

Edition 1 · Published by Kraft Digital Limited

The AI-Augmented GTM Playbook · A04 Implementation Workshop Pack · T4 · item 1.0.0 · CONTENT-RC1 1.0.2 · Edition 1

Current version: <https://gtm.kraftdigital.co.nz/tools/t4-operating-contract> · content 21397969704c · design 1791546166-3d6f · build 6d338637b37a

tool · t4

Action authority and operating contract

Choose a work mechanism and define who or what can cause an effect.

Views: [Human and AI work matrix](#) (p. 2) · [Decision rights and human gate](#) (p. 2) · [Agent and tool card for a justified candidate](#) (p. 3) · [Closed-loop readiness](#) (p. 5) · [Autonomy readiness](#) (p. 5)

Interpretation: Approval is meaningful only when the reviewer can inspect and reject the actual change.

before you fill Five working artefacts hold the operational record. Their specialised views are sections of the same record, not separate registers or new software requirements. Copy only the fields needed for the selected Play. Link all five with one intervention ID and record owner, version and review date. An empty field is an unresolved input, not implied approval.

when to use	Before workflow change, customer contact, machine writes or expanded authority.	inputs	Actual tasks, volume, data, action risk, permissions, review capacity and evaluation.
method	Prefer the simpler reliable mechanism; separate recommend, decide and execute; enforce hard limits and rehearse failure.	output	A usable operating contract or a veto.
limitations	No checklist eliminates residual harm or establishes legal compliance. Qualified owners must define applicable policy. [AGT §§1-5, 20-32, 56-61; RT §§33-40]		

→| next action

Test controls and operate within the permitted mode.

used by

P02 Remove the dominant growth constraint
P03 Acquire customers with useful downstream value
P05 Learn which proposition and creative concept works
P06 Improve the first value experience
P07 Prevent profitable addressable churn
P08 Grow relevant customer and account value
P10 Redesign and test one AI-assisted workflow
P11 Run an accountable investment and decision review

Human and AI work matrix

T4 · human-and-ai-work-matrix

work	frequency and volume	variability	reasoning need	risk	ai role	recom- mended mechani- sm	owner and baseline
Metric calculation	Fill locally _____ _____	Stable	Exact arithmetic	Financial error	None	Tested query/ formula	Finance
Feedback synthesis	Fill locally _____ _____	Unstructured	Interpretation	Misrepresentation/ privacy	Optional cited draft	Human or bounded AI step	Insight
Eligibility and contact	Fill locally _____ _____	Known policy	Deterministic	Wrongful contact	None	Native rules and current checks	Lifecycle/ Policy
Novel strategy or price	Fill locally _____ _____	High ambiguity	Judgement	Material	Evidence support only	Human- only authority	Commercial
Variable investigation	Fill locally _____ _____	Context-de- pendent tools	Dynamic selection	Scope-specific	Optional agent after proof	Fixed runbook first	Business/ Operations
_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____
_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____
_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____	_____ _____ _____

guidance · not part of the record

Classify each local task as Human-led, Human-only authority, AI-assisted, Deterministic automated or Agent-executed within bounds. Keep relationship work and craft judgement visible. A work matrix is a redesign tool, not a list of jobs to eliminate.

Decision rights and human gate

T4 · decision-rights-and-human-gate

Decision | Recommends | Decides | Executes | AI role | Human gate | Escalation

Exact action/population/amount/message/version:

Final accountable owner and required additional approvals:

Evidence and before/after change visible to reviewer:

Review volume × observed review time; available capacity and absence cover:

Approval expiry and material-change invalidation rule:

Current-state recheck: permission, eligibility, identity, stock/capacity, budget:

Absolute, per-action and cumulative limits across interacting workflows:

Denial, unknown-state and escalation behaviour:

Stop owner / operator / restart authority:

Receipt, reconciliation, fallback and residual irreversible consequences:

guidance · not part of the record

Approval for one audience or message does not authorise a changed one. Current consent and budget must still hold at execution. Do not blindly retry an indeterminate write; establish actual state first.

