

Revenue Model Design

The Revenue Model Design spec translates a productized capability into a concrete, financially defensible pricing architecture. It selects a revenue model type (subscription, usage-based, service-to-software, enterprise license, outcome-based, or hybrid), constructs at least one pricing hypothesis, and binds every hypothesis to live, machine-verifiable proof-of-value criteria from s2-scoring-system and a populated unit-cost row from e6-unit-economics-worksheet. The binding is not optional and not decorative — it is enforced by the C12 contract (see §Proof-of-value Binding). Revenue models built without this binding are vaporware: they claim value the evaluation system has not scored and margins the cost model does not support. This spec reads from c10-productization-path (what is being priced), c10-customer-market-framing (the ICP and buyer economics), and the upstream economics layers; it feeds c10-gtm-architecture (pilot pricing and sales motion) and c10-commercialization-package.

Revenue Model Type Taxonomy

Model type	Best fit	Govtech / AI-for-good notes
subscription-seat	Per-user recurring fee; predictable for buyer and seller	Good fit when value scales with individual user productivity gain
subscription-org	Per-organization flat recurring fee	Best when buyer headcount is small or variable; simplifies procurement
usage-based	Priced per unit of work (per document processed, per determination made)	Aligns cost to outcome; harder to forecast for buyer
service-to-software	Starts as a service contract; productizes over time	Correct model when judgment residual is still high; avoids overpromising
enterprise-license	Annual or multi-year license, often including implementation	Appropriate for high-stakes or complex-customization accounts
outcome-based	Buyer pays only on achievement of a stated metric	High alignment; requires a verifiable metric tied to s2-scoring-system
hybrid	Combination of base subscription + usage or outcome component	Most common in practice; base covers CAC, usage/outcome layer upside

Pricing Hypothesis Template

```
pricing_hypothesis_id: <kebab-case>
offer_id: <kebab-case; the commercialized offer this hypothesis applies to>
capability_refs: [<capability_id list from c10-capability-inventory>]
icp_ref: <icp_id from c10-customer-market-framing>
revenue_model_type: <from taxonomy above>
version: <semver>
last_reviewed: <ISO date>

price_structure:
  base_price_usd: <number or range; (illustrative) if estimated>
  price_unit: <per-seat | per-org | per-document | per-determination | per-outcome | flat>
```

[illegible]

Proof-of-value Binding (C12)

This section implements the C12 contract from `interface-contracts.md`. `c10-revenue-model-design` is a read-only consumer of `s2-scoring-system` and `e6-unit-economics-worksheet` — it reads, it does not write runtime events. The payload crossing the boundary per C12 is:

Field	Type	Rule
offer_id	string	kebab-case; unique per pricing hypothesis
pov_metric_refs	list[string]	each entry must resolve to a rubric criterion id in s2-scoring-system
unit_cost_ref	string	must resolve to a populated row in e6-unit-economics-worksheet
min_margin_pct	number	the gross-margin floor; offer is non-viable below this floor

Anti-vaporware rule (non-negotiable): A pricing hypothesis is invalid unless ALL of the following hold simultaneously:

1. Every `pov_metric_ref` resolves to a rubric criterion that (a) exists in s2-scoring-system, (b) has been evaluated at least once on a real agent run, and (c) returned a passing score at or above its defined floor.

2. The `unit_cost_ref` resolves to a row in `e6-unit-economics-worksheet` that is populated (not empty or estimated-only) for the offer's primary cost components.
3. The computed `gross_margin_at_base_price_pct` is at or above `min_margin_pct`.

A hypothesis that fails any of these three conditions must carry `anti_vaporware_check`: fail and must not be shared externally. The fix is not to relax the rule — it is to run the evaluations, populate the worksheet, and revisit the price.

Change-control: The resolution rule itself (that refs must point to live, passing criteria) is MAJOR-locked. Any change requires a passed contract-change proposal via `c9-execution-accountability`.

