EXECUTIVE INTELLIGENCE SERIES

PUBLICATION EDITION AI WORKFLOW SYSTEMS

A Business Guide to AI Prompts, Skills, Scripts, Hooks, Connectors, Plugins, MCP Servers, Harnesses, and Marketplaces

A practical executive reference for turning repeated AI work into reliable business workflows.

PREPARED BY

Time to Revenue

FORMAT

A4 PDF Guide

01 TIME TO REVENUE

Use this guide as an operating reference.

It is educational business guidance, not legal, financial, security, procurement, or regulatory advice.

AI systems, connectors, plugins, and agent workflows should be reviewed against your own data policies, vendor contracts, risk tolerance, and operational controls before live deployment.

Product capabilities and platform rules change. Validate permissions, retention, logs, security posture, and approval flows before relying on any tool in a live business workflow.

02 TIME TO REVENUE

AI value appears when repeated work becomes reliable work.

The point is not to make every team more technical. It is to give leaders a clear vocabulary for deciding when a prompt is enough, when a workflow needs structure, and when governance becomes non-negotiable.

Use the simplest reliable component for the job. That principle protects momentum and governance at the same time.

Time to Revenue

03 TIME TO REVENUE

In this guide

A curated publication edition for business owners and operators who want reusable AI workflows without unnecessary complexity.

04 TIME TO REVENUE

Use the simplest reliable component for the job.

AI becomes commercially useful when repeated work is made reusable, controlled, and governed.

01
Capability is not the outcome

A model can draft and reason, but business value comes from the structure around the model.

02
Structure should scale with risk

Prompts are enough for one-off drafting. Workflows that affect customers, money, records, or compliance need controls.

03
Separate responsibilities

Instructions, procedures, tool access, exact checks, guardrails, packaging, governance, and distribution are different jobs.

04
Avoid premature complexity

Do not build a connector, plugin, or agent harness when a saved prompt or skill will solve the problem.

05 TIME TO REVENUE

For leaders turning AI experiments into operating assets.

This guide is written for non-technical business users who need repeatability, safety, and clarity.

01
Business owners and founders

Decide where AI should stay informal and where it needs process.

02
Operations leaders

Reduce manual work while keeping consistency and accountability.

03
Consultants and advisors

Package AI-enabled services and assessments in a way clients can trust.

04
Sales, marketing, HR, finance, legal, and support teams

Standardize recurring work without pretending every workflow needs an agent.

05
Analysts and project managers

Turn repeated reports, briefs, and summaries into reliable routines.

06
Teams evaluating agents and connectors

Ask sharper questions about permissions, logs, approvals, data access, and ownership.

06 TIME TO REVENUE

From prompt to registry

This is not a ladder every task must climb. It is a map for matching structure to repetition, value, and risk.

01 Prompt	02 Saved Prompt	03 Skill
04 Script	05 Hook	06 Connector / MCP Server
07 Plugin	08 Harness	09 Marketplace / Registry

07 TIME TO REVENUE

Good AI workflows separate responsibilities.

The mistake is putting everything into the prompt. Better systems assign each responsibility to the right layer.

01
Instruction
Tells AI what to do.

02
Procedure
Tells AI how your team does the work.

03
Tool access
Lets AI use approved systems.

04
Exact logic
Handles checks AI should not guess.

05
Guardrails
Prevent risky actions or missing steps.

06
Packaging
Makes the workflow shareable.

07
Governance
Controls permissions, approvals, cost, logs, and memory.

08
Distribution
Helps teams find approved reusable components.

08 TIME TO REVENUE

The ten components at a glance

Use this page to choose a starting point before adding more structure.

01
Prompt
Temporary, low-risk work where a human reviews the result.

02
Saved Prompt
Simple recurring outputs that need consistency without a full workflow.

03
Skill
Repeated work where the team has a known method or checklist.

04
Script
Exact checks, calculations, formatting, validation, and transformations.

05
Hook
Guardrails that must run automatically at a workflow control point.

06
MCP Server
Reusable, standardized access to approved tools, files, databases, and internal systems.

07
App Connector
Fast connection to a trusted SaaS app such as CRM, helpdesk, calendar, or storage.

08
Plugin
A complete workflow that should be installed and reused by a team.

09
Harness
Risk-bearing workflows that need approvals, logs, permissions, and live-execution controls.

10
Marketplace / Registry
Discovery, reuse, ownership, permissions, and review of approved components.

SECTION 01 THE COMPONENTS

The building blocks of reliable AI work.

