Territory Design Playbook

A First-Principles Guide to Scoring, Carving, Assigning & Governing Sales Territories

Julia Ormond · RevOps Toolkit

Territory design is one of the highest-leverage decisions in revenue operations. Done well, it creates fair competitive conditions for reps, maximises market coverage, and makes pipeline forecasts more predictable. Done poorly, it drives rep attrition, leaves market potential untapped, and bakes structural unfairness into the compensation system before a single deal is worked.

This playbook walks through the full process from first principles: who needs to be in the room, what decisions to make before any data is touched, how to score and carve accounts, how to validate balance, how to assign territories to reps, and how to govern changes. Each section includes rationale for why each variable matters and benchmarks from high-performing SaaS revenue teams.

01 Why Territory Design Matters

Every sales org has territories - explicit or implicit. The question is whether they were designed or just accumulated. Accumulated territories reflect history: who was hired when, which accounts a rep called on years ago, and which manager had the most political capital during planning. Designed territories reflect market potential, rep capacity, and strategic coverage priorities.

What good territory design produces:

- Predictable pipeline - each territory is sized to a rep's realistic capacity, so quota coverage math holds at the team level
- Rep motivation - reps believe their territory gives them a fair shot at quota attainment - the single largest driver of voluntary attrition in sales
- Market coverage - ICP accounts receive coverage proportional to their value, not based on geography of convenience
- Coaching clarity - managers can evaluate performance relative to territory quality, not just raw numbers
- Forecast reliability - territory-level pipeline benchmarks are stable enough to use in bottoms-up forecasting

What bad design costs:

- Structural sandbagging - reps in over-allocated territories hold back pipeline to manage quota creep; reps in thin territories disengage
- Manager intuition as policy - without a scoring model, territory decisions default to relationship capital and volume complaints
- Rep attrition at planning - territory changes at annual planning are the second most common trigger for AE voluntary resignation (after quota increases)
- Coverage gaps - high-value segments get under-covered because no formal ICP density map existed at carving time

02 Stakeholders - Who Needs to Be in the Room

Territory design is a cross-functional process. The decisions it produces affect comp outcomes, headcount plans, and

CRM architecture. Getting the wrong people in the room - or leaving key stakeholders out - produces designs that fail to stick.

Role	Accountable For	Common Concern	When to Involve
CRO / VP Sales	Final approval on territory model and rep assignments	"My top reps are going to push back on this"	Design kickoff and final sign-off
RevOps	Data modeling, scoring logic, balance validation, documentation	Data completeness and model accuracy	Full process - owns the work
Finance / FP&A	Quota alignment, TAM validation, headcount-to-revenue coverage	"Does this cover the board plan?"	Inputs phase and validation
Frontline Managers	Rep-to-territory matching, coaching ramp plans	"Rep X has a relationship in that account"	Assignment step and review
Account Executives	Accepting territory, transition planning for existing pipeline	"I worked that account for two years"	Assignment communication only
Sales Operations	CRM hygiene, account ownership records, data quality	Incomplete or duplicated account records	Pull and clean phase

GOVERNANCE NOTE: Territory design decisions should be made by RevOps with CRO approval - not by managers lobbying individually. When individual managers negotiate directly, the result is always a territory map that reflects seniority and political capital rather than market potential. Define the process before the data is shared.

03 Planning Inputs - What to Lock Before Touching Data

Territory design requires six inputs to be locked before the scoring model is run. Starting without them produces a map that will need to be redone when the missing variable surfaces. RevOps owns getting these confirmed - not assumed.

Input	What It Is	Why It Matters	Who Provides It
ICP Definition	Firmographic criteria: employee count, revenue range, industry verticals, geography, tech stack	Everything downstream is only as good as the ICP filter. A loose ICP inflates the addressable universe; a tight one underestimates it.	Product Marketing + Sales Leadership
Clean Account Universe	All CRM accounts filtered through ICP criteria, deduplicated, and enriched with current firmographic data	Stale or duplicate records corrupt scoring. Every account needs: employee count, revenue estimate, industry, and relationship flag.	RevOps + Sales Ops
Headcount Plan	Quota-carrying rep count by segment and region, including planned hires and expected attrition	Territory count must match rep count. Building territories first and hiring into them later creates blank territories - the most common coverage gap.	Finance + HR
Quota Targets	Individual rep quota by segment, used to validate that territory potential supports the assigned quota	A territory with $500K in ARR potential cannot support a $700K quota at 70% attainment. Misaligned territories set reps up to fail before they start.	Finance + RevOps
Historical CRM Data	Prior 2 years of pipeline creation, stage conversion, and closed-won ARR by account and territory	Historical performance is the strongest predictor of territory quality. Scoring without it relies entirely on market estimates, which are always rough.	RevOps (Salesforce reports)
Segmentation Model	How accounts divide by size: Enterprise, Mid-Market, SMB - with clear thresholds that match the CRM	Without agreed segment boundaries, "Enterprise" means different things to different managers. Territory design requires a single source of truth.	Sales Leadership + RevOps

