ENTERPRISE OPERATING MANUAL

Transformative AI Enterprise Adoption

Using Forward Deployed Agents

A field guide for executives, operators, and technical teams implementing internal agent-enabled workflows in real enterprise conditions.

THREE OPERATING PRINCIPLES

1. **Make work legible** — Map reality before automation.
2. **Build with controls** — Design governance before autonomy.
3. **Drive adoption** — Treat buy-in as the core system.

Trace Cohen | ValueAdd VC
valueaddvc.com

HOW TO USE THIS MANUAL

How to use this manual

This manual is built for real enterprise conditions: fragmented ownership, legacy systems, uneven data quality, competing budgets, and users who may resist more visibility into their work. Use it sequentially for the first deployment and as a reference after.

EXECUTIVES
Set mandate, select workflows, fund the program, inspect progress.

OPERATORS
Map actual work, validate future-state, expose edge cases, lead adoption.

TECHNICAL TEAMS
Define agent scope, access, integrations, approval gates, governance.

Do not start with models, tools, or vendors. Start with a painful workflow that can be described, owned, measured, and changed.

The Core Sequence

#	Stage	What happens	Timeline
0	**Mandate + scope**	Define the business problem, sponsor, budget owner, and boundaries.	2–4 wk
1	**Workflow selection**	Score candidates for impact, feasibility, risk, and adoption readiness.	2–3 wk
2	**Reality mapping**	Interview operators, observe work, collect artifacts, map exceptions.	4–8 wk
3	**Systems + data**	Map sources, systems of record, permissions, APIs, and gaps.	4–8 wk
4	**Process redesign**	Create future-state workflow with owners, states, exception handling.	3–6 wk
5	**Agent scope**	Define precise agent tasks, approvals, tools, and prohibited actions.	2–4 wk
6	**Build + pilot**	Build MVP, test, deploy to controlled group, measure adoption.	10–24 wk
7	**Govern + scale**	Operate with logs, reviews, audits, incident process, expansion roadmap.	Ongoing

90-DAY PROOF OF METHOD
Stages 0–3 deliver a Proof of Method milestone at day 90: validated reality map, scored workflow, system feasibility report, and go/no-go decision artifact.

EXECUTIVE SUMMARY

The practical thesis

Most enterprise AI efforts fail because the organization is not ready to be automated. Processes are undocumented, data is inconsistent, ownership is fragmented, incentives are misaligned, and users resist tools that increase visibility without improving their experience.

Process before agent

Agents cannot execute a workflow the company cannot clearly describe.

Adoption before automation

A technically correct solution that users avoid is a failed deployment.

Governance before autonomy

The more powerful the agent, the more explicit the controls must be.

Narrow before scale

One adopted workflow is worth more than ten impressive pilots.

Incentives before mandates

Users need a reason to prefer the new system beyond being told to use it.

Evidence before expansion

Scale only after usage, quality, and operating value are proven.

The right starting question is not "What can AI do?" It is "Which workflow is painful, measurable, owned, and ready to be made more structured?"

COMMON FAILURE PATTERNS

What enterprises get wrong

Leadership announces an AI priority, a team builds a promising demo, stakeholders agree it looks useful, the system enters pilot, and real usage never materializes.

Assumption	Reality	Response
Process is documented	Real process lives in email, spreadsheets, memory.	Reality mapping with operators.
Data is accessible	Critical fields missing, inconsistent, or owned elsewhere.	Data lineage + permission map.
Users want automation	Users fear visibility, workload increase, or job risk.	Adoption + incentive plan.
Engineering can build it	Engineering may not own systems, data, or budget.	Decision rights + budget ownership.
Pilot = rollout ready	Pilot users may be unusually motivated.	Validate default behavior; retire legacy.
Vendor will handle it	Vendors optimize for deployment, not adoption.	Separate vendor from adoption ownership.

