OBED INDUSTRIES · PRESENTATION

Skills For
Enterprise Execution

Reusable AI Workflow Packages for Enterprise —
converting institutional knowledge into repeatable AI execution.

OBED INDUSTRIES · OBEDINDUSTRIES.AI

Reusable Workflow Packages

An **AI Skill** is a self-contained workflow package that encodes domain expertise into a form that any AI agent can execute — consistently and repeatedly. Think of a skill as institutional knowledge made executable.

CORE CONCEPT

- **Reusable:** Write once, run on any supported AI platform
- **Portable:** Works with Anthropic (Claude) and OpenAI (GPT) — no lock-in
- **Progressive:** From simple prompts to full multi-agent workflows
- **Versioned:** Skills evolve like software — tracked, reviewed, released

SKILL ANATOMY

REQUIRED

`SKILL.md`

Defines the workflow, context, and behavior

🤖 **`agents/`**
Sub-agent definitions

📚 **`references/`**
Domain knowledge

⚙️ **`scripts/`**
Automation helpers

🗂️ **`assets/`**
Templates & examples

💡 **Progressive Disclosure**

Start with a one-page SKILL.md. Add agents and references as the workflow matures. Skills grow with your needs — not the other way around.

Standard Work for Knowledge Tasks

Domain knowledge in most organizations is **fragmented and inaccessible to AI**. SOPs are buried in SharePoint. Experts hold know-how in their heads. Tribal knowledge disappears with turnover. Skills change that.

THE PROBLEM

Knowledge is Scattered

SOPs, spreadsheets, expert intuition, and tribal know-how all live in different places — and none of it is AI-ready.

THE GAP

Generic AI ≠ Domain AI

Off-the-shelf AI tools lack company context. Every engineer reinvents the wheel. Output quality varies wildly by user.

THE SOLUTION

Skills Fill the Gap

Skills encode company knowledge into AI-executable form. Once built, anyone on the team gets expert-level output — every time.

"Standard Work for Knowledge Tasks"

Manufacturing perfected standard work for physical tasks. Skills bring the same discipline to knowledge work — consistent process, measurable quality, continuous improvement.

Full Product Lifecycle Coverage

Skills are applicable across the full product lifecycle — anywhere domain knowledge drives decisions and outputs.

PHASE 1	PHASE 2	PHASE 3	PHASE 4	PHASE 5
Design & Research	**Development**	**Test & Validation**	**Manufacturing**	**Field & Service**
Requirements review	Code review	Test readiness review	Release checklists	Issue triage
Design review	Documentation gen	Anomaly triage	Build instruction gen	RCA support
Risk analysis	Architecture review	Report generation	Nonconformance triage	Customer response
Trade study support	Test case creation	Pass/fail analysis	First article support	Warranty analysis

EXAMPLE SKILL PORTFOLIO

- 📋 **Requirements Review**
- 🔍 **Design Review**
- ✅ **Test Readiness**
- 🏭 **Manufacturing Release**
- 🚨 **Field Issue Triage**
- **+ Your workflow here →**

Measurable Outcomes

Skills don't just improve individual output quality — they create measurable, team-wide performance gains across every metric that matters.

WHAT IMPROVES

- **First-pass output quality:** Domain context means fewer wrong turns from the start
- **Execution speed:** No need to re-explain context every time; the skill carries it
- **Consistency:** Same process, same quality — regardless of who runs it
- **Onboarding speed:** New team members execute at expert level on day one
- **SME leverage:** Experts encode knowledge once instead of answering the same questions repeatedly

KEY METRICS TO TRACK

CYCLE TIME

Time from task start to reviewed, accepted output

FIRST-PASS QUALITY

% of outputs accepted without major revision

REWORK REDUCTION

Hours saved from fewer revision cycles

SME TIME SAVED

Expert hours freed from repeated coaching

Measure Before and After

Establish baseline metrics before your first skill deployment. Track changes over 30–90 days. The data builds the case for broader rollout — and shows where skills need refinement.