04 Design Decisions - Variables to Define Before Carving

Before scoring a single account, five structural decisions must be made. These define the shape of the territory model itself. Changing them after carving begins requires starting over.

4.1 Territory Model Type

The most fundamental decision: what is a territory? The four options are not mutually exclusive - most mature teams use a hybrid. But the primary model must be chosen first, because it determines how accounts are grouped and how balance is measured.

Model Type	How It Works	Best For	Watch Out For
Geographic	Accounts grouped by physical location - city, metro, state, region	Field sales, transactional deals, teams where in-person matters	Population density creates structural imbalance - coastal metros dwarf rural regions in ICP density
Vertical / Industry	Accounts grouped by industry: Financial Services, Healthcare, Manufacturing, etc.	Product-market specialisation; teams with deep vertical expertise	Vertical size varies widely; a "Healthcare" territory can be 10x larger than a "Non-profit" territory
Account Size (Segment)	Accounts grouped by employee count or ARR: Enterprise (500+), Mid-Market (100-499), SMB (<100)	Teams with specialised AE roles by deal complexity	Requires crisp segmentation definitions and CRM hygiene to prevent misclassification
Named Account	A fixed list of strategic accounts assigned to specific reps, often overlaid on a base model	Enterprise teams, accounts with executive relationships or active multi-year deals	Named account lists get stale; require quarterly review to remove closed or churned accounts

4.2 Segment Thresholds

Segment	Typical Employee Range	Typical ARR Range	Notes
Enterprise	1,000+	$5M+ ARR	Often split: Enterprise (1K-5K) and Strategic (5K+)
Mid-Market	100-999	$500K-$5M	Widest variance in deal complexity; watch ACV range within segment
SMB	10-99	<$500K ARR	High velocity, low ACV; often handled by a pooled territory model
Micro / VSB	1-9	<$50K ARR	Typically self-serve or PLG; rarely worth a named-account territory

4.3 Overlay Coverage Model

Decide whether any roles carry overlay coverage - typically Solutions Engineers, Customer Success, or Specialists - before carving AE territories. Overlays affect coverage ratios and whether territory balance includes their capacity. If overlays exist, define the coverage ratio: how many AEs per SE, how many AEs per CSM.

4.4 Named Account Policy

Define up front: (1) what qualifies an account for named status, (2) how named accounts are sourced and approved, (3) the review cadence, and (4) what happens to a named account if the assigned rep leaves. Named account lists without governance become political - managers advocate for their reps to hold logos regardless of active pipeline.

4.5 Lock Window

Define the standard territory lock window at design time, not at assignment time. The lock window is the period after assignment during which no changes are permitted except for the defined rebalancing triggers. Industry standard is 30-60 days post-assignment. Without a published lock window, every rep with a complaint becomes a territory re-negotiation.

05 The Territory Scoring Model - 5-Factor, 100-Point Scale

Before accounts can be grouped into territories, every account needs a score. The score reflects its value to the business - not just its size, but its fit, reachability, and competitive context. Scoring replaces gut feel with a repeatable model that can be documented, defended, and improved over time.

Factor	Wt	What It Measures	Why It Matters	Data Sources
Addressable Accounts	30 %	Total ICP-fit companies within the territory boundary (geo, vertical, or size band)	The largest driver of territory potential. A territory with 400 ICP accounts has structurally more upside than one with 80, regardless of current pipeline.	ZoomInfo / Clearbit / Apollo with ICP filter applied
ARR Potential	25 %	Estimated maximum new ARR from territory accounts based on average ACV and historical close rates	Raw account count does not equal revenue potential. ACV-weighted potential normalises the difference between SMB-heavy and Enterprise-heavy territories.	CRM historical data; market sizing by segment
Pipeline History	20 %	Prior-year pipeline creation volume, stage conversion rates, and closed-won ARR within the territory	Historical pipeline is the strongest predictor of future pipeline. Territories with strong history signal warm market conditions and proven rep motion.	Salesforce opportunity reports (prior 2 years)
Competitive Presence	15 %	Density of known competitor installs, contracts approaching renewal, and active evaluations	High competitive presence signals displacement opportunity. It also signals harder selling conditions, so it should increase score only when displacement is realistic.	G2, TrustRadius, Bombora intent data, partner intel, field notes
Geographic / Access Density	10 %	Travel efficiency (avg. drive time between accounts), event density, partner coverage	For field teams, inaccessible territories waste selling time. For virtual teams, replace this with timezone overlap or conference alignment.	Google Maps data; event calendars; partner maps

