <https://www.fiercepharma.com/marketing/crossing-threshold-more-than-half-of-physicians-restrict-access-to-sales-reps>

- [48] Gartner press release (Mar. 2026): 67% of B2B buyers prefer a rep-free buying experience (Aug.–Sep. 2025 survey; Gartner notes rep-free buyers show higher purchase regret). <https://www.gartner.com/en/newsroom/press-releases/2026-03-09-gartner-sales-survey-finds-67-percent-of-b2b-buyers-prefer-a-rep-free-experience>
- [49] Gartner press release (Aug. 2025): by 2030, 75% of B2B buyers will prefer sales experiences that prioritize human interaction over AI. <https://www.gartner.com/en/newsroom/press-releases/2025-08-25-gartner-says-by-2030-that-75-percent-of-b2b-buyers-will-prefer-sales-experiences-that-prioritize-human-interaction-over-ai>
- [50] McKinsey B2B Pulse 2024, “Five fundamental truths: How B2B winners keep growing” — the “rule of thirds” across in-person, remote, and digital self-serve. <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/five-fundamental-truths-how-b2b-winners-keep-growing>
- [51] Digital Commerce 360 (Sept. 2024), on McKinsey B2B Pulse — 39% of B2B buyers willing to spend \$500K+ per order via self-serve or remote channels. <https://www.digitalcommerce360.com/2024/09/12/more-b2b-buyers-are-willing-to-spend-big-bucks-per-online-order/>
- [52] McKesson Medical-Surgical SupplyManager — self-serve ordering portal that surfaces contracted vs. open-market items per logged-in customer. <https://mms.mckesson.com/content/inventory-management/supplymanager/>
- [53] Henry Schein e-commerce and Special Markets — self-serve ordering with negotiated/contracted customer programs. <https://www.henryschein.com/us-en/specialmarkets/default.aspx>
- [54] FDA guidance, “Identifying Trading Partners Under the DSCSA” — prescription drugs may only be transacted with authorized trading partners (licensed/registered); verification records retained. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/identifying-trading-partners-under-the-drug-supply-chain-security-act>
- [55] FDA, DSCSA waivers and exemptions beyond the stabilization period — wholesale-distributor exemption expired Aug. 27, 2025; requirements fully enforceable. <https://www.fda.gov/drugs/drug-supply-chain-security-act-dscsa/waivers-and-exemptions-beyond-stabilization-period>
- [56] Henry Schein, DEA “Know Your Customer” welcome packet — a real distributor onboarding artifact: questionnaires, DEA/state license verification, site visits, annual re-verification. <https://www.henryschein.com/us-en/images/medical/dea-welcome-packet-kyc.pdf?cdivId=medical>
- [57] Ohio Board of Pharmacy, Suspicious Order Monitoring and Due Diligence guidance (updated Dec. 2024) — 21 C.F.R. § 1301.74(b) obligations. https://www.pharmacy.ohio.gov/Documents/Pubs/Special/SUSPICIOUS_ORDER/Suspicious%20Order%20Monitoring%20and%20Due%20Diligence.pdf
- [58] DEA CSA Registration Database (formerly via NTIS; in-house at DEA since Nov. 2020) — registrant validation tool and downloadable registrant dataset for use in in-house verification systems. <https://www.deadiversion.usdoj.gov/drugreg/registration.html>
- [59] NABP licensure services (e-Profile, Clearinghouse, NABP Verify) — multistate pharmacy licensure and disciplinary data. <https://nabp.pharmacy/members/programs-services/licensure/>
- [60] NCPDP dataQ — continuously verified pharmacy database (NCPDP Provider ID, NPI, licensure documentation) with programmatic access (vendor description). <https://dataq.ncdpd.org/>
- [61] Open Credentialing Initiative — interoperable digital Authorized Trading Partner credentials for DSCSA. <https://www.oc-i.org/dscsa-authorized-trading-partners>

- [62] Oldroyd, McElheran & Elkington, “The Short Life of Online Sales Leads”, Harvard Business Review (Mar. 2011) — 1.25M leads: firms contacting within an hour were ~7x as likely to qualify the lead as those waiting even an hour more, and 60x+ vs. waiting 24 hours. <https://hbr.org/2011/03/the-short-life-of-online-sales-leads>
- [63] McKinsey, “Winning corporate clients with great onboarding” — corporate-client onboarding averages up to 100 days; compression is a revenue lever. <https://www.mckinsey.com/industries/financial-services/our-insights/winning-corporate-clients-with-great-onboarding>
- [64] CertifyOS / Verisys / MedTrainer — vendor-published claims that manual healthcare credentialing takes ~120 days and automated primary-source verification compresses it to hours (vendor claims). <https://www.certifyos.com/resources/blog/manual-vs-automated-provider-credentialing>
- [65] HighRadius / Billtrust — vendor-published claims on B2B credit-application automation (e.g., 2x faster credit decisions; high automation of manual steps) (vendor claims). <https://www.highradius.com/product/automated-credit-application-software/>
- [66] Kumar, Bhagwat & Zhang, “Regaining ‘Lost’ Customers”, Journal of Marketing 79(4) (2015) — 8 years of data: win-back is profitable; first-lifetime behavior predicts who returns; service-fix + offer combinations work best (peer-reviewed). <https://journals.sagepub.com/doi/10.1509/jm.14.0107>
- [67] Griffin & Lowenstein, Customer Winback (2001), citing a Marketing Metrics study — probability of selling: existing customers 60–70%, lapsed customers 20–40%, new prospects 5–20% (widely cited heuristic; primary methodology never published). <https://www.amazon.com/Customer-Winback-Recapture-Customers-Loyal/dp/0787946672>
- [68] Gallo, “The Value of Keeping the Right Customers”, Harvard Business Review (Oct. 2014) — acquiring a new customer is 5–25x more expensive than retaining one, depending on study and industry. <https://hbr.org/2014/10/the-value-of-keeping-the-right-customers>
- [69] Proton.ai executive brief — estimates B2B distributors lose ~4.3–15.1% of annual revenue to customer churn (vendor estimate). <https://welcome.proton.ai/hubfs/Executive%20Brief%20-%20AI%20in%20Wholesale%20Distribution-2.pdf>
- [70] Zilliant case studies — MRO distributor grew accounts \$60M+ using churn/wallet-share alerts on ERP transaction data; electronics distributor +~20% customer revenue (vendor case studies). <https://resources.zilliant.com/case-studies/mro-distributor-reduces-customer-churn-grows-revenue-more-than-60million>