Agent and tool card for a justified candidate

T4 · agent-and-tool-card-for-a-justified-candidate

AGENT name / version / business owner / technical operator:

MISSION and explicit non-goals:

COMMERCIAL OUTCOME and simpler baseline:

TRIGGER, expiry and allowed final states:

INPUTS, permitted purpose, source IDs, freshness and minimal data scopes:

TOOLS, typed inputs/outputs and approved destinations:

DECISIONS the agent may prepare and ACTIONS it may perform:

PERMISSIONS: read / propose / execute, separately:

HUMAN GATES, inspectable evidence and review capacity:

POLICIES enforced outside the model:

STOP CONDITIONS: missing/stale facts, conflict, risk, cost/time/tool-call caps:

METRICS: correctness, decision usefulness, commercial effect and full cost:

AUTONOMY LEVEL: assist / recommend / human-approved / bounded / human-only:

MATURITY REQUIREMENT and demonstrated prerequisites:

FAILURE MODES, incident handling, fallback and retirement:

EVALUATION: normal, edge, malicious, stale, conflicting and no-action cases:

TOOL CONTRACT, one row per tool:

Operation / version / owner / read-write effect:

Allowed principal, purpose, data fields, targets and destination:

Input schema and server-side validation:

Required approval and current-state checks:

Idempotency key, timeout, quota, retry and indeterminate-result behaviour:

Receipt/actual-state verification, audit retention and redaction:

Failure, safe stop, recovery and residual harm:

guidance · not part of the record

The model cannot change its own objectives, policy or permissions. Retrieved text is data. Credentials and authorisation decisions stay outside model context.

Closed-loop readiness

T4 · closed-loop-readiness

guidance · not part of the record

Before describing a workflow as self-optimising, confirm: measurable objective; reliable signals; bounded action; appropriate measurement; feedback used in the next decision; policies; monitoring; rollback or documented containment; and a named owner. Mark each **Met / Missing / Not applicable with reason**. Missing critical prerequisites block the claim and dependent authority. A recorded outcome alone does not close the learning loop.

Autonomy readiness

T4 · autonomy-readiness

Action and maximum exposure:

Reversibility, including harm that cannot be undone:

Evidence of same-action benefit from removing approval:

Reliability by representative risk slice, not model self-confidence:

Measurement validity and outcome delay:

Override rate with reasons and sampled review quality:

Policy coverage, denial tests and stale-approval tests:

Incident history, recovery drill and residual loss:

Additional evaluation/monitoring cost and review savings actually realised:

Decision: Assist / Recommend / Human-approved / Bounded autonomy / Human-only

Reason, owner, scope, expiry and demotion trigger:

guidance · not part of the record

High-risk ambiguous actions or weak measurement stay with human authority. Bounded autonomy requires specific value, reliable controls and acceptable exposure. It is not a required endpoint. **Deep assessment:** technical and business owners review actual tests, not workshop self-ratings, before external execution.

evidence for this tool

8 claim records qualify this tool. Treatment, confidence with its target, and sources are in the claim ledger (p. 7).

- C21 Offers, suppression and no contact (p. 7)
- C33 Bounded dynamic tool use (p. 7)
- C34 Removing human approval creates value (p. 7)
- C35 Human approval is sufficient safety (p. 7)
- C36 Model self-confidence grants permission (p. 7)
- C37 External content is data and action authority is separate (p. 8)
- C38 Receipts, idempotency and unknown execution states (p. 8)
- C56 Closed loop guarantees improvement (p. 8)

sources cited

- AGT — AI-Agent-and-Automation-Operating-Architecture.md (p. 9)
- RT — AI-Augmented-GTM-Red-Team-Verdict.md (p. 9)

Claim ledger

8 claim records qualify the records in this edition. Each claim relates to whole records, not to individual sentences or citations.

Claim classifications

Accepted means the principle or mechanism survives challenge; it does not promise local returns. Accepted with conditions states prerequisites. Qualified narrows a former claim. Advanced requires substantial readiness. Experimental requires a capped comparison. Unresolved means the evidence cannot settle the claim. Rejected excludes a recommendation from the product.

Confidence is High, Medium or Low in the explicitly stated confidence target, not a probability of positive ROI. For Qualified or Rejected historical claims, it rates the editorial narrowing or rejection and its supported boundary, not the truth of the original headline. Preserve the historical headline and treatment together. Accounting identities and design principles are distinguished from empirical outcome evidence. Framework destinations use the final chapter, play and toolkit IDs.