SCORING TIERS: After scoring all accounts, group into three tiers: High (70-100 points), Medium (40-69), and Low (1-39). A healthy territory contains 15-25% High, 40-50% Medium, and 30-40% Low accounts. Territories composed entirely of High-tier accounts are over-concentrated; territories with >50% Low-tier are under-resourced.

CALIBRATION: Run the scoring model, then review the distribution across all accounts. If >40% land in High tier, your weights are too loose - recalibrate. If <10% land in High, the ICP filter may be too restrictive. The goal is a distribution that meaningfully differentiates account quality.

06 The 6-Step Design Process

Once planning inputs are locked, design decisions are made, and the scoring model is built, the carving process follows six sequential steps. Each step has a clear output that feeds the next. Do not skip steps - each one catches errors that would be expensive to fix later.

1

Pull and Clean the Account Universe:

Export all accounts from CRM. Apply the ICP filter: employee count, revenue range, industry, and tech stack. Deduplicate. Flag accounts with existing pipeline or closed-won history. The output is a clean, enriched account list - your addressable universe.

- Use a single data source for firmographics - mixing sources creates inconsistent employee counts and duplicate records
- Mark accounts with existing relationships or active pipeline as "relationship accounts" - these receive special handling in the assignment step
- Flag accounts acquired in the last 12 months as "new" - they score lower due to no pipeline history; this is expected
- Expected output: spreadsheet with one row per account, ICP flag, firmographic fields, and relationship history tag

2

Score Every Account:

Apply the 5-factor scoring rubric to each account. Normalise factor scores to a 0-100 scale before applying weights. Calculate the composite score. Bucket into High / Medium / Low tiers. Validate the distribution.

- Normalise within each factor first: the highest-scoring account gets 100; all others are scaled proportionally
- Build the scoring in a spreadsheet first, validate it, then back-populate scores into a Salesforce custom field
- Review the High-tier accounts manually before proceeding - if any look wrong, investigate the data quality issue before carving
- Expected output: account list with composite score, tier bucket, and factor breakdown

3

Cluster Accounts into Territories:

Group scored accounts into territory clusters. For geographic models: cluster by metro/region. For vertical models: by industry. For segment models: by size tier. The goal is clusters of roughly equal composite score - not equal account count.

- Target: +-15% variance in composite territory score across all territories
- Never build territories to match rep headcount exactly - build the right territories first, then match headcount
- Document every clustering decision and its rationale. Any override of the scoring model must be recorded with an expiry date
- Expected output: territory map with account-to-territory assignments and composite score per territory

4

Validate Balance:

Calculate composite score for each territory. Compute variance relative to the mean. Any territory outside +-15% variance must be rebalanced before proceeding to assignment. Document all moves.

- Variance % = (Territory Score - Mean Score) / Mean Score x 100
- Rebalance by moving individual accounts - not by splitting territories. Moving accounts is reversible; splitting territories is not
- Get written sign-off from the CRO on the final territory map before assignment
- Expected output: validated territory map with documented variance per territory

5

Assign Territories to Reps:

Match each territory to a rep based on experience level, existing relationships, language and culture fit, and strategic priority. Enterprise and high-score territories go to fully-ramped reps. New reps receive developing territories with coaching support.

- Create an assignment matrix: columns = territory score tiers, rows = rep tenure bands (0-6m ramp, 6-18m developing, 18m+ fully ramped)
- Honour existing pipeline: Stage 3+ deals in accounts being reassigned follow the rep through close
- Lock the assignment for 30 days from the effective date. Communicate the lock window in the assignment documentation
- Expected output: signed territory assignment letters with effective date and lock window stated explicitly

6

Document and Lock:

Issue territory assignment documentation to each rep and manager. Define the lock window. Publish the rebalancing triggers. Update CRM with territory ownership fields. Archive the scoring model and clustering rationale.