AI Maturity Model

Level	State	Description	Readiness
1	**Fragmented**	Work in email, spreadsheets, meetings, memory.	🔴 Not ready
2	**Documented**	Process described but docs don't match reality.	🔴 Not ready
3	**Structured**	Standard intake, owners, states, escalation paths.	🟠 Assist-ready
4	**Agent-assisted**	Agents retrieve, summarize, draft under human review.	🟢 Active
5	**Governed**	Agents with controls, monitoring, auditability.	🟢 Mature

STAKEHOLDERS + OPERATING MODEL

Who to involve and how decisions get made

Stakeholder	Why they matter	Key question
Executive sponsor	Priority, cover, funding, escalation.	What outcome justifies cross-functional cooperation?
Process owner	Owns workflow, adoption, result.	Who is accountable if this process fails today?
Operators	Know how work actually happens.	Show me the last real example.
Engineering / IT	Systems, integrations, constraints.	What is technically possible and maintainable?
Security / risk	Access, permissions, audit controls.	What could the agent do that creates unacceptable risk?
Finance	Funding and resource allocation.	Whose budget pays for build and ongoing usage?
Change mgmt	Turns pilot into operating behavior.	What training and adoption support are required?

Operating model

Role	Accountability	Must own
Executive sponsor	Strategic priority + escalation	Outcome, funding, conflict resolution
Program lead	End-to-end execution	Timeline, dependencies, decision log
Process owner	Workflow success	Target process, user adoption, metrics
Agent engineering	Technical delivery	Architecture, integrations, agent behavior
Security / compliance	Control framework	Access, logging, approvals, incidents
Change lead	Behavior change	Training, communications, legacy retirement

No production deployment proceeds unless the process owner, engineering lead, security lead, and executive sponsor have all approved scope, controls, and adoption.

SECTION 2

The Step-by-Step Implementation System

Each step has an objective, stakeholders, readiness gate, key questions, common trap, and working detail.

STEP 1 Define mandate and scope

Translate executive interest into one measurable operational problem. Define what will be solved, who owns the result, what is out of scope, and what success looks like.

Who to involve Executive sponsor, program lead, process owner, finance.	**Required outputs** Problem definition, scope boundary, sponsor memo, metric target.
READINESS GATE Sponsor can describe the problem without generic AI language.	**COMMON TRAP** A broad AI mandate substituting for a scoped business problem.

Questions to ask: • What business outcome matters enough to fund this? • Which process is in scope vs. out? • Who owns the result if adoption fails? • What risk boundaries are non-negotiable?

Working detail: Write a one-page mandate naming the workflow, business owner, target users, expected improvement, exclusions, risk boundaries, and funding source.

STEP 2 Select the first workflow

Choose a workflow that is painful, measurable, owned, and realistic. Build confidence first — do not start with the hardest political process.

Who to involve Executive sponsor, process owner, operators, engineering, security.	**Required outputs** Process selection scorecard, priority ranking, named first workflow.
READINESS GATE Best combination of impact, ownership, feasibility, and adoption likelihood.	**COMMON TRAP** Choosing the most strategic workflow when it has weak ownership or poor data.

Questions to ask: • Where is work slow, manual, or error-prone? • Can baseline metrics be captured? • Is there one accountable business owner? • Will users see direct benefit?

Working detail: Compare candidates across pain, frequency, measurability, data readiness, risk, ownership, and adoption likelihood.

STEP 3 Extract current reality

Map how work actually happens today through interviews, observation, and artifact review. Documentation is useful but not the source of truth.

Who to involve Operators, managers, process owner, program lead.	**Required outputs** Current-state process map, exception tree, manual work inventory, delay map.

READINESS GATE	COMMON TRAP
Operators review the map and confirm: this is what actually happens.	Accepting the documented workflow without testing it against real examples.

Questions to ask: • Walk me through the last real instance. • What do you track outside official systems? • Where do you copy and paste? • What breaks most often?

Working detail: Capture spreadsheets, email handoffs, side trackers, manual decisions, delays, and points where the system does not reflect reality.