Each component has a job. The publication edition uses the same compact structure for every chapter: what it is, when to use it, when not to use it, a business example, and a governance note.

01 TIME TO REVENUE

Prompt

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A one-time instruction given to an AI system for a specific task, question, draft, summary, or exploration.

WHEN TO USE IT

Use a prompt for low-risk, temporary work where a human will review the result and no live system access or exact validation is required.

WHEN NOT TO USE IT

Do not rely on prompts when the same instruction is reused often, a standard method must be followed, exact checks are required, or the task affects customers, money, legal obligations, privacy, or records.

BUSINESS EXAMPLE

A sales lead asks AI to rewrite a proposal paragraph for one specific prospect before a human edits it.

GOVERNANCE NOTE

Keep sensitive data out of unapproved tools. Ask the model to state assumptions and uncertainty. Convert repeated prompts into saved prompts or skills.

02 TIME TO REVENUE

Saved Prompt

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A reusable prompt template for simple recurring work, often with fields the user fills in each time.

WHEN TO USE IT

Use saved prompts for repeated but straightforward outputs: weekly reports, meeting summaries, job descriptions, client recap emails, or status updates.

WHEN NOT TO USE IT

Avoid using a saved prompt as the only structure for multi-step procedures, live data access, approvals, strict checklists, or exact calculations.

BUSINESS EXAMPLE

A project manager uses an approved weekly-status prompt with fields for completed work, blockers, priorities, and decisions needed.

GOVERNANCE NOTE

Store approved prompts centrally, assign owners, include data-handling rules, and review high-use prompts when business processes change.

03 TIME TO REVENUE

Skill

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A reusable operating procedure that tells AI how your team performs a repeated task, including steps, rules, examples, and quality checks.

WHEN TO USE IT

Use a skill when quality depends on following a known method: complaint triage, proposal review, hiring scorecards, client assessment reports, or campaign briefs.

WHEN NOT TO USE IT

Do not create a skill for one-off work, simple templates, live-data problems, exact calculations, or workflows that really need permissions and logs.

BUSINESS EXAMPLE

A support team creates a complaint-triage skill with severity levels, escalation triggers, response standards, and output format.

GOVERNANCE NOTE

Keep skills focused, test them against real examples, define escalation rules, and review them when policies, products, or regulations change.

04 TIME TO REVENUE

Script

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A small deterministic program that checks, transforms, validates, formats, or calculates something exactly.

WHEN TO USE IT

Use scripts for required fields, totals, dates, duplicate records, file names, valid JSON, CSV cleanup, and other checks where the answer should be the same every time.

WHEN NOT TO USE IT

Do not use scripts for interpretation, nuance, tone, ambiguous rules, or work where a person must decide what good means.

BUSINESS EXAMPLE

Before a lead import, a script confirms every row has name, email, source, consent status, and no duplicate email address.

GOVERNANCE NOTE

Document the rule, test valid and invalid examples, log failures for important workflows, and update scripts when business rules change.

05 TIME TO REVENUE

Hook

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

An automatic check or action that runs at a specific point before or after an AI workflow step.

WHEN TO USE IT

Use hooks when a control must happen every time: before sending, before publishing, before deleting, after tool use, or after completion.

WHEN NOT TO USE IT

Avoid hooks when work is occasional and low-risk, the rule is unclear, or the workflow is still changing so quickly that controls create more friction than safety.

BUSINESS EXAMPLE

Before an AI-drafted refund email is sent, a hook checks whether the amount exceeds policy and pauses for manager approval.

GOVERNANCE NOTE

Define trigger, pass/fail behavior, notification, logging, override rules, and owner. Test both allowed and blocked paths.

06 TIME TO REVENUE

MCP Server

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A standardized way for an AI agent to access approved tools, data, resources, or prompts through the Model Context Protocol.

WHEN TO USE IT

Use an MCP server when multiple AI clients need reusable access to internal documents, files, databases, tickets, repositories, browser tools, or custom business systems.

WHEN NOT TO USE IT

Do not build or enable MCP access when static context is enough, a trusted app connector solves the need, or the team cannot manage authentication, permissions, and logs.

BUSINESS EXAMPLE

A sales assistant queries approved service descriptions and case studies through a read-only MCP server before drafting a proposal.

GOVERNANCE NOTE

Start read-only, use least privilege, separate test and production access, log tool calls, and require approval for write actions.

07 TIME TO REVENUE

App Connector

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A prebuilt integration that lets an AI tool or automation platform connect to a specific business application.

