AI Integration Playbook (C2O Edition)

Design AI initiatives with shared ownership

Scan for full guide with examples, checklists, and updates:
contribution2outcome.com/kb/foundations/what-is-c2o

contribution2outcome.com

December 17, 2025

Version 1.0

Owner: Stark Burns

Table of Contents

Executive Overview 3

Use-Case Intake Canvas 3

- Outcome Definition 3
- Lifecycle Assessment 4
- Data & Technical Requirements 4
- Risk Assessment 4
- Stakeholder Map 4

Role Mapping Patterns for AI Initiatives 4

- Standard AI Role Matrix 4
- Key Role Definitions for AI 5
- Distributing Validation Work 5

Bias & Risk Safeguards 5

- Bias Taxonomy 5
- Validation Cadence 5
- Red-Team Reviews 6
- Rollback Protocol 6

Implementation Sprint (4 Weeks) 6

- Week 1: Discover & Define 6
- Week 2: Build & Validate 7
- Week 3: Instrument & Decide 7
- Week 4: Run & Adopt 7

Metrics Dashboard 8

- Leading Indicators 8
- Lagging Indicators 8
- Behavioral Indicators 8

Governance Integration 8

- AI-Specific POP Additions 8
- Gate Evidence Requirements 9

Anti-Patterns to Avoid 9

Worked Example: Customer Service AI Assistant 9

- Outcome 9
- Role Matrix 9
- Validation Distribution 10

Metrics 10
Next Steps 10
See Also 10

A practical guide for launching AI initiatives with shared ownership using the C2O model. Designed to address the expert productivity paradox where AI tools increase individual productivity but create validation bottlenecks.

Executive Overview

The Problem: AI tools can boost individual productivity by 30–50%, but organizations often see only marginal gains because:

- Expert bottlenecks form around validation and review
- Governance overhead consumes the productivity gains
- Adoption stalls due to unclear ownership
- Risk management becomes reactive rather than proactive

The Solution: C2O distributes Drive/Contribute/Enable work across lifecycle phases, reducing the 19% expert productivity paradox by shifting validation to Contributors and embedding governance from day zero.

This playbook covers:

- Use-case intake and outcome definition
- Role mapping for AI initiatives
- Bias and risk safeguards
- 4-week implementation sprint
- Metrics and governance

Source: contribution2outcome.com/kb/foundations/what-is-c2o

Use-Case Intake Canvas

Before starting an AI initiative, answer these questions collaboratively:

Outcome Definition

Table 1: Question / Your Answer

Question	Your Answer
What outcome will change if we succeed?	
How will we measure success?	
What is the target timeline?	
Who benefits from this outcome?	

Lifecycle Assessment

Table 2: Question / Your Answer

Question	Your Answer
Which lifecycle phase is the current bottleneck?	Discover / Decide / Build / Run / Adopt
What manual work will AI augment or replace?	
What human oversight is required?	
What is the expected productivity gain?	

Data & Technical Requirements

Table 3: Question / Your Answer

Source: contribution2outcome.com/kb/foundations/what-is-c2o

Your Answer	
What data sources are required?	
What is the data quality status?	
What model type is being considered?	
What infrastructure exists?	

Risk Assessment

Table 4: Question / Your Answer

Question	Your Answer
What bias risks exist? (representation, measurement, historical)	
What hallucination/accuracy risks exist?	
What compliance requirements apply?	
What rollback protocol is needed?	

Stakeholder Map

Table 5: Question / Your Answer

Question	Your Answer
Who must Enable success from day 0?	
Who has veto power?	
Who are the end users?	
Who champions adoption?	

Role Mapping Patterns for AI Initiatives

Standard AI Role Matrix

Table 6: C2O Contribution Mode Matrix

Phase	Drive	Contribute	Enable	Advise	Inform
Discover	Product / Business Owner	Data Science, Domain Leads	Legal, Security, Finance	Ethics Council	Executive Sponsor
Decide	Business Sponsor	AI Lead, Finance	Enterprise Architecture	Compliance, Risk	PMO
Build	Engineering Lead	Data Science, SMEs, QA	Platform, DevOps, MLOps	Risk & Controls	Support
Run	Operations Lead	SRE, Support, Product	Data Governance, MLOps	CISO, Compliance	Business Teams
Adopt	Change Lead	Champions, Training	HR, Communications	Customer Success	All Contributors

Key Role Definitions for AI

Table 7: Role / AI-Specific Responsibilities

Role	AI-Specific Responsibilities

AI Lead	Model selection, architecture, performance
Data Science	Training, validation, feature engineering
MLOps	Pipeline automation, monitoring, deployment
Ethics Council	Bias review, fairness assessment
Risk & Controls	Compliance, audit, risk mitigation
Domain SME	Validation of outputs, edge case identification

