

Capability Inventory

The Capability Inventory is the first artifact produced in the commercialization track. It enumerates every technical capability in the portfolio, forces an honest separation between demo-tier and product-tier potential, and flags which capabilities carry a genuine moat / defensibility advantage versus which are commodity executions. No commercialization hypothesis is sound until this register is populated: skipping it produces the classic vaporware failure mode of pricing a capability before understanding what it actually is and what problem it reproducibly solves. Every downstream artifact — the ICP definition in c10-customer-market-framing, the productization path in c10-productization-path, and the full package in c10-commercialization-package — reads capability ids from this register.

Schema

```
- capability_id: <kebab-case unique id>
  name: <short human title>
  description: <two to three sentences; what it does, not what it could be>
  current_form: <demo | service | internal-tool | product>
  problem_solved: <one sentence; the specific pain the capability addresses>
  evidence_of_demand: <cite observed demand signal: ticket count, user quote, pilot request,
market study>
  repeatability: <low | medium | high>
  repeatability_notes: <what makes it low or high; dependencies, judgment calls, data
requirements>
  moat_potential: <none | weak | moderate | strong>
  moat_source: <proprietary-workflow | data-advantage | distribution | integrations |
compliance-moat | operational-expertise | none>
  commercialization_tier: <shelf | explore | accelerate | ship>
  tier_rationale: <one to two sentences>
  blocking_gaps:
    - <what must be resolved before this capability can advance one tier>
  cross_refs:
    - <ids from e4-workflow-artifacts, t7-govtech-package, or other canonical artifacts
this capability depends on>
```

Tier definitions

Tier	Meaning
shelf	Insufficient demand signal or repeatability; archive for now.
explore	Demand signal exists; repeatability and moat are unproven; investigate further before investing.
accelerate	Demand signal confirmed, repeatability medium or higher, at least one moat source identified; invest in productization.
ship	Standard workflow defined, ICP confirmed, unit economics positive; move to active go-to-market.