WHEN TO USE IT

Use connectors for common SaaS systems such as CRM, helpdesk, calendar, email, cloud storage, accounting, analytics, or task management.

WHEN NOT TO USE IT

Avoid connectors that request broad permissions, expose sensitive data unnecessarily, lack activity logs, or cannot be limited by folder, role, workspace, or project.

BUSINESS EXAMPLE

A support assistant uses a helpdesk connector to read ticket history before drafting a response for human review.

GOVERNANCE NOTE

Prefer narrow read-only access at first, document who approved the connector, review activity, and revoke unused access.

08 TIME TO REVENUE

Plugin

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A packaged workflow or capability that may include instructions, skills, scripts, connectors, hooks, templates, permissions, and documentation.

WHEN TO USE IT

Use a plugin when a complete workflow should be installed and reused by a team without each person reconstructing the setup manually.

WHEN NOT TO USE IT

Do not create or install a plugin when a prompt, saved prompt, or single skill is enough, the workflow is experimental, or the source and permissions are not trusted.

BUSINESS EXAMPLE

A sales proposal plugin combines a proposal-writing skill, CRM connector, approved service library, pricing hook, and formatting template.

GOVERNANCE NOTE

Review source, permissions, data access, update behavior, owner, logs, removal process, and approval rules before broad deployment.

09 TIME TO REVENUE

Harness

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

The control layer around one or more agents that manages permissions, approvals, logs, memory, cost, tool access, and live execution.

WHEN TO USE IT

Use a harness when AI can affect customers, records, money, HR, legal issues, finance, sensitive data, or live business systems.

WHEN NOT TO USE IT

A harness may be unnecessary for casual drafting, low-risk experiments, one-off work, or workflows with no external tools and no live actions.

BUSINESS EXAMPLE

A finance agent can prepare variance commentary, but the harness logs spreadsheet changes and requires review before anything is distributed.

GOVERNANCE NOTE

Define authority boundaries, approval gates, logs, forbidden actions, cost limits, memory rules, rollback paths, monitoring, and a human owner.

10 TIME TO REVENUE

Marketplace / Registry

A compact operating reference for when this component is enough and when the work needs more structure.

WHAT IT IS

A place to find, distribute, approve, and manage reusable AI components such as prompts, skills, plugins, connectors, MCP servers, scripts, and templates.

WHEN TO USE IT

Use a registry when multiple people or teams reuse components and the organization needs approved versions, ownership, permissions, review dates, and documentation.

WHEN NOT TO USE IT

A registry may be unnecessary for one person using a few casual prompts, but becomes valuable once components are shared.

BUSINESS EXAMPLE

A company keeps an internal catalog of approved skills, connectors, review dates, owners, risk levels, and documentation links.

GOVERNANCE NOTE

Separate experimental and approved components, track permissions, review external sources, remove unused items, and do not let convenience override security review.

SECTION 02 PUTTING THE COMPONENTS TOGETHER

Most useful workflows combine more than one component.

The goal is not to assemble every layer. The goal is to add only the structure required by repetition, value, access, and risk.

SCENARIO 1 TIME TO REVENUE

Customer Support Automation

Support repeatedly asks AI to draft replies to customer complaints.

01 Prompt	02 Saved prompt	03 Complaint triage skill
04 Helpdesk connector	05 Required-field script	06 Refund/legal/safety hook
07 Support workflow plugin	08 Approval harness	09 Internal registry

Why it works

AI handles drafting and summarization while scripts, hooks, approvals, and logs protect sensitive customer outcomes.

SCENARIO 2 TIME TO REVENUE

Sales Proposal Generation

Sales staff spend too much time turning discovery notes into proposals.

01 Prompt	02 Saved proposal template	03 Proposal method skill
04 CRM connector	05 MCP access to approved service library	06 Required-section script
07 Pricing approval hook	08 Team plugin	09 Harness and registry

Why it works

The team gains speed and consistency without giving AI permission to invent pricing or send proposals without review.

SCENARIO 3 TIME TO REVENUE

Internal Policy Assistant

Employees ask HR and operations the same policy questions repeatedly.

01 Prompt	02 Policy explanation template	03 Policy-answering skill
04 Document connector	05 MCP policy resources	06 Source-check script
07 Sensitive-topic hook	08 Assistant plugin	09 Role-aware harness

Why it works

Employees get faster answers, but the AI does not become the final authority on sensitive HR or legal matters.

