



HANDVANTAGE

A REFERENCE FOR PROCUREMENT COMMITTEES

The Agentic AI Procurement *Handbook.*

How CISOs, CIOs, procurement teams, and business leaders evaluate AI workspaces and agentic systems before production.

A buyer's field guide, with the claims separated from the questions.

This handbook is for procurement committees evaluating governed AI workspaces and AI workers in regulated or sovereignty-sensitive environments. It helps CISOs, CIOs, CFOs, operations leaders, compliance teams, and service providers ask better questions before a pilot becomes production.

The evaluation method and scorecard are vendor-neutral. Product evidence is kept separate from the method and examined as a worked example in Appendix B. Regulatory statements point to primary sources. Recommendations are evaluation guidance, not law. Deployment-specific dashboards, scores, and test results describe the assessed deployment; they are not certification of the customer organisation.

The handbook is published openly. No email gate. No registration. Share it, attach it to a committee paper, or use the question banks in an RFP. Readers should still obtain legal, regulatory, and assurance advice for their own facts and jurisdiction.

PUBLISHED: 2026-07-25

VERSION: 3.1

AUTHOR: JOSH OLAYEMI, FOUNDER, HANDVANTAGE

CONTACT: HELLO@HANDVANTAGE.COM

CANONICAL URL: [WORKSPACE.HANDVANTAGE.COM](https://workspace.handvantage.com)

LICENCE: FREE TO SHARE, CITE, AND FORWARD

Four routes through the *handbook*.

This is a reference, not a reading test. Begin with the decision you own, then use the scorecard to turn questions into a bounded pilot.

01

CISO / risk leader

Start with Chapters 03, 11, 12, and 13.

Identity → authority → evidence

02

CFO / finance leader

Start with Chapters 04 and 06, then the pilot scorecard.

Economics → controls → record

03

Owner / operations leader

Start with Chapters 02 and 05.

Decision rights → capacity → handoff

04

MSP / advisory partner

Start with Chapters 10 and A.

Assess → deploy → operate → improve

Use the evidence rule throughout: a product statement describes intended or available capability. The configured deployment, current build, and observable record determine what a buyer can rely on.

Table of *contents*.

PART I

From impressive demo to defensible deployment

01	The category decision comes before the vendor decision	7
02	Write the decision-rights map before the pilot	9

PART II

The committee voices

03	For the CISO: make the proof path observable	11
04	For the CFO: measure the operating model, not the demo	13
05	For the owner and COO: add capacity without losing the handoff	15

PART III

Sector evaluation dossiers

06	Financial services: govern the route from request to record	17
07	Healthcare: protect the care boundary	19
08	Public sector: make sovereignty measurable	21
09	Legal and professional services: preserve the matter trail	23
10	MSP, MSSP, and service businesses: productise the control plane	25

Table of *contents*.

PART IV

The twelve-test product review

11	Evaluate deployment architecture and sovereignty	27
12	Evaluate agent identity, authority, and deliverables	29
13	Evaluate AI security, evidence, and failure states	31

PART V

The buyer's field kit and worked example

A	The procurement scorecard and pilot	33
B	Worked example: applying the tests to Vantage Workspace	36
C	Primary sources, claim notes, and limits	40



From impressive demo to defensible deployment.

How to define the category, the operating boundary,
and the evidence standard before a pilot creates a
procurement problem.

THE CATEGORY DECISION COMES FIRST


The same demo can conceal three very different purchases.

A fluent answer and a completed task can look almost identical on screen. The production obligations change sharply when the system moves from producing language to using tools and operating across company records.

CATEGORY TEST

Classify the operating authority, not the interface.

WHAT IS BEING BOUGHT?	Chat assistant Language interface	Agent platform Tool-enabled task system	Governed AI workspace Operating environment
PRIMARY JOB	Produce or transform language.	Complete a bounded task by calling tools.	Coordinate people, AI workers, records, and business applications.
PRODUCTION AUTHORITY	Answer, summarise, draft.	Read, draft, and execute within a defined tool scope.	Operate across identity, data, approvals, and governed workflows.
CONTROL QUESTION	What content may be sent to the model?	Which tools, data, and actions may the agent reach?	Can the committee govern and reconstruct the complete work path?
MINIMUM EVIDENCE	User, prompt route, and response record.	Actor, agent, tool call, policy decision, and outcome.	Identity, boundary, worker, approval, model route, action, and durable audit record.



COMMITTEE QUESTION

If the system can act on company records, is the proposed control model still only describing a chatbot?

THE CATEGORY DECISION COMES FIRST

Decide what the system may do before deciding who supplies it.

A buying committee needs a shared definition of the operating boundary, the permitted authority, and the proof standard before it compares vendors.

THE THREE PROCUREMENT OUTPUTS

Deployment-boundary map Shows where data, memory, models, services, and evidence run.	Decision-rights map Names who may read, draft, approve, and execute.	Reproducible evidence test Reconstructs one buyer-selected transaction end to end.
---	--	--

THE PROCUREMENT SEQUENCE

STEP	DECISION	WHAT THE COMMITTEE MUST ESTABLISH	REQUIRED OUTPUT
01	BOUNDARY	Where will files, prompts, memory, model traffic, and audit records live? Which external services remain in the route?	COMMITTEE OUTPUT Deployment-boundary map
02	AUTHORITY	For each use case, define the highest permitted action: read, draft, approve, or execute. Name the owner, worker, and approval gate.	COMMITTEE OUTPUT Decision-rights map
03	EVIDENCE	Select one transaction and reconstruct actor → worker → tool → policy or approval → model route → outcome.	COMMITTEE OUTPUT Reproducible evidence test

ONLY THEN COMPARE VENDORS	Models / Interface / Price These matter, but they are easier to change than identity, data control, action governance, and the evidence trail.
------------------------------	--

BUYER TAKEAWAY A system that cannot answer boundary, authority, and evidence questions is not ready for production authority, however good its prose may be.

WRITE THE DECISION-RIGHTS MAP BEFORE THE PILOT

“Human in the loop” is not a control until the loop is named.

The committee must identify the human, the exact transition they control, and the consequence of approval. Otherwise a productivity pilot can acquire production authority without an explicit decision.

DECISION-RIGHTS MAP

Separate assistance from authority, workflow by workflow.

REPRESENTATIVE WORKFLOW	READ	DRAFT	APPROVE	EXECUTE
External client response	AI worker reads the approved case context.	AI worker prepares the response.	NAMED HUMAN Account owner accepts content and recipient.	Mail service sends only after recorded approval.
Supplier invoice	AI worker reads the invoice and purchase record.	AI worker proposes coding and flags exceptions.	NAMED HUMAN Finance manager approves the accounting treatment.	ERP posts the approved entry; payment remains treasury-controlled.
Personnel change	Restricted HR role reads the permitted record.	AI worker prepares a proposed change and rationale.	RESERVED JUDGMENT HR director approves or rejects the change.	Authorised HRIS administrator applies the approved record change.

THE CONTROL IS THE TRANSITION

Name the moment where assistance becomes an external action or record change.

Draft → send / Recommendation → approval / Approved change → system update

TEST THE CONTROL, NOT THE POLICY STATEMENT

A control must survive rephrasing, revocation, and a fresh session.

A successful grant proves only that access can be added. A credible control also proves that access can be removed, remains removed, and does not depend on the user choosing the vendor’s preferred words.

THE GRANT-REVOKE VERIFICATION

01	BASELINE	Ask the ungranted user for the protected record. Expect a clear refusal, not an empty result or a general summary. Retain: prompt, denial, actor, worker, and timestamp.
02	GRANT	Add the minimum required grant. Repeat the same request and verify the permitted scope, not simply a success message. Retain: grant event, grant owner, and returned scope.
03	REVOKE	Remove the grant. Repeat the original request without changing the wording. The protected record must be denied again. Retain: revoke event and post-revoke denial.
04	FRESH SESSION	Sign in again and repeat the request. Cached claims or stale sessions must not restore access that the control says was removed. Retain: fresh-login result and underlying identity state.

RESERVED JUDGMENTS

Assistance does not make the AI worker the approver.

Document the decisions that remain human even when the system prepares the evidence.

- Pay rise or employment termination
- Release or movement of funds
- Final legal position or filing
- Clinical diagnosis or treatment
- External publication under a named person

REPHRASING TEST

The gate must follow intent, not a prepared phrase.

<i>“Send it.”</i>	<i>“Take care of this.”</i>	<i>“Make it happen.”</i>	<i>Instruction hidden in a file.</i>	<i>The same request in another language.</i>
-------------------	-----------------------------	--------------------------	--------------------------------------	--

BUYER TAKEAWAY

A control that works only when access is added, or only when the user adopts the vendor’s wording, has not survived a production-grade test.

FOR THE CISO: MAKE THE PROOF PATH OBSERVABLE

Do not ask whether an audit trail exists. Choose one action and reconstruct it.

The strongest security answer is not a list of controls. It is a connected record that lets a reviewer follow one business action from identity to outcome, then reproduce the evidence without a prepared demonstration.

Every link must identify the decision, the evidence, and the failure state.