Distributing Validation Work

Anti-pattern: All AI outputs routed to senior experts for review (creates bottleneck)

C2O pattern: Validation distributed to Contributors by confidence level:

Table 8: Confidence Level / Validation Owner / Action

Confidence Level	Validation Owner	Action
High (>0.9)	Automated checks	Auto-approve with audit trail
Medium (0.7–0.9)	Contributor	Review with checklist; escalate exceptions
Low (<0.7)	Domain Expert	Full review required

Bias & Risk Safeguards

Bias Taxonomy

Table 9: Bias Type / Definition / Detection

Bias Type	Definition	Detection	Mitigation
Representation	Training data doesn't reflect target population	Demographic analysis of data	Augment data; stratified sampling
Measurement	Features proxy for protected attributes	Correlation analysis	Remove proxies; fairness constraints
Historical	Training data encodes past discrimination	Outcome analysis by group	Rebalancing; fairness-aware training
Automation	Over-reliance on AI without human check	Usage patterns	Mandatory human review for high-stakes

Validation Cadence

Table 10: Lifecycle Phase Reference

Phase	Validation Frequency	Focus
Build	Daily	Model accuracy, bias metrics
Run (first 2 weeks)	Daily	Drift, accuracy degradation
Run (ongoing)	Weekly	Drift alerts, incident trends
Adopt	Per cohort	User feedback, outcome correlation

Red-Team Reviews

Conduct red-team reviews at each phase gate:

Table 11: Gate / Red-Team Focus / Participants

Gate	Red-Team Focus	Participants
Discover ’Decide	Use-case validity; data availability	Ethics, Legal, Domain SME
Decide ’Build	Architecture risks; bias potential	Security, ML Lead, Risk
Build ’Run	Edge cases; failure modes	QA, Domain SME, SRE
Run ’Adopt	Scaling risks; user impact	Support, Champions, Risk

Rollback Protocol

Document in POP before launch:

```
## AI Rollback Protocol

### Triggers
- [ ] Accuracy drops below [threshold] for [duration]
- [ ] Bias alert exceeds [threshold]
- [ ] User complaints exceed [count/timeframe]
- [ ] Compliance incident

### Actions
1. Automated rollback to previous version
2. Notify: [Driver], [Sponsor], [Risk]
3. Root cause analysis within [timeframe]
4. Post-mortem within [timeframe]

### Decision Authority
- Automated rollback: MLOps
- Rollback decision: Driver
- Re-launch approval: Sponsor
```

Implementation Sprint (4 Weeks)

Week 1: Discover & Define

Table 12: Day / Focus / Activities

Day	Focus	Activities	Outputs
1–2	Outcome definition	Intake workshop; outcome canvas	Outcome statement; success criteria
3	Stakeholder alignment	Role mapping session	Contribution matrix
4	Risk assessment	Risk pre-mortem	Risk register; mitigation plan
5	POP finalization	Review and sign-off	POP v1 published

Week 1 Checklist:

- Outcome defined with measurable criteria
- Drivers named per phase
- Bias taxonomy documented
- Rollback protocol drafted
- POP signed by Sponsor

Week 2: Build & Validate

Table 13: Day / Focus / Activities

Day	Focus	Activities	Outputs
6–7	Prototype development	Model training; initial validation	Prototype v1
8–9	Validation rounds	Bias testing; accuracy checks	Validation report
10	Acceptance review	Test against criteria	Go/no-go recommendation

Week 2 Checklist:

- Prototype meets accuracy threshold
- Bias metrics within tolerance
- Validation checklist completed
- Edge cases documented
- Build ’ Run gate pre-read prepared

Week 3: Instrument & Decide

Table 14: Day / Focus / Activities

Day	Focus	Activities	Outputs
11–12	Instrumentation	Monitoring setup; dashboards	Monitoring live
13	Gate review	Build ’ Run evidence review	Gate decision
14	Enablement prep	Runbooks; support training	Support readiness
15	Pilot launch	Limited rollout	Pilot live

Week 3 Checklist:

- Monitoring dashboards operational
- Alert routing configured
- Runbooks published
- Support team briefed
- Pilot launched to first cohort

Week 4: Run & Adopt

Table 15: Day / Focus / Activities

Day	Focus	Activities	Outputs
16–17	Pilot monitoring	Signal review; incident response	Signal report
18	Expansion decision	Pilot review; expansion gate	Go/hold decision
19	Adoption launch	Full rollout; adoption campaign	Full launch
20	Retrospective	Sprint retro; lessons captured	Improvement plan

Week 4 Checklist:

- Pilot signals within thresholds
- Expansion gate passed
- Full rollout complete
- Adoption campaign launched
- Retrospective conducted

Metrics Dashboard

Leading Indicators

Table 16: Metrics Overview

Metric	Definition	Target	Owner
Decision Latency	Time from AI output to human decision	<4h for high-priority	Driver
Validation Coverage	% of outputs receiving review	>95% for low-confidence	QA
Drift Alerts	Count of model drift threshold breaches	<5/week	MLOps
Bias Metrics	Fairness metrics by protected group	Within tolerance	Data Science

Lagging Indicators

Table 17: Metrics Overview

Metric	Definition	Target	Owner

Value Realization	Productivity gain vs. baseline	>20%	Product
Adoption Rate	% of target users actively using	>80%	Change Lead
Incident Volume	AI-related incidents per week	Trending down	SRE
Accuracy (Production)	Model accuracy on live data	>95%	MLOps

Behavioral Indicators

Table 18: Metrics Overview

Metric	Definition	Target	Owner
CPO	Collective Psychological Ownership score	+0.5 vs baseline	Sponsor
Collaboration Density	Cross-functional touchpoints	Targeted increase	Driver
Escalations Avoided	Decisions resolved without escalation	>85%	Driver
SME Load	Expert hours on validation	-30%	Driver

Governance Integration

AI-Specific POP Additions

Add this section to your standard POP:

```
## AI Governance Rules

### Model Ownership
- Model Owner: [Name]
- Data Owner: [Name]
- Ethics Reviewer: [Name]

### Thresholds
| Trigger | Threshold | Action |
|---------|-----------|--------|
| Accuracy drop | >5% from baseline | Investigate; hold expansion |
| Bias alert | Any protected group | Immediate review |
| Drift | >10% distribution shift | Retrain trigger |

### Review Cadence
- Daily during pilot
- Weekly after full launch
- Quarterly bias audit
```

Gate Evidence Requirements

Table 19: Gate / AI-Specific Evidence

Gate	AI-Specific Evidence
Discover ’Decide	Data quality report; bias risk assessment
Decide ’Build	Architecture review; ethics approval
Build ’Run	Validation report; bias metrics; runbooks
Run ’Adopt	

	Pilot results; support readiness; adoption plan

Anti-Patterns to Avoid

Table 20: Anti-Pattern / Symptom / Fix

Anti-Pattern	Symptom	Fix
Expert bottleneck	All outputs queue for senior review	Distribute validation by confidence level
Governance theater	Reviews happen but don't impact decisions	Tie reviews to gates; require evidence
Pilot purgatory	Never graduate from pilot	Set expansion criteria upfront
Adoption afterthought	Build complete but no one uses it	Embed adoption plan from day 0
Bias blindspot	No systematic bias monitoring	Implement bias taxonomy; automate alerts
Rollback panic	No plan for reverting	Document rollback protocol in POP

Worked Example: Customer Service AI Assistant

Outcome

"Customer service agents resolve tier-1 issues in d5 minutes with e90% accuracy, reducing escalation rate by 30%."

Role Matrix

Table 21: C2O Contribution Mode Matrix

Phase	Drive	Contribute	Enable	Advise
Discover	CS Director	AI Lead, CS Leads	Platform	Legal, Compli-ance
Decide	CS Director	Finance, AI Lead	Architecture	Risk
Build	AI Lead	Engineers, CS SMEs	MLOps	CS Director
Run	CS Operations	SRE, Support	MLOps	AI Lead
Adopt	Change Lead	Champions, Training	HR, Comms	CS Director

Validation Distribution

Table 22: Confidence / Action / Volume

Confidence	Action	Volume
High (>0.9)	Auto-apply with audit	70%
Medium (0.7–0.9)	Agent review + approve	25%
Low (<0.7)	Route to human	5%

Metrics

- Leading: Response accuracy, time to resolution, agent override rate
- Lagging: Customer satisfaction, escalation rate, resolution time
- Behavioral: Agent confidence scores, adoption utilization

Next Steps

1. **Complete intake canvas** with key stakeholders
2. **Map roles** using the standard AI matrix
3. **Define bias taxonomy** and validation distribution
4. **Draft POP** with AI-specific governance rules
5. **Schedule Week 1** activities (outcome workshop, risk pre-mortem)
6. **Set up metrics dashboard** with leading indicators
7. **Trigger assessment** after Week 4 to capture wins and lessons

See Also

- Pre-emptive Ownership Pact (POP)
- Contribution Mapping Canvas
- Evidence Review Checklist
- Metrics Dictionary

Playbook from C2O Knowledge Base

Related Documents

- Pre-emptive Ownership Pact (POP)— C2O Template