STEP 4 Map systems, data, and permissions

Document every field, system, permission, integration, and data quality issue before the agent role is designed.

Who to involve	Required outputs
Engineering, IT, data owners, security, compliance.	System map, data lineage table, permission matrix, integration feasibility.

READINESS GATE	COMMON TRAP
Team knows what data exists, where it lives, who owns it, and whether it can be safely accessed.	Assuming dashboard data and source-system data mean the same thing.

Questions to ask: • Where does each data field originate? • Which system is the system of record? • Who can read, write, and approve access? • What data cannot be used by an agent?

Working detail: For every field: source system, owner, update frequency, quality level, access method, permission boundary.

STEP 5 Redesign the process

Before adding agent capabilities, simplify. Remove redundant steps, standardize intake, define owners, create states, clarify exception handling.

Who to involve	Required outputs
Process owner, operators, managers, engineering, compliance.	Future-state workflow, state model, owner map, approval model, exception framework.

READINESS GATE	COMMON TRAP
The future-state workflow is clear enough to build, train, measure, and govern.	Automating the old workflow just because it is familiar.

Questions to ask: • Which steps don't create value? • Where can intake be standardized? • What needs approval vs. not? • What must remain human-owned?

Working detail: Remove duplicate entry, unclear handoffs, redundant approvals. Redesign before automation.

STEP 6 Define the agent role

Specify exactly what the agent will and will not do. Start constrained: retrieval, summarization, drafting, classification, or prepared actions requiring approval.

Who to involve Agent engineer, process owner, security, compliance, operators.	**Required outputs** Agent responsibility matrix, allowed/prohibited actions, approval policy.
READINESS GATE Every agent behavior is documented, testable, permissioned, and owned.	**COMMON TRAP** Broad access or vague instructions because the demo performed well.

Questions to ask: • What task is repetitive enough for agent assistance? • What actions are prohibited? • When is human approval mandatory? • What happens when the agent is wrong?

Working detail: Each skill needs: trigger, inputs, allowed tools, prohibited actions, confidence threshold, approval requirement, output format, fallback.

STEP 7 Build the business case

Convert the redesigned workflow into a funded implementation plan with baseline metrics, expected improvement, build cost, operating cost, and adoption investment.

Who to involve Finance, executive sponsor, program lead, process owner.	**Required outputs** Business case, ROI model, funding plan, resource allocation, timeline.
READINESS GATE Funding and ownership are explicit before engineering capacity is committed.	**COMMON TRAP** Innovation funding creating a pilot with no long-term operating owner.

Questions to ask: • What is the current cost of delay, error, and rework? • What improvement is realistic in 90–180 days? • Whose budget pays? • What costs are hidden in adoption and maintenance?

Working detail: Use conservative assumptions. Show what improves, what doesn't, what risks remain.

STEP 8 Design architecture and controls

Design a secure system with identity, permissions, logging, workflow state, agent orchestration, human approvals, monitoring, and rollback.

Who to involve Engineering lead, agent engineer, security, compliance.	**Required outputs** Architecture document, security model, audit plan, monitoring plan, incident process.

READINESS GATE	COMMON TRAP
Engineering, security, compliance, and process owner approve design before build.	Treating auditability and human override as enhancements rather than requirements.

Questions to ask: • What will the agent read/write? • How will permissions be enforced? • Where will logs live? • Can the system be rolled back?

Working detail: The design should be explainable to a skeptical security or compliance reviewer.

STEP 9 Build, test, and pilot

Build the minimum usable workflow and test with real scenarios. The pilot validates adoption, reliability, risk controls, and measurable improvement.

Who to involve	Required outputs
Agent engineer, engineering, QA, operators, process owner, change lead.	MVP, test results, pilot plan, issue log, adoption baseline.

READINESS GATE	COMMON TRAP
Pilot users complete real work and measurable improvement is visible.	Mistaking a successful controlled demo for a successful operational pilot.