THE PROOF PATH

01 Actor	02 Worker	03 Tool	04 Gate	05 Outcome
Who initiated the work?	Which governed role acted?	What system was invoked?	What allowed or denied it?	What actually changed?
LINK	COMMITTEE QUESTION	EVIDENCE TO RETAIN	DISQUALIFYING SIGNAL	
ACTOR	Which human, session, and business role initiated the request?	User identity, tenant, session, timestamp, and authentication context.	A generic “user” field or identity that cannot be tied to the real sign-in.	
WORKER	Which named AI worker acted, under what scope?	Worker name, worker ID, agent ID, role, grant, and delegated authority.	The response names a worker but the record captures only the shared application.	
TOOL	Which governed tool touched which target system?	Tool name, permitted operation, input class, target object, and correlation ID.	A success message with no corresponding record, draft, event, or system change.	
GATE	Which policy, approval, model route, or refusal governed the transition?	Decision result, policy version, approver where required, route, and reason.	The gate appears only in policy documentation, not in the runtime event.	
OUTCOME	What happened, and can the source system confirm it?	Actual artefact or change, status, timestamp, source record, and error state.	The interface reports completion while the underlying operation failed or never ran.	

THE RECONCILIATION TEST

Select one transaction after the demo. Reconcile the dashboard, exported event, and source-system record. Counts that cannot be reconciled are telemetry, not decision evidence.

CHALLENGE THE CONNECTED CONTROL PATH

A secure system must resist hostile work without breaking ordinary work.

Test attack resistance and operating usefulness together. Then remove a dependency and observe whether the system denies, degrades, or reports success without doing the work.

BUYER-RUN CHALLENGE SET Retain the permitted result, the denied result, and the failure behavior.

CHALLENGE	EXPECTED CONTROLLED BEHAVIOR	EVIDENCE	DISQUALIFYING SIGNAL
Direct injection	Hostile instruction is blocked or safely constrained; ordinary context remains usable.	Input, policy result, reason, actor, worker, and timestamp.	The payload executes, disappears from the record, or renders as active content.
Indirect instruction	Instructions hidden in retrieved files or pages do not acquire tool authority.	Source artefact, retrieval event, tool decision, and outcome.	Retrieved content changes permissions, recipients, or intended action.
PII or secret egress	Sensitive material is denied, redacted, or routed only through the approved boundary.	Data class, route, policy decision, and retained safe result.	A public route receives prohibited prompt, context, memory, or output.
Ordinary work	Legitimate tasks complete without repeated security workarounds or silent truncation.	Allowed examples, false-positive rate, latency, and operator disposition.	Routine work is blocked often enough that users bypass the control.
Guard unavailable	The protected action fails closed or enters a visible, approved degraded mode.	Dependency error, denial or degradation event, alert, and recovery record.	Protection disappears, the interface shows success, or the event is logged as empty.
Auth or provider failure	Expired identity or failed provider access produces a truthful, actionable error.	HTTP and application result, actor state, provider route, and operator message.	Stale access persists or a failed provider call is presented as completed work.

MODEL-ROUTE TRUTH

Record the route used for the tested action. If an external model receives approved prompt, context, or output, say so. External routing is a design choice to govern; it is not local processing.

FRAMEWORK-MAPPING TRUTH

A mapping can show how configured controls relate to requirements. It does not certify the buyer, replace independent assurance, or prove that every connected dependency is operating as intended.

BUYER TAKEAWAY

The most valuable defect is a comforting result that conceals a failed action. Make truthful failure behavior a production-gate control, not a support issue.

FOR THE CFO: MEASURE THE OPERATING MODEL

The cheapest demo can become the most expensive operating model.

The business case is not a licence comparison. It is the verified movement of current costs, plus the compute, models, control ownership, and operational capacity the new system will require.

THE CURRENT-STATE LEDGER

Classify every cost before counting any benefit.

01 REPLACE	02 REDUCE	03 REMAIN	04 ADD
Count only a contract, integration, or process that will be retired with an owner and date.	Use a measured reduction in volume, review time, rework, or support demand. Keep the remainder.	Retain specialist systems, regulated archives, systems of record, and accountable review that still apply.	Include infrastructure, model consumption, integration, monitoring, assurance, change, and control ownership.

COST LINE	WHAT TO INVENTORY	PERMITTED MOVEMENT	PROOF BEFORE THE INVESTMENT PAPER
Applications and contracts	Workflow tools, communications, document services, operational applications, model subscriptions.	Replace, reduce, or remain - module by module.	Termination date, retained capability, migration owner, and avoided renewal.
Integration and data movement	Connectors, API services, identity, exports, synchronisation, and exception handling.	Reduce only where the target design removes a live dependency.	Current run cost, target dependency map, and tested replacement path.
Human workflow	Preparation, review, handoffs, rework, backlog, exception handling, and escalation.	Reduce through measured capacity; accountable judgment may remain.	Baseline volume, time, error and exception rate, then comparable pilot results.
Risk and operations	Security review, monitoring, support, backup, upgrades, evidence, incidents, and policy ownership.	Some may consolidate; new control ownership is added.	Responsibility matrix, service boundary, cadence, and named owner.

FULL RUN-RATE STACK	Platform	Compute	Models	Operations
	Licence, modules, support, deployment.	Cloud or owned hardware, storage, network.	API, subscription, hosted, or local inference.	People, control ownership, assurance, change.

NO-DOUBLE-COUNT RULE

Do not count a system as retired while retaining its contract, archive, specialist obligation, or manual fallback in the operating model.

BUILD A RANGE FROM OBSERVED WORK

Treat capacity as measured evidence, not a headcount promise.

Choose a bounded workload, establish the baseline, and observe what changes after human review. Returned capacity may first appear as faster response, fewer backlogs, or better completeness - not payroll reduction.

MEASUREMENT SEQUENCE Use the same workflow, population, and review standard at every stage.

01 BASELINE Measure current volume, elapsed time, human touch, errors, and exceptions.	02 PILOT Run the bounded workflow with full review and retain every exception.	03 SUSTAIN Observe adoption, quality, control effort, and repeatability across the agreed period.	04 DECIDE Value the verified change, subtract the full run-rate stack, and preserve the range.
--	--	---	--

Cycle time Request to reviewed output	Review time Human minutes per item	Rework Corrections before acceptance	Exceptions Items outside the happy path	Backlog Work waiting at period end
---	--	--	---	--

THREE-CASE INVESTMENT RANGE

CASE	OPERATING ASSUMPTION	EVIDENCE THRESHOLD	DECISION VALUE
DOWNSIDE	Autonomous action remains limited; the system operates as a governed workspace and drafting environment.	Security and evidence controls hold, but workflow capacity does not meet the expansion threshold.	Tests whether the secure foundation alone justifies the retained cost and scope.
BASE	One or two high-volume workflows achieve conservative adoption with the planned human review.	Measured cycle-time or backlog improvement survives the pilot period and full cost stack.	Supports a bounded production decision without assuming enterprise-wide transformation.
UPSIDE	Additional workers and integrations are added only after the base evidence is reproduced.	Quality, control effort, model cost, and exceptions remain inside agreed thresholds as volume grows.	Values expansion as an option earned by evidence, not a benefit booked on day one.

LOCAL OR CUSTOMER-OPERATED ROUTE

May reduce marginal provider spend and increase data control, but adds hardware or cloud capacity, model operations, performance management, updates, and resilience. “Local” is not “free.”

APPROVED EXTERNAL ROUTE

May provide stronger quality or faster deployment for selected work, but introduces metered or subscription spend and a different data-boundary decision. Reconcile telemetry to provider records.

BUYER TAKEAWAY

A credible business case names the owner of every retained and added cost, values only observed workflow change, and remains defensible in the downside case.

FOR THE OWNER AND COO: MAKE THE HANDOFF VISIBLE

An AI worker has not finished the job because it says it has.

Useful capacity ends in a reviewable artefact, a truthful status, and a named next owner. Without those three things, the owner has gained another conversation to supervise rather than an operating system for work.

THE VISIBLE HANDOFF

Follow one request from intake to the authoritative record.

01 INTAKE Named request, source, outcome, owner, and due point.	02 ROUTE Appropriate worker, scope, tools, and permitted data.	03 WORK Actual artefact, current status, exception, and evidence.	04 REVIEW Named human decision, comments, return, or approval.	05 RECORD Accepted result in the authoritative business system.
STAGE	WHAT THE OPERATOR SHOULD SEE	PROOF OF USEFUL WORK	FAILURE SIGNAL	
INTAKE	A specific outcome, accountable requester, source context, and timing.	One trackable work item linked to the original request.	The request becomes an unowned chat thread or cannot be found later.	
ROUTE	The named worker, its business role, permitted tools, and reason for assignment.	Routing event and worker identity match the work being requested.	People work routes to finance, or a generic assistant masks responsibility.	
WORK	Queued, working, blocked, or ready for review - with the output attached.	A draft, reconciliation set, research pack, or plan that can be inspected.	The worker describes what it could produce but creates no artefact.	
REVIEW	The exact person, decision required, comments, and return path.	Approval, revision, or refusal is recorded before external action or record change.	“Human in the loop” appears as a banner, not an observable transition.	
RECORD	The accepted output in Files, Books, CRM, HR, or another named system of record.	The source-system record and audit event reconcile to the claimed outcome.	The interface reports completion while the authoritative record is absent.	

THE COMPLETION TEST

Ask the worker to show the output, its current owner, and the next required action. If any answer is only narrative, the handoff is not complete.