DECISION FLOW TIME TO REVENUE

Choose the lightest component that can still be trusted.

Move down the flow only when the answer is yes.

01
One-time task?
Use a prompt.

02
Repeated instruction?
Use a saved prompt.

03
Repeatable method?
Use a skill.

04
Exact check or transformation?
Use a script.

05
Automatic guardrail?
Use a hook.

06
External app or data needed?
Use a connector or MCP server.

07
Full workflow to share?
Use a plugin.

08
Risk, approvals, logs, or live actions?
Use a harness.

09
Discovery and reuse across teams?
Use a registry.

RULE TIME TO REVENUE

Let the task shape the component.

This simple mapping catches most decisions before they become technical debates.

IF THE TASK IS	START WITH
Occasional	Prompt
Repeated	Saved prompt
Procedural	Skill
Exact	Script
Safety-critical	Hook
Data-dependent	Connector or MCP server
Workflow-shaped	Plugin
Risk-bearing	Harness
Shared across teams	Registry

MARKETING

Campaign email: prompt. Monthly outlines: saved prompt. Brand-consistent briefs: skill.

SALES

Call summary: prompt. Proposal from CRM: skill, connector, hook, and harness.

FINANCE

Clean exports: script. Variance commentary: skill plus script.

COMPLIANCE

Review before publishing: hook plus harness.

OPERATING PRINCIPLES TIME TO REVENUE

Build for usefulness first, then reliability, then scale.

The strongest AI workflows are boring in the right ways: narrow access, clear ownership, tested checks, and obvious review paths.

1. Start simple and add complexity only when the business need is clear.
2. Reuse only work that repeats often, creates value, or carries risk.
3. Separate AI judgment from exact checks; use scripts for deterministic validation.
4. Use hooks when a guardrail must run every time.
5. Use connectors only when live data is truly needed.
6. Give agents the minimum permissions necessary, starting read-only where possible.
7. Require human approval for sensitive actions.
8. Log workflows that affect customers, records, money, compliance, or decisions.
9. Review reusable components regularly and assign owners.
10. Treat external plugins and connectors as supply-chain risk.
11. Document ownership, maintenance, versioning, and retirement.
12. Test failure paths, not only happy paths.

AVOID TIME TO REVENUE

Most AI workflow risk comes from the wrong level of structure.

The common pattern is either too little control for a risky workflow or too much machinery for a simple task.

01
Overbuilding
Turning every prompt into a complex workflow before repetition or risk justifies it.

02
Wrong control
Using AI judgment where a formula, script, or validation rule is required.

03
Excessive access
Giving agents broad permissions because the connector makes it easy.

04
Skipped approvals
Letting AI affect customers, money, HR, legal exposure, or public content without review.

05
Unverified components
Installing plugins or connectors without checking source, permissions, support, and data access.

06
Stale assets
Failing to review saved prompts, skills, hooks, and plugins as policies and products change.

07
Access mistaken for governance
Assuming a connector provides approvals, logs, and accountability by itself.

08
No audit trail
Using AI in important workflows without a record of data, tools, outputs, approvals, and actions.

WORKSHEET TIME TO REVENUE

Before creating a reusable AI workflow, answer these questions.

Use this page as an intake worksheet before building a saved prompt, skill, connector, plugin, or governed agent workflow.

Task and value

- [ ] What task are we improving?
- [ ] Who does it today?
- [ ] How often does it happen?
- [ ] What outcome should improve?
- [ ] How will success be measured?

Method and output

- [ ] Does it require judgment, calculation, or both?
- [ ] Is there a known method?
- [ ] What should the output look like?
- [ ] What examples define good work?

Data and access

- [ ] What data is needed?
- [ ] Can static context work?
- [ ] Is read-only enough?
- [ ] What data must never be used?

Risk and governance

- [ ] What could go wrong?
- [ ] Does a human need to approve it?
- [ ] What gets logged?
- [ ] What actions are forbidden?

Ownership and testing

- [ ] Who owns it?
- [ ] How often is it reviewed?
- [ ] How are old versions retired?
- [ ] Have messy data and failure cases been tested?

FIVE QUESTIONS TIME TO REVENUE

A practical governance model does not need to be complicated.

It needs to make authority, data, approval, logging, and ownership explicit.

01

Allowed actions

Draft only, read only, create records, suggest responses, prepare reports, or stage changes for review.

02

Data boundaries

Approved folders, assigned accounts, no payroll, no medical data, and no personal customer data unless approved.

03

Approval triggers

Refunds, legal language, HR decisions, public content, mass email sends, pricing changes, or deletions.