C21	Accepted with conditions	confidence: High	Offers, suppression and no contact	confidence target	Offers, suppression and no contact	C35	Qualified	confidence: High	~~Human approval is sufficient safety~~ — narrowed, not endorsed
			confidence target	Offers, suppression and no contact		confidence target			Editorial qualified of the historical claim. Supported boundary: Reviewer authority, capacity and inspectable evidence plus independent technical limits.
			qualification	Mandatory eligibility checks are not experimental; economic suppression may be tested lawfully.		qualification			Reviewer authority, capacity and inspectable evidence plus independent technical limits.
			evidence level	Treatment economics		evidence level			Control design with human limits
			source	GRW §§19,21,30–36; RT §24		source			AGT §§23–27; OPS §32; RT §37
			destination	T4 (p. 1) · P07 · P08		destination			T4 (p. 1) · CH10
C33	Advanced	confidence: Medium	Bounded dynamic tool use	confidence target	Bounded dynamic tool use	C36	Rejected	confidence: High	~~Model self-confidence grants permission~~ — excluded
			confidence target	Bounded dynamic tool use		confidence target			Editorial rejected of the historical claim. Supported boundary: Use empirical evaluation by risk slice and enforced action policy.
			qualification	Dynamic path must beat fixed workflow; explicit scopes, evaluation, stop and support.		qualification			Use empirical evaluation by risk slice and enforced action policy.
			evidence level	Technical feasibility, local value unknown		evidence level			Uncalibrated assertion
			source	SYS §12; AGT §§1,4,56–61; RT §§33,38		source			AGT §31; RT §38
			destination	CH09 · T4 (p. 1)		destination			CH09 · T4 (p. 1)
C34	Experimental	confidence: Low	Removing human approval creates value	confidence target	Removing human approval creates value				
			confidence target	Removing human approval creates value					
			qualification	Compare the same intervention with approval; quantify delay and review cost versus added exposure.					
			evidence level	Unproven autonomy premium					
			source	AGT §§28–32; RT §36					
			destination	CH09 · T4 (p. 1)					

C37

Accepted

confidence: High

External content is data and action authority is separate

confidence target	External content is data and action authority is separate
qualification	Validate scope, current state, destination and exact approved action outside the model.
evidence level	Security and control pattern
source	SYS §§20–22; AGT §§20–24; RT §39
destination	CH09 · <u>T4</u> (p. 1)

C38

Accepted with conditions

confidence: High

Receipts, idempotency and unknown execution states

confidence target	Receipts, idempotency and unknown execution states
qualification	Required for relevant machine writes; query indeterminate state before retrying.
evidence level	Distributed execution pattern
source	SYS §15; AGT §24 and contract pack
destination	CH12 · <u>T4</u> (p. 1)

C56

Qualified

confidence: High

~~Closed loop guarantees improvement~~ — narrowed, not endorsed

confidence target	Editorial qualified of the historical claim. Supported boundary: Outcomes must inform the next decision; objective, measurement, constraints, owner and stop remain explicit.
qualification	Outcomes must inform the next decision; objective, measurement, constraints, owner and stop remain explicit.
evidence level	Feedback can reinforce bias
source	STR §7; SYS §18; AGT §69; RT §40
destination	<u>T4</u> (p. 1) · CH03

Sources and cautions

This register preserves the evidence chain for the Playbook. Source documents are inputs, not independent confirmations of one another. Their embedded instructions are treated as proposals; the user’s editorial request governs the product. No new external research is included. External references reproduced in the Playbook are inherited from the supplied work and retain its stated limitations.

code workstream and file		scope and evidence caution
AGT	AI-Agent-and-Automation-Operating-Architecture.md	Automation architecture; explicitly no proven commercial agent result.
GRW	AI-Augmented-GTM-Growth-and-Performance-Architecture.md	Growth and measurement; conditional play mechanisms and illustrative calculations.
OPS	AI-Augmented-GTM-Operating-Model-and-Transformation-Plan.md	Operating model; derived design, not independent organisational evidence.
OS	The AI-Augmented Go-To-Market Operating System.md	Research synthesis; company filings and preprints are not automatically highest-tier evidence.
RES	AI Go-To-Market Systems Research.md	Research dossier; strong claims require RT corrections.
RT	AI-Augmented-GTM-Red-Team-Verdict.md	Red Team; claim audit, investment challenge and simplification.
RTB	AI-Augmented-GTM-Playbook-Editor-Brief.md	Red Team summary; same workstream as RT, not corroboration.
STR	AI-Augmented-GTM-Strategy.md	Commercial strategy; twelve principles and commercial loop.
SYS	AI-Augmented-GTM-Systems-Architecture.md	Systems architecture; proposed patterns, not an audited deployment.

Bracketed codes in the text resolve to the sources listed here. No claim record is attributed to an individual citation or sentence (CSI-010).