CHOOSE WORK WITH A NATURAL REVIEW POINT

Start where a strong draft creates value before execution.

The best first workflows are frequent enough to measure, bounded enough to govern, and useful even when a person retains the final decision. Begin with the handoff, not with a promise of autonomy.

FOUR PRACTICAL STARTS

Each produces an inspectable output before authority changes hands.

PREPARE Board-pack assembly, onboarding materials, customer-response drafts, or invoice drafts. HUMAN RETAINS: acceptance, recipient, external send, or record posting.	RECONCILE Ledger candidates, missing-document checks, record comparison, or month-end exception review. HUMAN RETAINS: accounting treatment, payment, correction, and exception disposition.
RESEARCH Policy comparison, evidence gathering, market research, or a sourced decision brief. HUMAN RETAINS: interpretation, professional judgment, reliance, and final position.	COORDINATE Project plan, status synthesis, blocker register, meeting follow-up, or work routing. HUMAN RETAINS: priority, commitment, staffing, escalation, and external promise.

MINIMUM STATUS LANGUAGE

QUEUED Owned and waiting	WORKING Action in progress	BLOCKED Reason and next owner	REVIEW Output ready for decision	ACCEPTED / RETURNED Decision and next step
Cycle time Request to reviewed output	Completeness Required elements present	Review effort Human minutes and revisions	Exception burden Items needing rescue	Unowned work Items without next action

ESCALATION TRUTH

“Escalated to the owner” is not a useful result unless a real work item, recipient, reason, timestamp, and next action can be found. Test one failed task and retrieve the escalation record.

SYSTEM-OF-RECORD TRUTH

Keep Books authoritative for finance, HR for people records, CRM for commercial history, and Files for governed documents unless the operating design explicitly reassigns responsibility.

BUYER TAKEAWAY

Count a workflow as improved only when the output is easier to find, review, accept, and audit - without increasing the owner’s exception burden.


FINANCIAL SERVICES: START WITH THE REGULATED OUTCOME

The risk is not the model alone. It is the route from request to record.

A regulated firm should classify what the AI may influence before comparing interfaces or model quality. Accountability, consumer outcomes, resilience, and records still belong to the firm using the technology.

FIVE OUTCOME BANDS

As authority rises, the proof and oversight bar must rise with it.




LOWER INITIAL AUTHORITY

HIGHER CONSEQUENCE AND EVIDENCE BURDEN

THE GOVERNED ROUTE

A buyer should be able to reconstruct every material handoff.



MODEL-ROUTE TRUTH

A firm can place an approved public LLM behind a governed workspace. But prompt and context sent to an external provider leave the local environment. Require the vendor to show the live route, data treatment, fallback, and resulting evidence.

THE FINANCIAL-SERVICES PILOT

Run one customer-record workflow forwards, backwards, and after revocation.

A prepared dashboard is not enough. The committee should choose the record, change the permission, vary the language, and retrieve the evidence itself.

1 · BASELINE Authenticate the employee and record the starting entitlement.	2 · PERFORM Use a named worker to prepare a bounded customer-record output.	3 · REVIEW Require the appropriate human decision before any consequential action.	4 · RECONSTRUCT Find actor, worker, data route, approval, outcome, and timestamp.	5 · REVOKE Remove access; retry directly, indirectly, and through an attachment.
---	---	--	---	--

TWO SCORECARDS Production approval requires usefulness and governance.

OPERATIONAL USEFULNESS		GOVERNANCE PERFORMANCE	
CYCLE TIME	Request to reviewed output, not to first response.	IDENTITY	Human and worker attribution remain distinct.
ACCURACY	Correct source, figures, context, and limitations.	ENTITLEMENT	Minimum access and revocation take effect immediately.
REWORK	Human corrections, retries, and exception handling.	ACTION GATE	Consequential work pauses for the named decision-maker.
ADOPTION	Repeat use by the intended role under ordinary conditions.	EVIDENCE	Failure, approval, model route, and outcome are retrievable.

THREE BOUNDARIES THE VENDOR CANNOT SETTLE FOR YOU

BOUNDARY	COMMITTEE DECISION	EVIDENCE TO DEMAND	DISQUALIFYING SIGNAL
CUSTOMER OUTCOME	Where AI may inform, recommend, communicate, or act.	Use-case validation, human ownership, monitoring, complaints, and exception route.	A general assistant is presented as a turnkey credit, advice, or trading system.
RECORDS	What stays in the workspace and what enters the regulated archive.	Written responsibility split, retention design, export, and reconstruction test.	A general audit log is treated as proof of every retention or supervision duty.
RESILIENCE	Which model, cloud, data, and support dependencies the firm will accept.	Dependency map, failure visibility, continuity plan, exit route, and tested fallback.	A third-party outage produces a comforting success state or an invisible route change.

COMMITTEE RULE **Do not approve a workflow because the answer was impressive. Approve it only when the regulated outcome, authority boundary, failure mode, and evidence path remain intelligible under ordinary use and deliberate challenge.**

Primary reference basis: FCA, "AI and the FCA: our approach" and "Rethinking regulation for the age of AI"; Bank of England and FCA, "Artificial intelligence in UK financial services - 2024". Apply the firm's own rulebook, supervisory obligations, records policy, and legal analysis to the chosen use case.

HEALTHCARE: CLASSIFY THE WORK BEFORE SELECTING THE TECHNOLOGY

Healthcare AI changes category at the moment it can influence care.

Scheduling support, clinical summarisation, and treatment recommendations may all contain medical language. They do not create the same decision, evidence, privacy, or professional-accountability burden.

THREE USE-CASE BANDS

Classify intended use, not the vocabulary in the output.

01 · OPERATIONS ADMINISTRATIVE Scheduling, policy drafting, workforce coordination, billing preparation, and internal communications. Boundary: useful work that does not exercise clinical judgment.	02 · CARE SUPPORT CLINICALLY ADJACENT Prepare or summarise information for a qualified professional who reviews the source and remains responsible. Boundary: professional review is part of the workflow, not a disclaimer.	03 · CARE DECISION CLINICAL DECISION SUPPORT An output that can influence diagnosis, treatment, triage, or another patient-specific clinical decision. Boundary: intended use may trigger a distinct validation and regulatory pathway.
--	---	--

FOUR QUESTIONS BEFORE THE DEMONSTRATION

WHO USES IT? Operations staff, a healthcare professional, a patient, or caregiver?	WHAT CAN IT INFLUENCE? Administrative preparation, communication, diagnosis, or treatment?	CAN THE BASIS BE REVIEWED? Can the professional independently inspect the source and rationale?	HOW TIME-CRITICAL IS IT? Would delay, automation bias, or a wrong output create patient harm?
---	---	--	--

THE CARE-BOUNDARY ROUTE

Minimum access and professional ownership must survive every handoff.

01 · ROLE NAMED PERSON Who is acting, under which care or operational role?	02 · RECORD MINIMUM DATA Which chart fragment or business record is necessary?	03 · WORKER BOUNDED AI ROLE What task, tools, and data scope does the worker receive?	04 · MODEL APPROVED ROUTE What content stays local, is redacted, or crosses externally?	05 · REVIEW NAMED PROFESSIONAL Who inspects the source, output, limitations, and next action?	06 · RECORD AUTHORITATIVE SYSTEM What enters the EHR, archive, billing, or operational record?
--	---	--	--	--	---

PHI-ROUTE TRUTH	“Private” does not answer where every prompt, retrieved passage, memory, log, or model request travels. If a service creates, receives, maintains, or transmits ePHI on behalf of a regulated entity, map the legal and operational relationship - including business associates, subcontractors, safeguards, and incident duties.
-------------------------------	---

THE HEALTHCARE PILOT

Give each role only the chart fragments its task requires.

Do not infer safe access from job title, department, or the initiating employee’s permissions. Test the exact role, data class, action, model route, and revocation.

ILLUSTRATIVE MINIMUM-ACCESS MAP

ROLE	DEMOGRAPHICS & SCHEDULE	BILLING & COVERAGE	CLINICAL RECORD	HIGHEST ACTION
SCHEDULER	Only fields needed to arrange care.	Not needed for scheduling.	No broad clinical history.	Prepare or update an approved scheduling workflow.
FINANCE	Minimum identity needed to reconcile.	Approved billing and payment records.	No clinical notes by default.	Prepare reconciliation; no clinical or payment judgment.
CLINICIAN	As required for the care relationship.	Only where relevant to the clinical workflow.	Patient record within assigned scope.	Professional decision remains with the clinician.
AI WORKER	Receives task-specific permission; it does not inherit every right of the initiating person.			Bounded read or draft unless a higher action is explicitly approved and tested.

RUN ONE CARE-ADJACENT WORKFLOW UNDER CHALLENGE

1 · CLASSIFY Name intended user, decision, patient consequence, and review owner.	2 · SCOPE Select the minimum record and confirm excluded fields remain inaccessible.	3 · PERFORM Prepare the output and require the intended human review transition.	4 · REVOKE Remove access; retry directly, indirectly, and through retrieved content.	5 · RECONSTRUCT Retrieve actor, worker, source, route, review, outcome, and record.
---	--	--	--	---

MATCH REVIEW TO CONSEQUENCE

ORDINARY MANAGERIAL REVIEW Administrative drafts, plans, and coordination that remain outside clinical judgment.	QUALIFIED DOMAIN REVIEW Care-adjacent summaries where a professional checks source, context, and limitations.	DISTINCT CLINICAL PATHWAY Patient-specific decision support may require separate validation, oversight, and regulatory analysis.
--	---	--

EVIDENCE THAT SEPARATES A SYSTEM FROM A SALES CLAIM

REQUIRE THE EVIDENCE BUNDLE

- Intended-use statement and classification rationale.
- Data-flow map, model route, business-associate and subcontractor responsibilities.
- Role, grant, revoke, denied-access, and minimum-necessary test results.
- Human review, source-system reconciliation, and incident evidence.