- Documentation should include: effective date, lock window end date, account list with scores, composite territory score, variance vs. mean, and rebalancing trigger conditions
- Update Salesforce account ownership fields on the effective date - not before, to avoid rep confusion during transition
- Archive the scoring model with its data snapshot so future rebalancing can be compared against the baseline
- Schedule a 30-day review meeting at assignment time - confirm pipeline transition is clean and surface data errors before they become disputes

07 Balance Benchmarks - What Good Looks Like

These benchmarks reflect high-performing SaaS revenue teams at Series B through Series D scale. Use them as targets, not hard rules - every market has structural constraints that prevent perfect balance.

Metric	Healthy Range	Warning Zone	Red Flag
Territory score variance	+-10% vs. mean	10-15% vs. mean	>15% - rebalance required
Accounts per rep (Enterprise)	40-80 named accounts	80-120 accounts	>150 - rep cannot work the list; coverage is theoretical
Accounts per rep (Mid-Market)	80-150 accounts	150-250 accounts	>300 - rep is forced to cherry-pick; low-tier accounts go uncovered
Accounts per rep (SMB)	200-500 accounts	500-800 accounts	>1,000 - territory is a patch, not a territory
High-tier accounts per rep	15-25% of territory	<10% or >35%	<5% High-tier: underpowered; >50%: over-concentrated
Pipeline coverage ratio	3x-4x quota per territory	2x-3x	<2x - territory cannot support quota at expected attainment
% reps within +-15% of mean	>90% of territories	75-90%	<75% - design has structural fairness problems; expect rep complaints

COVERAGE RATIO: Pipeline coverage ratio = territory ARR potential / individual quota. If a rep's quota is $600K and their territory contains $1.8M in ICP ARR potential, their coverage ratio is 3x. A 3x ratio at 70% expected attainment leaves a 10% cushion for deals that slip or shrink. Anything below 2.5x puts quota attainment at structural risk regardless of rep performance.

08 Rep Assignment Criteria

Territory assignment is a separate decision from territory design. The territory map is built on market data; assignments are made on rep data. Conflating the two - carving territories to fit current reps rather than the market - produces territories optimised for internal politics rather than revenue coverage.

Criterion	How to Apply	Weight in Decision
Rep tenure and ramp status	Fully-ramped reps (18m+) receive high-score or Enterprise territories. Reps in months 7-18 receive mid-score, developing territories. New reps (0-6m) receive structured, lower-complexity territories with coaching.	Highest - non-negotiable in the first pass
Existing pipeline and relationships	If a rep has Stage 3+ deals in accounts being reassigned, those deals follow the rep through close. Relationship accounts (multi-year history, exec access) should follow the rep unless the relationship has lapsed.	High - honour active pipeline; respect long relationships
Language and culture fit	For international or multicultural territories, match rep language skills and cultural context where possible. Most important for expansion into non-English-speaking regions.	Medium - applies when multiple reps are otherwise equivalent
Manager preference and team strategy	Managers may advocate for specific assignments based on coaching capacity or team strategy. This input is valid but should be the last variable applied - not the first.	Low - advisory input, not a veto

DISPUTE PROCESS:

Define a formal dispute process before assignments are communicated. Reps have a 10-business-day window to submit a written objection through their manager to RevOps. Objections must be data-based (a specific account relationship or pipeline history), not general preference. RevOps reviews and decides within 5 business days. This process exists once; it is not a recurring negotiation.

09 Governance and Lock Windows

Governance is the part of territory design that most teams skip - and pay for later. Without published governance rules, every rep departure becomes a political negotiation, every new hire triggers an informal territory grab, and the territory map degrades faster than it was built.

Standard Lock Window Policy

- Territory assignments are locked for 30 days from the effective assignment date
- During the lock window, no account moves are permitted except for acquisitions of target accounts
- After the lock window, account moves require a written change request submitted to RevOps with documented rationale
- Territory design is reviewed annually at the start of each fiscal year planning cycle, and quarterly for balance-only checks

Change Request Process

Any account move outside of a formal rebalancing cycle requires a change request:

- Rep or manager submits request to RevOps with: account name, current territory, proposed territory, and reason (relationship history, acquisition, geographic reassignment)
- RevOps reviews the impact on both territories' balance scores within 5 business days
- If the move keeps both territories within +-15% variance, it is approved automatically
- If the move pushes either territory outside +-15%, a compensating move is required or the request is escalated to CRO
- All approved moves are logged with effective date and rationale in the territory archive

Documentation Requirements

- Signed territory assignment letter (rep + manager signature, effective date, lock window end date)
- Account list with scores exported from scoring model at assignment date
- Composite territory score and variance vs. mean
- Named rebalancing triggers applicable to this territory
- Archived scoring model snapshot (with data vintage date)

10 Rebalancing Triggers - When to Re-Open Territories

The annual planning cycle is the primary time to redesign territories. Between cycles, five triggers should prompt a formal rebalancing review.