04

Logs

User, data source, tool calls, output, approval decision, final action, and timestamp.

05

Owner

A named human responsible for review, updates, incidents, and retirement.

ROADMAP TIME TO REVENUE

A small business does not need to jump straight into complex agents.

Move from personal productivity to governed execution only as the workflow proves value and risk increases.

01
Individual productivity
Prompts for drafting, summarizing, brainstorming, and rewriting.

02
Team consistency
Saved prompts for repeated team tasks.

03
Repeatable procedures
Skills for high-value methods.

04
Reliability checks
Scripts for exact validation and formatting.

05
Guardrails
Hooks that run important checks automatically.

06
Data access
Connectors or MCP servers for approved context.

07
Packaged workflows
Plugins that make mature workflows reusable.

08
Governed execution
Harnesses for approvals, logs, permissions, memory, cost, and live actions.

09
Reuse and scale
Registries that help teams find approved components.

EXAMPLE TIME TO REVENUE

From manual lead review to governed lead intake.

This is how a business moves from AI helps me write to AI supports an approved business process.

Manual task

A business owner reviews website leads, checks fit, drafts a reply, and adds follow-up tasks to the CRM.

01 Prompt: draft one reply	02 Saved prompt: reusable lead reply	03 Skill: fit criteria and tone
04 Script: validate lead fields	05 Hook: check before CRM creation	06 Connector: form, CRM, email
07 Plugin: packaged intake workflow	08 Harness: logs, approvals, sending limits	09 Registry: approved lead intake automation

REFERENCE TIME TO REVENUE

Compact glossary

Use these plain-language definitions to align business and technical discussions.

AI agent	An AI system that can pursue a task across multiple steps, often using tools or data.
Approval gate	A point where work pauses for a human or policy decision.
Deterministic	Rule-based and consistent for the same input.
Harness	The control layer around agents.
Human-in-the-loop	A design where a person reviews or approves important steps.
Marketplace	A place to find and install reusable components.
Plugin	A packaged workflow or capability.
Registry	An internal catalog of approved reusable components.
Script	A small program for exact checks or transformations.
Tool use	When AI calls an external tool, system, function, or data source.
App connector	An integration between an AI tool and a specific business app.
Audit log	A record of what happened in a workflow.
Guardrail	A control that prevents or flags unwanted behavior.
Hook	An automatic check or action at a workflow point.
Least privilege	The minimum access needed to complete the task.
MCP server	A standardized server exposing tools, data, or resources to AI clients.
Prompt	A one-time instruction to an AI system.
Saved prompt	A reusable prompt template.
Skill	A reusable procedure for AI-assisted work.
Workflow	A sequence of steps that turns input into business output.

SOURCES TIME TO REVENUE

Clean source list for deeper review.

Platform-specific details change. Review current documentation directly before implementation.

Prompt Engineering

- OpenAI Prompt Engineering Guide
- OpenAI Help Center: Prompt Engineering Best Practices
- Anthropic Prompt Engineering Overview
- Google Gemini Prompting Strategies

Agents, Tool Use, and Guardrails

- OpenAI Agents SDK
- OpenAI Agents SDK Guardrails
- Anthropic Agent Skills

Model Context Protocol and Connectors

- Model Context Protocol Specification
- Microsoft 365 Copilot Connectors

AI Governance, Risk, and Responsible AI

- NIST AI Risk Management Framework
- NIST AI RMF Playbook
- Microsoft Responsible AI

Security, Supply Chain, and Component Trust

- OWASP Top 10 for LLM Applications
- OpenSSF Software Supply Chain
- GitHub Marketplace for Apps

The most advanced component is not always the right component.

A mature AI workflow is built in layers, and each layer should earn its place.

Start with prompts for one-time work. Turn repeated prompts into saved prompts. Turn repeated procedures into skills. Use scripts when the answer must be exact. Use hooks when a check must happen automatically.

Use connectors or MCP servers only when AI needs approved access to tools or data. Use plugins when a full workflow should be packaged and shared. Use a harness when the workflow involves risk, approvals, memory, cost, logs, permissions, or live execution. Use a registry when teams need to discover and reuse approved components.

The goal is not to use the most advanced component. The goal is to make repeated AI work reliable, useful, and safe.

NEXT STEP AI EFFICIENCY ASSESSMENT

Want to identify where AI can reduce operational drag in your business?

Request a Time to Revenue AI Efficiency Assessment.

timetorevenue.com