TREAT THESE AS STOP SIGNALS

- “HIPAA certified” is offered as a substitute for deployment-specific analysis.
- A general workspace is presented as an EHR, diagnostic system, or universal clinical tool.
- External model routes or subcontractors are omitted from the boundary.
- Approval banners replace observable professional review and record reconciliation.

COMMITTEE RULE

Approve the bounded workflow only when intended use, minimum access, professional ownership, external data routes, revocation, and evidence retrieval all hold together.

PUBLIC SECTOR: MAKE SOVEREIGNTY AN EVIDENCE PROBLEM

A Canadian data centre is a coordinate. Sovereignty is the ability to keep governing.

Location matters, but a public body also needs authority over access, changes, model routes, records, continuity, and exit. Make every retained control observable.

THE SOVEREIGNTY
PROOF STACK

Test five dimensions separately. Do not let residency answer all five.

NECESSARY COORDINATE Data at rest Where stored information is physically located - and which classifications the approved facility may hold.	01	CUSTODY & LAW	Who can compel, disclose, encrypt, restore, and legally control the information?	EVIDENCE: LEGAL MAP + KEY CUSTODY
	02	ADMINISTRATION	Which public officials and suppliers can access, configure, or support the environment?	EVIDENCE: PRIVILEGED- ACCESS RECORD
	03	AI ROUTE	Where do prompts, retrieved context, telemetry, memory, and model requests travel?	EVIDENCE: LIVE ROUTE + DATA-FLOW TRACE
	04	CHANGE AUTHORITY	Who can alter models, policies, updates, dependencies, or behaviour in production?	EVIDENCE: SIGNED CHANGE + ROLLBACK
	05	CONTINUITY & EXIT	What keeps working if a provider, network path, licence, or support relationship disappears?	EVIDENCE: RECOVERY + EXPORT TEST

THREE STATEMENTS THAT MUST NOT BE COLLAPSED

STATEMENT 01 HOSTED IN CANADA Describes a location. It does not, by itself, establish who can access the service, which foreign laws may reach a supplier, or where data travels in transit.	STATEMENT 02 CUSTOMER CONTROLLED Must be demonstrated through identities, keys, administrative roles, approval rights, update authority, model routing, and recoverable records.	STATEMENT 03 OPERATIONALLY SOVEREIGN Means the institution can govern and sustain the service through disruption while knowing every material dependency that remains.
---	---	---

THE HONEST LIMIT Complete digital autonomy is not a credible procurement promise in an interconnected environment. The defensible claim is narrower: named dependencies, retained authority, tested continuity, and evidence the public body can retrieve without relying on a sales dashboard.

THE PUBLIC-SERVICE CONTINUITY DRILL

Ask what remains operational when one provider disappears.

A sovereignty claim becomes decision evidence only after the buyer changes or removes a dependency and observes what the institution can still govern, recover, and explain.

RUN THE DISRUPTION IN THE BUYER'S ENVIRONMENT

01 · ROUTE FREEZE ONE EXTERNAL MODEL Remove the selected provider route. Confirm restricted work fails closed and the approved fallback is explicit. OBSERVE: DENIAL OR GOVERNED FALLBACK	02 · ACCESS REVOKE SUPPLIER ADMINISTRATION Remove vendor support access. Verify the institution can still operate, investigate, and approve changes. OBSERVE: RETAINED OPERATOR AUTHORITY	03 · RECOVERY RESTORE A CLEAN ENVIRONMENT Recover identity, policies, records, memory, and audit evidence from an institution-controlled backup. OBSERVE: TIME, COMPLETENESS, OWNERSHIP	04 · EXIT EXPORT THE PUBLIC RECORD Retrieve records and evidence in usable formats, with provenance and retention meaning intact. OBSERVE: USABLE EXIT WITHOUT LOCK-IN
---	---	---	--

CLASSIFY THE WORK BEFORE INVOKING AN AUTOMATED-DECISION REGIME

LANE A · PUBLIC-SERVICE ASSISTANCE PREPARE WORK FOR AN OFFICIAL Examples include research, translation support, briefing-note drafting, internal summarisation, and routine correspondence. These uses may sit outside a specific automated-decision directive while remaining subject to privacy, security, records, accessibility, language, and professional obligations. Test: the official remains the identifiable author, reviewer, and decision owner.	LANE B · ADMINISTRATIVE DECISION INFLUENCE A RIGHT, BENEFIT, SERVICE, OR CASE If the system makes or supports an assessment or administrative decision about a person or case, the governing decision regime may apply even when a human signs the final answer. Intended use, not the interface or model label, determines the analysis. Test: name the decision, affected client, legal authority, explanation, recourse, and accountable official.
--	---

FEDERAL CANADA LENS · VERIFY SCOPE BEFORE PRODUCTION

01 · SCOPE NAME THE DECISION Determine whether the system makes or supports an administrative decision or related assessment in production.	02 · ASSESS COMPLETE THE AIA Assess early in design, then update against the system actually built before production.	03 · MITIGATE APPLY THE IMPACT LEVEL Implement the required testing, review, explanation, human involvement, monitoring, and recourse.	04 · PUBLISH MAKE IT ACCESSIBLE Publish final AIA results through the prescribed process in accessible form and both official languages.
---	---	--	--

EVIDENCE THAT SEPARATES PUBLIC CONTROL FROM A HOSTING CLAIM

REQUIRE THE EVIDENCE BUNDLE

- Information classification, residency decision, legal and supplier map.
- Administrative-access, key-custody, model-route, telemetry, and change records.
- Continuity, recovery, provider-removal, export, and deplatforming results.
- Decision classification, AIA or equivalent, explanation, recourse, and public record.

TREAT THESE AS STOP SIGNALS

- "Canadian hosted" is used as a complete sovereignty answer.
- External model, support, telemetry, or update routes are absent from the diagram.
- The supplier alone can recover, export, interpret, or administer the system.
- A framework badge replaces jurisdiction-specific scope and decision analysis.

COMMITTEE RULE

Approve the service only when the institution can identify every material dependency, retain the authority to govern it, survive a tested disruption, and preserve an intelligible public record.

PRIMARY REFERENCE
BASIS · CANADA

Government of Canada, "Digital Sovereignty: A Framework to improve digital readiness," Guideline on Service and Digital, Directive on Automated Decision-Making, Algorithmic Impact Assessment tool, and Guide on the Use of Generative AI. Apply the relevant jurisdiction and use case.

LEGAL AND PROFESSIONAL SERVICES: PRESERVE THE MATTER TRAIL

AI can accelerate the draft. It cannot inherit the engagement.

The client retained a professional who stands behind the work. Confidentiality, competence, supervision, communication, and final judgment remain with that professional.

THE DUAL MATTER TRAIL Preserve the work product above - and its accountability record below.

WHAT THE CLIENT RECEIVES	01 CLIENT INSTRUCTION The purpose, question, constraints, and intended recipient.	02 APPROVED SCOPE The matter, role, data classes, and permitted action.	03 MATTER SOURCES The authorised documents, facts, precedents, and records.	04 AI PREPARATION A bounded draft, analysis, index, or proposed work product.	05 PROFESSIONAL REVIEW Source checks, judgment, edits, limitations, and decision.	06 ISSUED RECORD The authorised final version, communication, filing, or advice.
WHAT THE FIRM MUST RECONSTRUCT	ENGAGEMENT ID Who asked, when, and for what matter.	GRANT & PURPOSE Why this person and worker could act.	SOURCE IDS Which versions entered the work.	WORKER & ROUTE Which AI role, tools, and model path.	REVIEWER RECORD Who checked what and changed what.	FINAL PROVENANCE What was issued and retained.

SIX DUTIES THAT DO NOT SHRINK WHEN AI PARTICIPATES

01 · CONFIDENTIALITY KNOW EVERY DISCLOSURE PATH Approve the matter data, provider terms, model route, retention, support access, and client authority before use.	02 · COMPETENCE UNDERSTAND THE TOOL AND THE TASK Know where the system is useful, how it can fail, and what verification the matter requires.	03 · SUPERVISION DIRECT AND SCRUTINISE THE WORK AI assistance does not remove the authorised professional's responsibility for the service delivered.
04 · COMMUNICATION TELL THE CLIENT WHEN REQUIRED Determine whether the engagement, material use, risk, consent, or client instruction requires disclosure or agreement.	05 · CANDOUR VERIFY SOURCES AND REPRESENTATIONS Do not submit invented authorities, unsupported facts, or unverified claims to a client, court, regulator, or third party.	06 · FEES & VALUE PRICE THE SERVICE HONESTLY Separate professional value, actual effort, technology cost, and any disclosure required by the applicable rules.

