Where Should Sales Really Focus

Idea In Short

Most sales organizations spend more time defending how territories were drawn than using them to win business. The fix is not another map redraw. It is a disciplined method for scoring accounts on realized value, latent potential and cost to serve, then assigning capacity to match. Leaders who get this right stop treating headcount as the primary growth lever and start treating allocation as one. The immediate decision for any sales leader reading this: audit whether your top quartile of accounts by potential receives your top quartile of selling time. If it does not, territory design is quietly capping revenue growth and no amount of coaching, tooling, or compensation redesign will fix a coverage model built on outdated assumptions about where the business actually is.

Sales leaders spend enormous energy debating quota-setting formulas and compensation plans while treating the underlying question of where sellers should spend their time as largely settled. It rarely is. Two accounts of identical size can carry dramatically different value once growth trajectory, margin and switching cost are accounted for, yet most coverage models still allocate time by headcount ratios and geographic convenience rather than by scored potential. Getting territory and account prioritization right does not require adding sellers. It requires a disciplined method for deciding which of the accounts already on the books deserve the next hour of selling time and which do not.

The economics of misallocated coverage

Territory design is frequently treated as an administrative exercise handled once a year during planning season, then left alone until complaints about workload imbalance force another redraw. That framing understates what is actually at stake. Optimizing how territories are structured, independent of any change to headcount, quota, or incentive design, can lift revenue meaningfully on its own, according to research from Andris Zoltners, Prabhakant Sinha and Sally Lorimer of ZS Associates published in Harvard Business Review1. The mechanism is straightforward: when territories are unbalanced, sellers in

overloaded regions ration attention across too many accounts, while sellers in underloaded regions have idle capacity chasing marginal opportunity. Both failure modes destroy value and neither shows up clearly in a standard pipeline report because the aggregate numbers can still look acceptable even as individual rep performance diverges sharply.

The deeper problem is that territory boundaries and account priority are usually decided independently, when they should be sequential decisions. A geographically balanced territory full of low-potential accounts is still a poorly designed territory. Bain and Company's 2025 survey of roughly 1,300 senior commercial executives found that companies achieving above-median revenue and margin growth were markedly more likely to have built a differentiated capability for deploying selling resources against the highest-value market opportunities, rather than spreading coverage evenly2. The gap between winners and laggards in that research was not effort or headcount. It was precision about where effort went.

Segmentation before scoring

Prioritization only works once accounts are grouped in ways that make comparison meaningful. Segmenting by industry, size, or buying complexity establishes the boundaries within which a scoring model can be applied consistently, because comparing a global manufacturing account against a regional distributor on the same scale produces noise rather than insight. Segments should be built around shared buying behavior and shared value drivers, not solely around org-chart convenience such as sales region.

Within each segment, scoring typically weighs three categories of input.

1. Realized value, meaning current revenue, margin and payment reliability
2. Latent potential, meaning total addressable spend the account could plausibly direct toward the business
3. Cost to serve, meaning the selling and delivery effort required to capture that potential profitably

Skipping segmentation and jumping straight to a single company-wide score is the most common design flaw in prioritization models, because it forces enterprise accounts and transactional accounts to compete on the same scale despite operating under entirely different economics.

Building an evidence-based ideal customer profile

An ideal customer profile [ICP] gives the scoring model its criteria. Rather than guessing at which firmographic traits matter, the more reliable approach starts from a company's best existing outcomes and works backward. Frank Cespedes and coauthors, writing in MIT Sloan Management Review, describe the failure mode this replaces: ventures that pursue growth without adequately defining who their ideal customers are, which caps growth at a certain scale because selling, general and administrative costs accelerate faster than revenue once the sales organization is chasing an undifferentiated mix of accounts3. Building the profile from evidence means examining accounts with strong retention, healthy margin and reference value, then identifying which observable traits preceded the deal closing rather than which traits were noticed afterward.

A profile built this way also has to account for how a company reaches its ideal customers, not just who those customers are. Forrester's research on go-to-market strategy argues that an ICP alone sharpens targeting but does not multiply reach and that pairing it with an understanding of the partners and channels that already have credibility with target accounts accelerates both pipeline velocity and win rates. As Leslie Vitrano Hubright, vice president of global IT channel and partner ecosystem at Schneider Electric's data center and networks division, put it in that research:

> Ideal Partner Profiles are the strategic counterpart to Ideal Customer Profiles, creating a foundation for a diverse and resilient route to market4

The practical implication for territory design is that scoring should not stop at "does this account look like our best customers." It should also ask whether the account is reachable through a channel the seller can actually use, because a high-scoring account with no viable path to engagement absorbs planning time without producing pipeline.