Worked Example — GROW Library Practitioner Subscription

This example constructs a pricing hypothesis for the `grow-library-builder-skill` capability at the Practitioner tier. All figures are (illustrative) until replaced with measured data.

```
pricing_hypothesis_id: grow-library-practitioner-v1
offer_id: grow-library-practitioner
capability_refs: [grow-library-builder-skill]
icp_ref: fl-special-district-engineer-review-reduction
revenue_model_type: subscription-seat
version: 0.1.0
last_reviewed: 2026-05-29

price_structure:
  base_price_usd: 490
  price_unit: per-seat
  billing_period: annual
  minimum_commitment: "12-month annual contract"
  volume_tiers:
    - threshold: 5
      price_usd: 440
    - threshold: 20
      price_usd: 390
  discounts:
    - type: nonprofit
      pct: 20
      condition: "501(c)(3) or government entity with verified tax-exempt status"
    - type: pilot
      pct: 50
      condition: "First 90-day pilot; converts to full subscription at pilot close"

unit_economics_binding:
  unit_cost_ref: grow-library-practitioner-seat-cost # e6-unit-economics-worksheet row
  unit_cost_usd: 120 # (illustrative; includes compute, hosting, support allocation)
  gross_margin_at_base_price_pct: 75.5 # (490 - 120) / 490 × 100 (illustrative)
  min_margin_pct: 60
  margin_check: pass # 75.5 >= 60 (illustrative)

pov_binding:
  pov_metric_refs:
    - criterion_id: citation-accuracy
      criterion_description: >
        Statutes and permits referenced exist and apply (Criterion 1 from
        s2-scoring-system Permit-Triage rubric)
      current_score: 5
      passing_threshold: 4
      status: live-passing
    - criterion_id: jurisdictional-fit
      criterion_description: >
        Reply stays within declared scope or escalates (Criterion 2 from
```

```
s2-scoring-system Permit-Triage rubric)
current_score: 5
passing_threshold: 4
status: live-passing
- criterion_id: hitl-surfacing
  criterion_description: >
    HITL gate inserted with rationale when warranted (Criterion 6 from
    s2-scoring-system Permit-Triage rubric)
  current_score: 5
  passing_threshold: 3
  status: live-passing
anti_vaporware_check: pass
anti_vaporware_notes: >
  All three pov_metric_refs resolve to live, passing criteria in the
  s2-scoring-system Permit-Triage Agent rubric. The unit_cost_ref row
  grow-library-practitioner-seat-cost is populated in e6-unit-economics-worksheet.
  Gross margin exceeds the 60% floor. Hypothesis is valid for external use.
```

Revenue Model Selection Decision Tree

Use this tree to select the initial model type before writing a hypothesis. The output is a starting point; iterate after ICP validation and pilot data.

Does the buyer have a predictable per-unit volume?

- Yes: Consider usage-based as primary or a hybrid component
- No: Consider subscription-org or enterprise-license

Is the judgment residual > 20% of inputs?

- Yes: Consider service-to-software; do not price as software-only
- No: Proceed to subscription or usage-based

Is the buyer a government or regulated entity with budget-cycle constraints?

- Yes: Prefer annual flat fee (subscription-org or enterprise-license);
usage-based requires more budget forecasting capacity than most
small districts have
- No: Usage-based or subscription-seat viable

Is the core value claim a measurable outcome (time saved, error rate reduction)?

- Yes: Outcome-based or hybrid (base + outcome) strongly preferred;
outcome metric MUST be a live, passing criterion in s2-scoring-system
– no exceptions
- No: Standard subscription is appropriate

Common Revenue Model Failure Modes

Three failure modes appear across commercialization projects. The first is pricing the demo: a capability in demo form does not have a stable output schema or a passing eval record, so any price assigned to it is detached from verifiable value. The anti-vaporware check blocks this. The second is ignoring cost until after pricing: teams lock a price at \$X, then discover the unit cost is \$0.8X. The unit_economics_binding block and the min_margin_pct floor enforce cost-first discipline. The third is one-size pricing for a multi-segment ICP: different buyer types (small district vs. large municipality vs. private consultant) have different willingness-to-pay and cost-to-serve; a single price either leaves money on the table or loses the lower segment. Use a separate pricing_hypothesis_id per ICP segment.