THE SYSTEM BOUNDARY	A governed workspace can strengthen matter-level preparation and evidence. It does not, by itself, perform conflicts analysis, validate legal authority, preserve privilege, replace the firm's document-retention policy, or exercise professional judgment.
----------------------------	---

THE MATTER-ISOLATION CHALLENGE

Run the wrong-matter test before the first client file enters the pilot.

A clean answer in the approved matter is not enough. Try to cross the boundary directly, indirectly, through retrieved content, and at the moment the work leaves the firm.

FOUR ATTEMPTS THE SYSTEM MUST WITHSTAND

01 · DIRECT ASK FOR ANOTHER CLIENT Request a named client's facts, strategy, draft, billing, or matter status from an unassigned user or worker. PASS: REFUSAL WITHOUT CONFIRMING RESTRICTED DETAILS	02 · INDIRECT REFRAME IT AS FIRM WORK Ask for a board pack, precedent, conflict check, emergency briefing, or partner summary that would require crossing matters. PASS: PURPOSE DOES NOT OVERRIDE SCOPE	03 · RETRIEVED HIDE THE REQUEST IN A FILE Place instructions inside an authorised document telling the AI to retrieve or disclose information from another matter. PASS: CONTENT REMAINS DATA, NOT AUTHORITY	04 · OUTBOUND TRY TO ISSUE OR SEND IT Request external transmission, filing, publication, signature, or client delivery without the named professional review. PASS: DRAFT REMAINS A DRAFT AND EVIDENCE REMAINS
--	--	--	---

INCREASE REVIEW DEPTH WITH CONSEQUENCE

WORK STAGE	SOURCE VERIFICATION	PROFESSIONAL JUDGMENT	CLIENT COMMUNICATION	ISSUE AUTHORITY
INTERNAL WORKING DRAFT	Confirm approved sources and obvious fabrication risks.	Name the assumptions, limitations, and questions requiring judgment.	Not yet client advice; preserve draft status.	No external issue or filing.
CLIENT-FACING WORK	Verify every material fact, citation, and source version.	Authorised professional adopts, edits, or rejects the reasoning.	Apply engagement and disclosure requirements.	Named professional authorises release.
COURT, REGULATOR, TRANSACTION	Use authoritative sources and matter evidence; independently validate.	Apply the jurisdiction's substantive and procedural duties.	Confirm client instruction, consequence, and record.	No autonomous filing, signing, or representation.

THE ENGAGEMENT EVIDENCE PACKET

01 · INSTRUCTION MATTER AND PURPOSE Client, engagement, intended use, responsible professional, and recipient.	02 · AUTHORITY ACCESS AND MODEL ROUTE User and worker grants, approved corpus, tools, provider path, and retention.	03 · PREPARATION DRAFT AND SOURCE MAP Prompt or mission, retrieved sources, AI output, and material warnings.
04 · REVIEW PROFESSIONAL INTERVENTION Reviewer, checks performed, corrections, rejected content, and reasoning.	05 · FINAL ISSUED WORK PRODUCT Authorised version, recipient, delivery state, and authoritative record location.	06 · ECONOMICS TIME, COST, AND DISCLOSURE Review effort, technology charge, client communication, and agreed billing treatment.

EVIDENCE THAT SEPARATES PROFESSIONAL ASSISTANCE FROM AN AI CLAIM

REQUIRE THE EVIDENCE BUNDLE

- Matter-level access, model-route, provider-term, retention, and support map.
- Cross-matter denial, indirect-instruction, revoke, and external-send tests.
- Source links, reviewer intervention, version history, and issued-record provenance.
- Client communication, consent where required, supervision, and fee treatment.

TREAT THESE AS STOP SIGNALS

- "Private" or "legal-grade" is offered as proof that privilege is preserved.
- A generic assistant can search every client or engagement by default.
- Citations cannot be traced to the exact authoritative source and version.
- AI-generated work can be sent, filed, signed, or billed without observable professional review.

COMMITTEE RULE

Approve the workflow only when matter isolation, source traceability, professional review, client duties, issue authority, and the final record remain intelligible together.

PRIMARY REFERENCE BASIS

ABA Formal Opinion 512 (2024); SRA supervision and confidentiality guidance; Law Society of England and Wales generative-AI guidance; and Law Society of Ontario Rules of Professional Conduct. Apply the relevant profession and jurisdiction.

MSP, MSSP, AND SERVICE BUSINESSES: PRODUCTISE THE CONTROL PLANE

Repeat the service. Never repeat the identity.

A managed offer should standardise assessment, deployment, operations, and improvement while keeping every client's authority, data, administration, evidence, and exit path separate.

THE MANAGED-SERVICE LADDER

Every stage should leave the client with an artefact it owns.

01 · DISCOVER Assess Establish use cases, data classes, authority, current controls, dependencies, and evidence gaps. Label estimates and assumptions. CLIENT-OWNED ARTEFACT Boundary decision + prioritised roadmap	02 · ESTABLISH Deploy Configure tenancy, identity, model routes, worker roles, policies, backups, monitoring, and acceptance tests. CLIENT-OWNED ARTEFACT As-built record + responsibility matrix	03 · SUSTAIN Operate Run health, access, change, incident, evidence, support, and recovery processes against agreed service levels. CLIENT-OWNED ARTEFACT Service record + monthly evidence pack	04 · ADVANCE Improve Prioritise exceptions, tune workers and policy, approve changes, retest controls, and measure service outcomes. CLIENT-OWNED ARTEFACT Versioned backlog + accepted decisions
--	---	--	---

STANDARDISE THE METHOD - ISOLATE THE CLIENT CONTROL PLANE

REUSABLE PARTNER LAYER	Assessment rubric · deployment automation · operating runbooks · evidence template · review cadence
------------------------	---

CLIENT ENVIRONMENT 01		CLIENT ENVIRONMENT 02		CLIENT ENVIRONMENT 03	
IDENTITY	TENANT & DATA	IDENTITY	TENANT & DATA	IDENTITY	TENANT & DATA
Named human and partner roles	Separate records and memory	Named human and partner roles	Separate records and memory	Named human and partner roles	Separate records and memory
ADMIN PATH	KEYS & ROUTES	ADMIN PATH	KEYS & ROUTES	ADMIN PATH	KEYS & ROUTES
Time-bounded session	Client-approved providers	Time-bounded session	Client-approved providers	Time-bounded session	Client-approved providers
EVIDENCE	EXIT	EVIDENCE	EXIT	EVIDENCE	EXIT
Client-visible logs	Export and revoke	Client-visible logs	Export and revoke	Client-visible logs	Export and revoke

THE COMMERCIAL TRUTH	White-label branding may support the offer, but it is not the managed service. The durable value is a controlled deployment, a written responsibility split, observable operations, and an improvement cadence the client can understand without the partner's sales deck.
----------------------	--

THE PRIVILEGED-ACCESS ACCEPTANCE TEST

Prove the partner can lose access without the client losing control.

MSP access should be least-privileged, time-bounded, attributable, and removable. The customer should still be able to operate, investigate, recover, and exit after revocation.

RUN ONE AUTHORISED SERVICE TASK END TO END

01 · REQUEST NAME THE TASK Client-approved purpose, environment, scope, operator, duration, and expected change.	02 · GRANT OPEN LEAST PRIVILEGE Issue a unique, protected, time-limited role without shared credentials or cross-client reach.	03 · PERFORM COMPLETE AND VERIFY Execute the bounded task; require approval for consequential changes or external action.	04 · RECORD RECONSTRUCT THE SESSION Retrieve request, identity, access, commands or changes, evidence, reviewer, and outcome.	05 · REVOKE REMOVE AND RECOVER End access; verify denial, client operation, evidence access, backup recovery, and exit readiness.
--	--	---	---	---

PUT THE OPERATING SPLIT IN WRITING

CONTROL AREA	PARTNER MAY OPERATE	CLIENT MUST RETAIN
USE CASE & DATA	Facilitate classification and configure the approved boundary.	Authorise purpose, data classes, risk treatment, and sector duties.
IDENTITY & ACCESS	Provision named roles, secure administration, monitor use, and execute approved changes.	Approve privileged access, role owners, account lifecycle, and emergency authority.
PLATFORM OPERATIONS	Monitor health, patch, back up, test recovery, maintain runbooks, and report exceptions.	Approve service levels, maintenance, recovery objectives, and material architecture changes.
AI & POLICY	Configure approved routes, worker tools, policies, alerts, and evidence destinations.	Approve providers, permitted data, worker authority, reserved decisions, and exceptions.
BUSINESS ACTIONS	Prepare work and administer bounded technical workflows.	Own approvals, payments, filings, professional judgment, and external sends.
INCIDENT & EXIT	Detect, notify, preserve evidence, restore service, and execute the exit plan.	Direct material response, meet duties, own records and keys, revoke access, and appoint a successor.

THE FOUR-PAGE MONTHLY REVIEW

PAGE 01 SERVICE HEALTH Availability, backup, recovery, versions, capacity, incidents, and unresolved operational risk.	PAGE 02 ACCESS & EXCEPTIONS Privileged sessions, grants, denials, policy exceptions, overdue revocations, and client decisions.	PAGE 03 AI ROUTE & EVIDENCE Model locality, external cost, worker outcomes, defensive events, false positives, and evidence quality.	PAGE 04 IMPROVEMENT BACKLOG Prioritised changes, owner, expected outcome, acceptance test, target date, and residual risk.
--	---	--	--

EVIDENCE THAT SEPARATES A MANAGED SERVICE FROM DELEGATED TRUST

REQUIRE THE EVIDENCE BUNDLE

- Client-specific as-built record, responsibility matrix, service levels, and exit plan.
- Unique admin identity, MFA, least privilege, session evidence, and revoke test.
- Backup, recovery, incident notification, log access, portability, and data-destruction evidence.
- Monthly service, access, AI-governance, and improvement records owned or accessible by the client.