Whitespace as the multiplier

Once accounts are scored, the next question is what specifically to sell them and that is where whitespace analysis earns its place in the process. Whitespace measures the gap

between what an account currently buys and what it could plausibly adopt across the full portfolio, based on peer accounts with similar profiles. Consider an illustrative scenario: a mid-market industrial supplier discovers, after mapping current purchases against a peer benchmark, that its most tenured accounts are buying only two of six available product lines despite spending patterns that closely resemble accounts buying five or six. That gap represents pipeline that requires no new logo acquisition at all, only a redirected conversation with an account that already trusts the supplier.

McKinsey's research on data-driven commercial growth found that outperforming B2B sales organizations consistently use expected value as the common currency for prioritizing opportunities, often starting with a small number of use cases such as entering these kinds of white spaces to build early momentum before scaling the approach5. Companies using this kind of data-driven growth engine reported EBITDA increases in the range of 15 to 25 percent, generated through a combination of sales growth and margin improvement rather than headcount expansion. That figure is worth sitting with, because it reframes prioritization from a productivity tactic into a margin lever, which tends to get more durable executive sponsorship than a productivity argument alone.

Redesigning territories around potential, not geography

Once accounts are scored and whitespace is mapped, territory boundaries should follow the resulting potential distribution rather than the reverse. Geography still matters for travel-dependent field roles, but it should be a constraint applied after potential is distributed evenly, not the primary organizing principle. A territory built purely on postal codes will inevitably clump high-potential accounts in some regions and starve others, regardless of how carefully headcount is matched to account counts.

Redesign should also account for cost to serve explicitly, because two territories with identical scored potential can require very different amounts of selling effort to capture it. An account with a long, committee-driven procurement process consumes far more seller time per dollar of realized revenue than a transactional account of the same size. Ignoring this distinction produces territories that look balanced on a spreadsheet and feel wildly unbalanced to the sellers working them, which is precisely the mismatch that drives attrition among high performers who feel their best accounts are diluted by low-value assignments.

Operationalizing prioritization in the field

A scoring model only creates value if sellers actually change behavior because of it and that requires embedding the output somewhere reps encounter daily rather than in a slide deck reviewed once during planning. Salesforce's State of Sales research notes that sales teams are increasingly deploying AI agents across the sales cycle specifically to help reps move faster from planning to execution, a trend consistent with the broader shift toward surfacing prioritization guidance inside the tools sellers already use rather than in a separate governance artifact6. Whether or not a company adopts AI tooling, the underlying principle holds:

> prioritization guidance that lives outside the seller's daily workflow gets ignored within a quarter, regardless of how sound the underlying model is

Frontline managers play the decisive role here. A scoring model handed to reps without manager reinforcement in pipeline reviews tends to get treated as a compliance exercise rather than a genuine reallocation of effort. Managers who ask, in every one-on-one, why a rep's calendar allocation matches or diverges from the account priority tier create the behavioral pressure that a dashboard alone cannot.

Governance and re-scoring cadence

Prioritization degrades quickly without a fixed rescoring cadence, because the inputs driving a score, growth trajectory, competitive activity, procurement timing, change constantly even when the account itself looks stable on the surface. A model scored once a year and never revisited effectively freezes a company's view of value at a single point in time, which becomes less accurate with every passing month.

Governance also needs a clear owner with authority to reallocate, not just to recommend. Revenue operations or sales operations typically hold this role, working with finance on margin inputs and with sales leadership on the qualitative judgment that pure data cannot fully capture, such as an account's political complexity or a champion's likely departure. Splitting ownership between functions without a tie-breaker produces models that quietly stop being updated because no single group feels accountable for the outcome. The organizations that sustain the discipline treat the score itself as a living asset, reviewed on a

set schedule, defended with evidence and adjusted as the underlying business changes.

- 1Why Sales Teams Should Reexamine Territory Design
- 2The B2B Growth Divide: What Sets Winners Apart
- 3How To Identify The Best Customers For Your Business
- 4Multiply B2B Growth: Why Your Ideal Customer Profile Alone Isn't Enough
- 5Insights To Impact: Creating And Sustaining Data-Driven Commercial Growth
- 6State Of Sales Report

Summary

Territory and account prioritization is not a one-time exercise completed during annual planning and then forgotten. It is a scoring discipline that separates growth from stagnation because it forces an honest answer to a question most organizations avoid: does selling capacity actually sit where the value is? Ideal customer profile criteria, whitespace analysis and potential-based territory design work together only when reassessed on a fixed cadence and governed by someone with authority to reallocate. Companies that treat this as infrastructure rather than a planning ritual consistently redirect underused capacity toward higher-value accounts without adding a single new hire. The account list a rep works from should look meaningfully different each cycle. If it does not, the scoring model has stopped doing its job and the same accounts are absorbing effort regardless of what they are actually worth.