Questions to ask: • Can users complete the workflow without workarounds? • Does the agent handle exceptions? • Do logs capture what happened? • What causes users to abandon the process?

Working detail: Choose a small user group with real work. Track daily, hold weekly feedback, categorize issues.

STEP 10 Drive adoption and retire legacy

Train users, support them, watch usage, fix friction, and gradually remove old paths. Adoption = retirement of alternatives.

Who to involve	Required outputs
Process owner, managers, change lead, operators, support.	Adoption plan, training schedule, champion network, legacy retirement plan.

READINESS GATE	COMMON TRAP
The new workflow becomes the default way work is done.	Relying on executive announcements to change daily operator behavior.

Questions to ask: • Why would users prefer this? • What old workarounds remain? • What manager behavior must change? • What threshold triggers rollout?

Working detail: Users adopt when the workflow is easier, clearer, supported, and reinforced. They resist when it feels like oversight.

STEP 11 Govern, measure, and scale

Operate the agent-enabled workflow like production infrastructure. Review performance, incidents, permissions, and user feedback before expanding.

Who to involve Governance board, process owner, engineering, security, finance.	**Required outputs** ROI report, governance dashboard, incident log, permission audit, scale roadmap.
READINESS GATE Expansion approved only after adoption, quality, and controls are proven.	**COMMON TRAP** Scaling because the pilot looked impressive rather than because the operating model works.

Questions to ask: • Are metrics improving vs. baseline? • Where are users overriding the workflow? • What incidents occurred? • Which workflow is ready next?

Working detail: The second workflow should reuse templates and operating rhythms from the first.

SECTION 3

Templates, Artifacts, and Control Tools

Create evidence, not just documentation.

ARTIFACT 1: PROCESS SELECTION SCORECARD

Score each candidate workflow from 1 to 5. Select for combined impact and feasibility.

Dimension	Question	Weight	Score
Operational pain	Is the process visibly slow, manual, or error-prone?	High	___
Measurability	Can baseline and post-launch metrics be captured?	High	___
Ownership clarity	Is there one accountable business owner?	Critical	___
Data readiness	Are required data sources knowable and accessible?	Medium	___
User adoption	Will users see direct benefit from the change?	High	___
Risk level	Could errors create legal, financial, or customer harm?	Negative	___
Time to value	Can value be demonstrated within 90–180 days?	High	___

ARTIFACT 2: REALITY MAPPING WORKSHEET

Field	What to capture
Trigger	What starts the process? Who receives the first signal?
Actors	Who touches the work directly or indirectly?
Systems	Which tools, spreadsheets, inboxes, dashboards are used?
Inputs	What information is needed to begin? Where from?
Decision points	Who decides what happens next? What criteria?
Exceptions	What happens when data is missing, late, wrong?
Delays	Where does work wait? Who is the bottleneck?
Shadow workflow	What happens outside official systems?
Output	How does the process end? What proves completion?

ARTIFACT 3: STAKEHOLDER INTERVIEW GUIDE

Audience	Opening question	Follow-up
Executive sponsor	What outcome would make this worth funding?	Resistance? Tradeoffs? What cannot fail?
Process owner	Where does this workflow break today?	Who owns each step? Which metrics?
Operator	Walk me through the last completed instance.	What do you track manually?
Engineering	What systems are involved?	APIs? Fragile points? What not to automate?
Security / risk	What data or actions create concern?	What requires approval? What to log?
Finance	Who pays for build and ops?	What evidence justifies funding?

ARTIFACT 4: AGENT ROLE DEFINITION

Item	Definition
Skill name	Clear operational name, not a vague AI label.
Business purpose	What workflow outcome does this skill improve?
Trigger	What event starts the skill?
Inputs	Required data, documents, fields, or user prompts.
Allowed tools	Systems, APIs, databases the agent may access.
Allowed actions	Actions the agent may take without and with approval.
Prohibited actions	Actions the agent may never take.
Human approval	When a person must review, approve, or override.
Confidence/fallback	What happens when uncertain or missing data?
Output format	Exact format the user or system receives.
Owner + version	Business owner, technical owner, version history.