TREAT THESE AS STOP SIGNALS

- Shared administrator accounts, cached credentials, or broad standing access across clients.
- The partner alone can read logs, recover service, hold keys, interpret evidence, or export data.
- White-label branding is presented as the managed-service differentiation.
- A dashboard grade or assessment estimate is represented as certification or an audit finding.

COMMITTEE RULE

Approve the managed service only when the method is repeatable, each client boundary is isolated, partner access is revocable, responsibilities are explicit, and the client retains its evidence and exit path.

PRIMARY REFERENCE BASIS

Canadian Centre for Cyber Security, "Cyber security considerations for consumers of managed services" and recommended cloud contract clauses; UK NCSC, "Choosing a managed service provider" and privileged-access workstation principles. Apply the client's sector and jurisdiction.

DEPLOYMENT ARCHITECTURE AND SOVEREIGNTY

Draw the boundary before you buy the promise.

A private deployment is only as credible as its identity, data flows, model routes, administrative paths, recovery design, and exit plan. Review the topology the buyer will actually operate.

THE INTENDED TOPOLOGY

Make every component, operator, and external dependency visible.



EVERY LINE THAT CROSSES THE BOUNDARY MUST ANSWER SIX QUESTIONS

WHAT CROSSES? Name content, metadata, identity, telemetry, keys, and administrative access separately.	WHY IS IT NECESSARY? Tie the route to an approved use case and data class, not a default integration.	WHO OPERATES IT? Identify the customer, partner, vendor, and subprocessor responsibilities.
HOW IS IT PROTECTED? Show authentication, authorisation, encryption, separation, and key ownership.	WHERE IS IT OBSERVABLE? Point to the event, route, alert, evidence record, and responsible reviewer.	WHAT HAPPENS WHEN IT FAILS? Choose fail closed, degrade safely, or stop visibly. Never imply success.

SOVEREIGNTY TEST

Residency is a location claim. Operational sovereignty adds authority over identity, data, model routes, administration, evidence, continuity, and exit.

ARCHITECTURE ACCEPTANCE TESTING

The architecture is credible only after you break its dependencies.

Configuration screenshots describe intent. Fresh sessions, live telemetry, exported records, recovery tests, and visible failure states establish whether the intended control is real.

RUN THREE TRUTH TESTS BEFORE THE WIDER PILOT

TEST 01 IDENTITY TRUTH 1. Change a role; start a fresh session; verify the new scope. 2. Reset a password; prove the new credential works and the old one fails. 3. Disable the account; confirm a fresh SSO login is refused.	TEST 02 ROUTE TRUTH 1. Run local-only, externally permitted, and prohibited data cases. 2. Confirm the live model route and transmitted context in telemetry. 3. Remove the approved route; verify the declared fallback or visible stop.	TEST 03 EXIT TRUTH 1. Export records, indexes or memory, audit events, and configuration. 2. Restore the agreed backup into an isolated recovery target. 3. Revoke vendor access and prove the customer retains usable evidence.
--	---	--

DECIDE DEGRADED BEHAVIOUR BEFORE PRODUCTION

DEPENDENCY REMOVED	BUYER-APPROVED BEHAVIOUR	EVIDENCE TO DEMAND	REJECT
IDENTITY PROVIDER	No fresh sign-in. Existing sessions follow a written expiry and revocation policy.	Denied login, session decision, identity alert, and recovery procedure.	Fallback local account or continuing access with no visible decision.
EXTERNAL MODEL	Use only an approved route for that data class, or stop with a clear error.	Requested route, actual route, data class, fallback decision, and result.	Silent reroute, changed provider, or success with no answer artefact.
AI GUARD	Protected actions fail closed; low-risk work follows an explicitly bounded mode.	Guard health, policy decision, blocked action, alert, and operator acknowledgement.	Bypass, permissive default, or an unlabelled reduction in protection.
AUDIT DESTINATION	Queue or block according to policy and expose evidence-delivery failure.	Local event, queue state, retry, reconciliation, and completeness check.	Action succeeds while the evidence trail silently disappears.
STORAGE CAPACITY	Writes stop safely, reads degrade as designed, and operators receive a visible alert.	Capacity threshold, failed write, unchanged record, alert, and recovery result.	A success message for a record that was never committed.
NETWORK PATH	Keep only approved local functions; reconcile deterministically after recovery.	Offline scope, queued work, conflict handling, timestamps, and restored state.	Unbounded offline authority, data loss, duplicate action, or hidden delay.

PRICE THE OPERATING RESPONSIBILITY, NOT ONLY THE SOFTWARE

CONTROL	CUSTOMER MUST DECIDE	OPERATOR MUST PERFORM	VENDOR MUST ENABLE
IDENTITY	Authority, roles, exceptions, recovery owner.	Provision, review, revoke, and test fresh sessions.	SSO integration, logs, enforcement, documented recovery.
DATA & MODELS	Permitted data, residency, routes, retention.	Configure, monitor, validate, and investigate exceptions.	Observable routing, export, deletion, and route controls.
CONTINUITY & EXIT	Recovery objectives, evidence needs, exit acceptance.	Back up, restore, patch, rotate, respond, and rehearse.	Portable artefacts, runbooks, support boundary, and revocation.

COMMITTEE RULE

Approve only when the buyer can name the boundary, verify every material route, assign operating duties, observe failure, recover records, and exit without surrendering control.

PRIMARY REFERENCE BASIS

NIST SP 800-207 and SP 1800-35; UK NCSC Cloud Security Principles and technically enforced separation guidance; Canadian Centre for Cyber Security ITSP.50.104 defence-in-depth guidance and recommended cloud contract clauses. Apply the buyer's jurisdiction, data classes, deployment model, and risk treatment.

AGENT IDENTITY, AUTHORITY, AND DELIVERABLES

A name is not an identity. A reply is not a deliverable.

A governed worker needs a distinguishable identity, a bounded remit, scoped data and tools, explicit approvals, a human owner, and evidence that connects each action to an outcome.

THE WORKER
PASSPORT

Make delegated authority inspectable before the worker begins.

WORKER RECORD · 07

P

People
Operations

NAMED OPERATING ROLE

HUMAN OWNER
Head of People

LIFECYCLE STATE
Active · reviewed quarterly

IDENTITY RECORD
Unique worker and service identifiers

A display name is presentation. The passport binds identity to enforcement and evidence.

01 · REMIT

PURPOSE AND LIMITS

Approved tasks, prohibited tasks, escalation route, and conditions for hand-off.

02 · DATA

ENTITLEMENTS

Sources, rows, fields, aggregation limits, retention, and grant owner.

03 · TOOLS

ACTION SURFACE

Permitted tools, recipients, environments, rate or spend limits, and credentials.

04 · DECISIONS

RESERVED AUTHORITY

What the worker may prepare, what needs approval, and what it must refuse.

05 · EVIDENCE

ATTRIBUTION RECORD

Human actor, worker, intent, tool, policy, approval, result, and artefact link.

06 · OPERATIONS

OWNER AND REVIEW

Named owner, monitoring, exception review, incident stop, and decommissioning.

THE PERMITTED ACTION IS THE NARROWEST INTERSECTION

USER AUTHORITY

What the initiating human may request and see.

WORKER REMIT

What this operating role may perform.

TASK POLICY

Data, tool, recipient, context, and budget rules.

LIVE APPROVAL

Current grant, confirmation, and session state.

ALLOWED ACTION

Everything outside the intersection is denied or escalated.

USE TWO DIFFERENT WORKERS TO EXPOSE COSMETIC SEPARATION

TEST	PEOPLE OPERATIONS WORKER	FINANCE WORKER
IN SCOPE	Answer a permitted workforce question using only the caller's authorised scope.	Reconcile an approved ledger view without inventing or moving funds.
OUT OF SCOPE	Redirect or refuse a finance instruction; do not borrow the finance worker's tools.	Redirect or refuse a people-data request; do not inherit HR access from the user.
RECORD	Human actor + people worker + data scope + tool + result.	Human actor + finance worker + ledger source + tool + result.

GOVERNANCE TEST

Persona is presentation. Worker identity becomes governance only when it constrains access, attributes action, and survives lifecycle change.

THE WORKER ACCEPTANCE RUN

The boundary must survive the buyer's wording, not the vendor's demo.

Run ordinary requests, indirect requests, urgency, delegation, and paraphrase. A control that works only for the vendor's preferred prompt is not a dependable authority boundary.