Worked Example — GROW / GoodSam Capability Stack

The following register entry illustrates the stormwater pre-screener agent — a real capability built for South Walton County special district operations — alongside the GROW library itself treated as a commercializable capability bundle. All figures marked (illustrative) are placeholders; replace with measured data before advancing a tier.

```
- capability_id: stormwater-permit-prescreener
name: Stormwater Permit Pre-Screener Agent
description: >
  Ingests a permit application (PDF or structured form), classifies it by
  stormwater impact category, surfaces required FDEP rule citations, and
  routes to clerk queue or engineer review. Built on the South Walton County
  special district portal; operates under FDEP delegated authority.
current_form: internal-tool
problem_solved: >
  District staff spend 2-4 hours per application triaging permit impact class
  and assembling the applicable rule citations before any substantive review
  can begin.
evidence_of_demand: >
  District commissioner request (2026-Q1); ~120 permit applications/year
  currently handled manually; engineer review backlog averaging 3.2 business
  days (illustrative).
repeatability: high
repeatability_notes: >
  Decision logic is fully deterministic for the top-80% impact classes; only
  the edge 20% (wetlands adjacency, grandfathered parcels) require engineer
  judgment. The workflow is encoded in e4-workflow-artifacts instance
  stormwater-permit-triage-v1.
moat_potential: strong
moat_source: compliance-moat
commercialization_tier: accelerate
tier_rationale: >
  Demand confirmed within one district; FDEP delegated-authority jurisdictions
  span 35+ FL counties [VERIFY], each with the same manual triage problem.
  Compliance moat derives from proprietary rule mapping and the audit-trail
  schema required by FS Chapter 119 [VERIFY].
blocking_gaps:
  - Multi-jurisdiction rule layer not yet built (see t7-multi-municipality-architecture)
  - Unit economics worksheet row not populated (e6-unit-economics-worksheet)
  - Pilot SLA and pricing hypothesis not tested
cross_refs:
  - e4-workflow-artifacts
  - t7-govtech-package
  - t7-compliance-workflow
  - s2-scoring-system

- capability_id: grow-library-builder-skill
name: GROW Builder Skill Library
description: >
  A structured library of ten inter-wired courses (Systems, Execution, Trust,
  Coordination) that equips practitioners to design, evaluate, and operate
  production-grade AI agent stacks. Each course is a set of fillable,
  contract-conformant Markdown artifacts with worked examples.
current_form: internal-tool
problem_solved: >
  Teams building AI agents repeatedly rediscover the same reliability,
  provenance, and commercialization failure modes. GROW encodes institutional
  knowledge as reusable, auditable artifacts that compress the design-to-deploy
  cycle.
evidence_of_demand: >
  10-course curriculum designed to spec; initial adoption by GoodSam.ai /
  Good Combinator internal teams; inbound queries from govtech and AI-for-good
  practitioners (illustrative).
repeatability: high
repeatability_notes: >
  Library artifacts are template-driven; any practitioner can instantiate a
  new agent build without inventing process steps. Inter-artifact contracts
  (C1-C12) enforce consistent wiring across teams.
moat_potential: moderate
moat_source: proprietary-workflow
```

```
commercialization_tier: accelerate
tier_rationale: >
  Proprietary workflow moat from the C1-C12 contract surface and the
  glossary-locked term set. Demand signal from govtech and AI-for-good
  verticals. Needs an ICP sharpening pass and a pilot cohort before
  advancing to ship.
blocking_gaps:
  - ICP definition not yet written (c10-customer-market-framing)
  - Pricing hypothesis not tested at external accounts
  - Proof-of-value metrics not yet bound to s2-scoring-system rubric criteria
cross_refs:
  - s2-scoring-system
  - e6-unit-economics-worksheet
  - c10-customer-market-framing

- capability_id: govtech-compliance-agent
name: GovTech Compliance Workflow Agent
description: >
  Wraps the t7-govtech-package artifacts into an agent that pre-screens
  municipal or special-district compliance decisions, cites governing statutes,
  and escalates expert-review-required determinations before any applicant-
  facing output is emitted.
current_form: demo
problem_solved: >
  Small municipalities and special districts lack the staff to maintain
  current awareness of rule changes and apply them consistently across
  applications.
evidence_of_demand: >
  FASD pilot interest expressed at 2026 annual conference (illustrative);
  direct district commissioner demand from South Walton operations.
repeatability: medium
repeatability_notes: >
  Core compliance-determination logic is repeatable; jurisdiction-specific
  rule layers require periodic maintenance via t7-rule-change-tracker. The
  not-legal-advice boundary requires a reliable expert-review gate before
  each jurisdiction can be marked production-ready.
moat_potential: moderate
moat_source: compliance-moat
commercialization_tier: explore
tier_rationale: >
  Demand signal preliminary; multi-jurisdiction architecture not yet built.
  Advance to accelerate after one live pilot with a non-South-Walton
  jurisdiction and a positive proof-of-value scorecard from s2-scoring-system.
blocking_gaps:
  - t7-multi-municipality-architecture not yet instantiated for a second jurisdiction
  - Expert-review SLA and liability boundary not contractualized
  - e6-unit-economics-worksheet row unpopulated
cross_refs:
  - t7-govtech-package
  - t7-compliance-workflow
  - t7-risk-expert-review-plan
  - s2-scoring-system
```

Usage Notes

Populate this register before writing any pricing or go-to-market artifact. The `capability_id` is the join key used by `c10-customer-market-framing` (ICP mapping), `c10-productization-path` (workflow standardization targets), and `c10-revenue-model-design` (offer construction). Treat `current_form` honestly: a capability in demo form priced as a product is a contractual and reputational risk. Advance tiers only when blocking gaps are resolved and evidence is updated — not on aspiration. Review this register quarterly or whenever a new capability reaches a proof-of-value milestone in `s2-scoring-system`.