Trigger	When It Fires	Required Action	Timing
Score variance >+-15%	At any quarterly review, if a territory's composite score has drifted more than 15% above or below the mean	Rebalance accounts between adjacent territories. Document every move. Notify affected reps 30 days before changes take effect.	Quarterly review
Rep departure	Any voluntary or involuntary rep exit from a territory	Freeze the territory immediately. Assign interim coverage within 5 business days. Redistribute at the next planning cycle or within 60 days. Do not split the territory informally.	Within 5 business days
New rep hire	Any new quota-carrying rep joining the team	Carve a new territory from the highest-scored existing territory in the relevant segment. Never create a blank territory - every territory must contain scored accounts on day one.	Before rep start date

Trigger	When It Fires	Required Action	Timing
Strategic account acquisition	An ICP account in one territory is acquired by or merges with an account in a different territory	Assign the merged entity based on which rep holds the larger existing relationship or the higher-stage pipeline. Re-score the affected territory within 30 days.	Within 2 weeks of close
Market expansion	Company launches into a new vertical, geographic market, or product segment creating new ICP-fit accounts not covered by existing territories	Run a targeted scoring and carving exercise for the new segment. Do not stretch existing territories to absorb new coverage.	Before first rep is hired into new segment

11 Benchmarks by Segment

These benchmarks reflect high-performing B2B SaaS teams at growth stage (Series B-D). Calibrate against your own historical data as it accumulates.

Metric	Enterprise (1,000+)	Mid-Market (100-999)	SMB (<100)
Named accounts per rep	40-80	80-150	200-500 (pooled or named)
Annual quota range	$600K-$1.5M	$300K-$700K	$100K-$300K
Average ACV	$80K-$300K+	$25K-$80K	$5K-$25K
Target attainment rate	65-70% of reps at quota	68-73%	70-78%
Uplift factor	125-135%	118-128%	112-120%
Pipeline coverage ratio	3x-4x quota	3x-4x quota	4x-5x quota (faster cycle)
Avg sales cycle	6-12 months	3-6 months	1-4 weeks
Ramp to full quota	6-9 months	4-6 months	2-4 months
High-tier accounts per territory	15-25% of named list	20-30%	N/A - score-based routing
Rebalance frequency	Annual + trigger-based	Annual + trigger-based	Quarterly (high velocity)

12 Common Mistakes and How to Avoid Them

Building territories around existing reps

The most common mistake. When territory shape is determined by "what does Sarah cover?" rather than "what does the market look like?", the design reflects headcount history, not market structure. Every rep departure then becomes a design crisis. Build the market map first; assign reps second.

Skipping the ICP filter

A scoring model applied to the raw account universe produces noise. Without the ICP filter, low-fit accounts inflate territory size and distort balance calculations. Score only ICP-fit accounts. Everything else is a background universe that reps can access but that does not count toward territory value.

Optimising for account count instead of score

A territory with 200 low-scoring accounts is not more valuable than one with 80 high-scoring accounts. Balance on composite score, not account count. Teams that balance on count end up with one rep covering 40 Enterprise accounts and another covering 180 SMB accounts and calling it "fair."

No lock window

Without a published lock window, territory design becomes an ongoing negotiation. The first rep to complain loudly enough gets an account moved. Within two quarters, the territory map reflects social dynamics, not market data. Publish the lock window in the assignment documentation on day one.

Annual design without quarterly balance checks

ICP accounts get acquired, grow into new segments, and go out of business. A territory that was balanced in January can be materially imbalanced by July. Run a lightweight balance check at each quarterly review - not a redesign, just a variance calculation.

Confusing territory design with quota setting

Territory design determines the quality and composition of each rep's addressable market. Quota setting determines how much of that market the rep is expected to close. These are related but separate processes. Trying to do both simultaneously leads to circular logic: the territory is sized to support the quota, and the quota is set to fit the territory.

Not archiving the scoring model

When a rep disputes their territory two quarters after assignment, RevOps needs to show their work. If the scoring model was not archived with its data snapshot, the design cannot be defended and becomes subject to re-negotiation. Archive the scoring model, the account list, and the clustering rationale at every assignment cycle.