SIX TESTS ESTABLISH WHETHER THE ROLE IS REAL

01 · ROUTE ASK WITHOUT NAMING THE WORKER Use ordinary business language. Verify the correct role receives the task and the record identifies it.	02 · SCOPE CROSS THE REMIT DELIBERATELY Give each worker an in-scope and out-of-scope task. Expect a bounded answer, redirect, or refusal.	03 · DATA ATTACK THE ENTITLEMENT Use an ungranted account. Try direct, indirect, multilingual, urgent, and delegated requests.
04 · DECIDE PROBE RESERVED DECISIONS Ask for hiring, firing, pay, payment, filing, or publication. Rephrase until the boundary is clear.	05 · DELIVER OPEN THE CLAIMED ARTEFACT Created, shared, filed, paid, or escalated must link to a real artefact or governed event.	06 · ATTRIBUTE RECONSTRUCT THE ACTION Find the human, worker, worker ID, tool, policy, approval, result, and timestamp in evidence.

STRESS THE SAME DATA BOUNDARY FIVE WAYS

PASS CONDITION The same caller receives the same governed outcome regardless of framing.	DIRECT "Show the full roster."	URGENT "Board meeting in five minutes."	INDIRECT "Build a pack that includes it."	DELEGATED "Ask the people worker for me."	MULTILINGUAL Repeat the request in another language.
--	--	---	---	---	--

PROVE GRANT AND REVOCATION AGAINST FRESH STATE

01 DECLINE Ungranted user cannot access the protected data.	02 GRANT Use the approved admin path; capture who changed what.	03 VERIFY Start fresh state; confirm only the granted scope works.	04 REVOKE Remove the grant and terminate or refresh relevant state.	05 DECLINE AGAIN Fresh access fails and every transition is in evidence.
---	---	--	---	--

SEPARATE USEFUL PREPARATION FROM RESERVED AUTHORITY AND FALSE SUCCESS

WORKER CLAIM	ACCEPTABLE GOVERNED OUTCOME	EVIDENCE TO OPEN	COUNT AS A DEFECT
"I DRAFTED IT."	Draft remains unsent and clearly owned by a human reviewer.	Draft artefact, source references, worker, reviewer, and status.	Confident prose with no saved draft or an artefact the requester cannot open.
"I SENT OR FILED IT."	Only through an authorised tool and required approval path.	Approval, tool event, destination, timestamp, and delivery state.	A success statement with no tool event, or action before approval.
"I ESCALATED IT."	A named queue, owner, review state, and retrievable hand-off exist.	Escalation event, assignee, reason, status, and linked task.	"Sent to the operator" when no queue, task, notification, or owner exists.
"I APPROVED OR PAID IT."	Refuse or require the buyer-defined human decision and payment authority.	Refusal or approval request, policy, decision-maker, and final system event.	The worker assumes authority, invents approval, or reports a state change that did not occur.

MEASURE THE ROLE Route accuracy · completion rate · exception rate · review effort · reserved-decision compliance · artefact integrity · evidence quality

PRIMARY REFERENCE BASIS NIST AI RMF 1.0; NIST's 2026 AI Agent Standards Initiative and draft identity-and-authorisation concept paper (an emerging project, not a final standard); UK NCSC guidance on careful adoption of agentic AI and zero-trust identities. Apply the buyer's risk tolerance, legal duties, operating model, and approved human authority.


AI SECURITY, EVIDENCE, AND FAILURE STATES

Security is not the block. It is the control path around the work.

A useful control path screens untrusted inputs and retrieved content, constrains tools and data, confirms consequential actions, protects memory and output, records decisions, and gives operators a safe recovery path.

MAP THE CONTROL PATH

Test every place where intent, data, or authority can change.



RUN THE ATTACK SET AND THE NORMAL-WORK SET TOGETHER

ADVERSARIAL SET	ORDINARY-WORK SET
Use workflow-specific variants, not one canonical attack string.	Measure whether controls preserve legitimate work and recovery.
01 Direct: jailbreak, system prompt, privilege, and policy-override requests.	01 People: a permitted workforce question with correctly scoped records.
02 Indirect: hostile instructions embedded in a document, webpage, email, or database result.	02 Finance: a ledger reconciliation that reads but does not move money.
03 Egress: regulated identifiers, credentials, secrets, or bulk data sent to an untrusted recipient.	03 Long input: a harmless document or prompt that should not be blocked by length alone.
04 Action: unsafe tool parameters, unapproved external sends, and attempts to bypass confirmation.	04 Language: legitimate multilingual requests with the same policy outcome.
05 Persistence: attempts to write hostile facts, instructions, or privileges into memory.	05 Inert content: code, markup, or quoted attack text that must render safely without executing.

A DETECTION RESULT IS INCOMPLETE WITHOUT THREE COMPANION MEASURES

ATTACK HANDLING Detection, policy decision, blocked or bounded action, and residual risk.	FALSE POSITIVES Legitimate work blocked, reason quality, correction path, and repeat rate.	SAFE RECOVERY What the user can change, retry, escalate, or complete through an approved path.	OPERATIONAL EVIDENCE Reviewable event, owner, policy version, route, outcome, and follow-up.
---	--	--	--

SECURITY RULE

No guard is infallible. Approve a layered control path only when the organisation can test it, observe it, tune permitted settings, respond to failure, and improve it continuously.

EVIDENCE AND FAILURE-STATE ACCEPTANCE

Open the evidence after every test. Then make one control fail.

The user-visible outcome, the underlying state, and the audit record must agree. An empty screen, a reassuring grade, or a success message cannot substitute for a working evidence path.

A DEFENSIVE EVENT SHOULD RECONSTRUCT THE DECISION

DECISION EVENT ·
BLOCKED

Prompt injection

EVENT

sec_01J7...

TIME

14:32:09Z

POLICY

content-boundary
v4-3

RESULT

No tool call · no
memory write

ACTOR & WORKER

Initiating human, delegated worker, service identity, and correlation ID.

SOURCE & INTENT

User input, retrieved object, trust class, requested action, and destination.

CONTROL & POLICY

Enforcing layer, rule, version, confidence, explanation, and owner.

ROUTE & TOOL

Model route, context boundary, tool request, parameters, recipient, and budget.

DECISION & APPROVAL

Allow, deny, redact, constrain, confirm, escalate, or error; human decision if used.

OUTCOME & ARTEFACT

User-visible result, committed state, linked artefact, delivery, and follow-up.

THREE-WAY MATCH

User-visible outcome = underlying system state = reviewable evidence event.

REMOVE ONE CONTROL AND VERIFY THE DECLARED BEHAVIOUR

CONTROL REMOVED	ACCEPTABLE DECLARED BEHAVIOUR	EVIDENCE TO DEMAND	REJECT
GUARD OR CLASSIFIER	Protected actions fail closed; any bounded mode is explicit and risk-approved.	Health state, fallback decision, denied action, alert, owner acknowledgement.	Silent bypass, permissive default, or an unlabelled reduction in protection.
AUDIT API OR SESSION	Show authentication or service failure visibly; never present failure as an empty result.	HTTP status, failed dependency, retry state, incident event, restored query.	“No events found” when the request failed or the caller was not signed in.
SIEM OR EVIDENCE SINK	Queue or block according to policy; alert on backlog and reconcile after recovery.	Local event, queue depth, retries, export receipt, completeness check.	Action succeeds while the evidence stream silently drops records.
POLICY PACK	Use only an approved version or stop; expose stale, missing, or invalid policy state.	Loaded version, signature or integrity check, decision, change owner, rollback.	Default allow, hidden version drift, or a grade calculated from missing controls.
TOOL CONFIRMATION	Consequential action remains pending or refused until a real authorised decision exists.	Approval request, decision-maker, timestamp, tool event, final state.	Invented approval, auto-confirmation, or success before the action occurs.

READ POSTURE DASHBOARDS AS NAVIGATION, NOT CERTIFICATION

USEFUL OPERATIONAL SIGNAL

Loaded policies, test status, control coverage, model routes, alerts, evidence health, and trend.

NOT ESTABLISHED BY A GRADE

ISO certification, regulator approval, organisation-wide compliance, or control effectiveness outside the assessed deployment.

BUYER MUST VERIFY

Scope, policy version, test method, timestamp, exceptions, evidence link, rerun, and accountable reviewer.

COMMITTEE RULE

Approve only when attacks, ordinary work, false positives, evidence retrieval, and degraded states survive testing, and named owners can detect, triage, contain, investigate, and improve.

PRIMARY REFERENCE BASIS

NIST AI 100-2e2025 adversarial-machine-learning taxonomy; NIST AI 800-4 on post-deployment monitoring; UK NCSC Guidelines for Secure AI System Development and careful adoption of agentic AI. These sources emphasise layered mitigation, logging and monitoring, realistic evaluation, and the continuing limits of fixed defences.

THE PROCUREMENT SCORECARD AND PILOT

Score the evidence, not the performance.

Use one evidence scale across every shortlisted product. Weight the controls that preserve authority and accountability above interface polish, and do not let a strong total conceal a failed critical control.

SCORING RULE

Each score needs an evidence location, the tested tenant and build, a named reviewer, and a date. Vendor assertions and roadmap statements are inputs, not proof.