Versioned Operational Assets

Skills are **operational assets** — not experiments. Treat them with the same discipline as software releases, controlled procedures, or calibrated instruments.

👤

Ownership

Every skill has a named owner — the subject matter expert responsible for content accuracy and approval of changes. No owner = no skill.

🔄

Review Cycles

Scheduled periodic reviews (quarterly or with each major process change). Skills go stale without maintenance — plan for it from day one.

📝

Change History

Version-controlled in Git (or equivalent). Every change is tracked: what changed, who approved it, and why. Full audit trail.

🔐

Role-Based Access

Not every skill is for every user. Provision by role — engineering, manufacturing, field service. Sensitive skills may require additional controls.

📋

Validation Gates

New skills require sign-off before production use. Validation includes output testing against known-good cases and SME review of a sample run.

📦

Release Management

Skills follow a release process: Draft → Review → Approved → Active → Deprecated. No ad-hoc edits to active skills without a change request.

⚙️

Governance Scales With Risk

A design review skill for conceptual work needs lighter governance than a manufacturing release skill tied to production hardware. Match the rigor to the stakes.

Know the Risks. Build the Controls.

Every powerful tool carries risk. Skills are no exception. The good news: each risk has a well-defined mitigation that can be built into your governance model from the start.

RISK	SEVERITY	DESCRIPTION	CONTROL / MITIGATION
Outdated Content	HIGH	Skill reflects old process; AI produces confidently wrong output	Version control + mandatory review cycles + owner accountability
Prompt Injection	HIGH	Malicious input manipulates skill behavior via crafted content	Input sandboxing, output validation, restricted execution scope
Overgeneralization	MED	One skill trying to cover too many contexts leads to poor fit	Modular variants — one skill per well-defined use case
Overreliance	HIGH	Users accept AI output without review; errors propagate downstream	Human approval gates for high-stakes outputs; clear output labeling
Quality Drift	MED	Skill performance degrades as model or context changes over time	Automated evals/benchmarks against known-good test cases

Crawl · Walk · Run

Don't build a skill library. **Build one skill first.** Prove value in a narrow workflow, then expand. The crawl-walk-run model applies perfectly here.

CRAWL

Single Skill

Days 1–30

- Pick one high-pain workflow
- Interview the SME
- Draft SKILL.md
- Run pilot with 2–3 users
- Measure baseline vs. output

WALK

Skill Set

Days 31–60

- Refine based on pilot feedback
- Add 2–3 adjacent skills
- Establish governance structure
- Expand to full team
- Formalize review cadence

RUN

Skill Library

Day 61+

- Cross-functional skill portfolio
- Automated eval pipeline
- Skills as onboarding assets
- External benchmark sharing
- Continuous improvement loop

Recommended 60-Day Pilot

Target one engineering-to-manufacturing handoff workflow — requirements review, design review, or test readiness. These are high-frequency, high-stakes tasks with measurable quality signals and a clear SME to interview. Perfect pilot candidates.

Not Vendor Lock-In. A Portable Investment.

Skills built on the **AgentSkills.io open standard** run on both Anthropic (Claude) and OpenAI (GPT-4o and beyond). Your workflow investment is yours — not tied to any one vendor's roadmap or pricing.

Anthropic / Claude

Native skill execution via Claude's Projects and agent APIs. Ideal for complex reasoning and long-context workflows.

OpenAI / GPT

Full skill compatibility with GPT-4o and Assistants API. Switch models without rewriting your skills.

AgentSkills.io

Open standard for skill authoring. Community-driven spec — interoperable, transparent, and free to implement.

WHY THIS MATTERS

- AI model landscape is evolving fast — portability is a hedge
- Best model for the task may change quarterly
- Enterprise procurement often spans multiple AI vendors
- Open standard = community validation and shared improvements

WHAT STAYS CONSTANT

- SKILL.md format and structure
- Governance and versioning process
- SME knowledge encoded in the skill
- Performance metrics and eval benchmarks

NEXT STEPS

Start with one workflow. Build from there.

Obed Industries helps teams design, build, and govern AI Skills for engineering and manufacturing workflows.

obedindustries.ai