ARTIFACT 5: ROI AND MEASUREMENT

Measure business change, not AI activity.

Category	Baseline	Target	Evidence
Efficiency	Cycle time; manual touches	30–50% cycle time reduction	Workflow logs
Quality	Error rate; rework	Lower rework, fewer errors	QA review
Adoption	Active users; completion	New workflow = default	Usage analytics
Financial	Cost per completion	Clear cost/risk reduction	Finance reports
User experience	Satisfaction; tickets	Users see it as easier	Surveys

ARTIFACT 6: GOVERNANCE CADENCE

Cadence	Participants	Purpose
Daily (pilot)	Program lead, agent engineer, support	Usage, errors, friction, blockers.
Weekly (rollout)	Process owner, managers, change lead	Adoption, feedback, exceptions.
Monthly (production)	Sponsor, engineering, security, finance	Metrics, ROI, incidents.
Quarterly (governance)	Sponsor, risk, compliance, business leads	Autonomy changes, scale roadmap.

ARTIFACT 7: ADOPTION RISK HEATMAP

Risk	Behavior	Mitigation
Fake buy-in	Agree in meetings, don't change behavior.	Written decision rights + manager accountability.
Operator fear	Avoid system due to increased visibility.	Position around reduced burden + clearer ownership.
Legacy persistence	Keep using email/spreadsheets.	Retire old paths gradually; monitor bypass.
Manager non-enforcement	Tolerate parallel processes.	Dashboards, scripts, escalation support.
Data distrust	Don't trust agent output.	Show source data, confidence, approval steps.
Support weakness	Small issues become adoption excuses.	Fast support loop during pilot + rollout.

Adoption = Retirement of Alternatives. The new workflow only becomes default when the old workflow is no longer available.

ARTIFACT 8: 90-DAY PROOF OF METHOD TRACKER

Week	Milestone	Deliverable	Owner
1–2	**Mandate confirmed**	One-page scope memo with sponsor signature	Program lead
3–4	**Workflow selected**	Completed scorecard with rationale	Process owner
5–8	**Reality mapped**	Current-state map validated by operators	Program lead
9–10	**Systems assessed**	Data lineage + integration feasibility	Engineering lead
11–12	**Go / no-go**	Proof of method report with recommendation	Sponsor

IMPLEMENTATION READINESS CHECKLIST

Final checklist

Review before go-live. Each unchecked item needs an owner and resolution date.

- [ ] Executive sponsor is active and can name the business outcome.
- [ ] Business process owner is accountable for adoption and results.
- [ ] First workflow is narrow, measurable, and feasible.
- [ ] Operators have validated the current-state reality map.
- [ ] Systems, data sources, owners, and permissions are documented.
- [ ] Future-state workflow has owners, states, approvals, and exception paths.
- [ ] Agent actions are narrow, explicit, testable, and permissioned.
- [ ] Security and compliance controls are approved before production.
- [ ] Business case includes build cost, operating cost, and adoption investment.
- [ ] Pilot plan includes training, support, metrics, and issue management.
- [ ] Legacy workflow paths identified and scheduled for retirement.
- [ ] 90-day proof of method artifacts complete and reviewed.
- [ ] Governance includes owners, logs, audits, reviews, and incident handling.
- [ ] Scale plan is based on proven adoption, not enthusiasm.

Do not deploy agents into organizational ambiguity. Make the workflow visible, structured, governed, adopted, and measurable. Then make it agent-enabled.

The companies that win will not simply deploy the most agents.

They will build the organizational muscle to make critical processes visible, structured, governed, adopted, and measurable.

1	2	3	4	5	6
Choose one workflow	Map real work	Structure the process	Define agent skills	Pilot and govern	Measure and scale

Trace Cohen | ValueAdd VC

valueaddvc.com