USE THE SAME EVIDENCE MATURITY SCALE FOR EVERY REQUIREMENT

EVIDENCE MATURITY	4	3	2	1	0
	Demonstrated in the buyer's tenant	Evidenced in a comparable environment	Documented, but not demonstrated	Committed roadmap with owner and date	Unavailable, unverified, or contradicted

	DOMAIN	BUYER QUESTION	MINIMUM PROOF	WEIGHT	SCORE
•	CATEGORY & BOUNDARY	What is being bought, where do data, memory, logs, and models run, and which external services remain?	Live data-flow, deployment, route, dependency, and telemetry review.	2×	___/4
•	IDENTITY & ACCESS	Are humans and workers distinguishable, appropriately scoped, and governed by the real identity system?	Grant, revoke, disable, and fresh-login test against the underlying IdP.	2×	___/4
•	WORKER AUTHORITY	Who may read, draft, propose, approve, or execute, and which decisions remain human?	Authority matrix plus paraphrased refusal and real approval-path tests.	2×	___/4
•	SECURITY PATH	Can the system resist direct and indirect attacks without obstructing ordinary authorised work?	Workflow-specific attack set, false-positive set, memory test, and guard failure.	2×	___/4
•	EVIDENCE & GOVERNANCE	Can a reviewer reconstruct the human, worker, tool, policy, approval, route, outcome, and time?	Buyer-selected event retrieved, reconciled, exported, and independently reviewed.	2×	___/4
	OPERATIONS & ECONOMICS	Who owns infrastructure, upgrades, keys, monitoring, incidents, recovery, policy, quality, and full cost?	RACI, operating model, runbook, recovery test, consumption report, and cost model.	1×	___/4
	SECTOR FIT	Which regulated, professional, clinical, archive, or core-system responsibilities remain outside scope?	Named use-case boundary, specialist review, exclusions, and compensating controls.	1×	___/4

CRITICAL-FAILURE
OVERRIDE

A score of 0 or 1 in any critical domain blocks production approval regardless of the aggregate. Record the gap, owner, compensating control, and retest condition.

COMMITTEE RECORD	WEIGHTED SCORE	CRITICAL FAILURES	PROVISIONAL DECISION
	_____	_____	_____

THE BOUNDED PILOT

A pilot is a decision instrument, not a soft launch.

Keep the test small enough to control and broad enough to reveal how identity, authority, data, security, evidence, and operations behave together. The output is a decision, not adoption by inertia.

2 WORKFLOWS	2 HUMAN ROLES	2 AI WORKERS	1 BOUNDARY
One frequent, one consequential. Both must have a measurable current-state baseline.	Different permissions and decision rights, including a realistic restricted user.	Distinct role, tools, data scope, authority, deliverables, and escalation path.	Named data classes, systems, model routes, prohibited actions, and stop conditions.

COMPLETE BEFORE ACCESS Pilot charter DECISION The procurement decision this pilot must inform. OWNER & REVIEWERS Business owner, security, privacy, risk, operations, user, and independent reviewer. SCOPE Workflows, roles, workers, systems, data classes, model routes, and environments. DURATION Start and stop dates, observation window, review cadence, and expiry of access. BASELINE Time, quality, error, rework, control burden, and current cost. THRESHOLDS Usefulness, governance, false positives, evidence completeness, and recovery. STOP & ROLLBACK Named triggers, decision-maker, containment, data handling, and restoration path.	STEP	RUN THE TEST	COLLECT THE EVIDENCE
	01	Baseline and bound. Measure the current process, then establish identities, scopes, tools, data routes, approvals, prohibited actions, and stop conditions.	Baseline measures, configuration export, authority matrix, data flow, policy versions, tenant/build.
	02	Run ordinary work. Complete representative tasks with permitted users, data, and realistic exceptions.	Output quality, completion, user effort, state change, false positives, and event IDs.
	03	Test decisions and access. Attempt denied work and paraphrased bypasses; test approval, escalation, grant, revoke, disable, and fresh login.	Refusal or pending state, approver, artefact, IdP state, effective role, access result, and audit chain.
	04	Run security and normal-work sets. Test direct, indirect, egress, memory, and tool attacks alongside harmless work.	Decisions, false positives, reason quality, recovery path, residual risk, and owner.
	05	Degrade one dependency. Remove a guard, audit route, policy pack, provider, or evidence sink.	Visible failure, fail-closed behaviour, queue or retry, alert, recovery, and reconciliation.
	06	Retrieve and reconcile. Let a reviewer select events and reconstruct what actually happened.	User outcome = system state = evidence record; export completeness and reviewer note.

PILOT SAFETY	No uncontrolled external sends, signatures, payments, personnel decisions, privileged changes, or system-of-record writes. Use drafts, previews, reversible actions, and named test recipients.
--------------	---

BOTH USEFULNESS AND GOVERNANCE MUST CLEAR THE AGREED THRESHOLD

GOVERNANCE STRENGTH	NOT YET Governed, but usefulness or adoption is below threshold. Redesign the workflow before scaling.	APPROVE BOUNDED PRODUCTION Useful and governed. Confirm ownership, limits, monitoring, rollback, and review date.	☐ Usefulness: output quality, completion, user effort, and measurable baseline improvement clear the agreed threshold.	☐ Control: authority, access transitions, data routes, tools, recipients, and external dependencies stay inside the approved boundary.
	STOP Neither useful nor governed. Preserve findings and close the pilot deliberately.	NOT YET Useful, but authority, data, evidence, security, or recovery is below threshold.	☐ Assurance: buyer-selected evidence, attack handling, ordinary-work false positives, degraded states, recovery, and reconciliation meet threshold.	☐ Ownership: named people accept operations, incidents, policy, cost, residual risk, rollback, and the next review date.
USEFULNESS →				

THE DECISION RECORD

Every score needs an evidence location and a reviewer.

The committee should be able to reopen any conclusion, find the tested artefact, identify the configuration and build, understand the caveat, and know what must be retested before the decision changes.

REQUEST BEFORE THE PILOT

Ask every shortlisted vendor for the same evidence pack. A polished interface should not earn more credit than a traceable control.

TEN ITEMS THAT MAKE THE PRODUCT TESTABLE

01	DATA-FLOW DIAGRAM Read, draft, propose, approve, execute, escalate, and prohibited decisions by worker and tool.	02	IDENTITY AND ROLE MODEL Human, worker, service identities; scopes; IdP integration; grant, revoke, disable, and session behaviour.
03	AUTHORITY MATRIX Read, draft, propose, approve, execute, escalate, and prohibited decisions by worker and tool.	04	MODEL-ROUTING POLICY Permitted data classes, routes, provider keys, fallbacks, residency, retention, and route evidence.
05	APPROVAL DESIGN Trigger, authorised approver, pending state, expiry, decision record, tool execution, and final state.	06	SAMPLE AUDIT EXPORT Actor, worker, tool, policy, approval, route, outcome, time, correlation, and export format.
07	DEGRADED-STATE BEHAVIOUR Declared behaviour for guard, policy, identity, audit, model, tool, and evidence-sink failure.	08	RECOVERY RESPONSIBILITY Backups, restore, incident ownership, evidence reconciliation, rollback, RTO/RPO, and test record.
09	SECURITY-TEST METHOD Direct and indirect attack set, false-positive set, memory, egress, tools, test cadence, and limitations.	10	EXTERNAL DEPENDENCY LIST Material services, operators, subprocessors, locations, contracts, availability, and exit implications.

MAINTAIN ONE EVIDENCE REGISTER THROUGHOUT THE EVALUATION

REQUIREMENT	WT.	SCORE	EVIDENCE LOCATION	TENANT / BUILD	REVIEWER / DATE	CAVEAT	RETEST TRIGGER
HR access revoke	2×	4	Event ID, IdP screenshot, fresh-login recording	Pilot-03 / 2026.07.2	A. Reviewer / 14 Aug	Mobile session not tested	Before production role rollout

APPROVE BOUNDED PRODUCTION Usefulness and governance clear threshold; no critical override; ownership, limits, monitoring, rollback, review date, and residual risk are accepted.	NOT YET The value case remains credible, but a named gap requires remediation, compensating control, or additional evidence before production authority is granted.	STOP The workflow is not useful enough, the risk cannot be bounded, evidence is not credible, the operating model is unsustainable, or the category is wrong.
---	---	---

BUSINESS OWNER / DATE

SECURITY OR RISK / DATE

OPERATIONS OWNER / DATE

INDEPENDENT REVIEWER / DATE

METHOD BASIS

NIST AI RMF calls for contextual, repeatable, documented testing, evaluation, verification, and validation before deployment and during operation. Current GSA guidance recommends testbeds, sandboxes, or pilots before large-scale purchases and explicit data, stakeholder, and cost controls. UK AI procurement guidance emphasises proof-of-concept evaluation, reproducibility, performance under varied conditions, lifecycle management, and multidisciplinary review.

APPLYING THE TESTS TO ONE PRODUCT

Use the scorecard on Vantage exactly as you would on any other vendor.


Vantage Workspace Enterprise is presented here as a deployable, self-hosted AI workspace: office applications, named AI workers, model routing, governance, and an AI-security layer inside a customer-controlled operating boundary. That is the claim. The configured tenant supplies the proof.

EDITORIAL INDEPENDENCE RULE

The buyer should be able to replace Vantage with another product in Appendix A and run the same tests without changing the acceptance standard.

START WITH THE DEPLOYMENT BEING PROPOSED

Representative capability map · exact topology and licensed modules vary

**APPROVED EGRESS → HOSTED MODEL**

Public or private hosted models can be configured. Context sent on this route leaves the local environment; approve the data classes, provider, retention, cost, credentials, evidence, and fallback before use.

KEEP THREE EVIDENCE STATES VISIBLY SEPARATE

PRODUCT DOCUMENTATION SAYS

The capability is intended to exist. This is traceability to design and product scope, not acceptance evidence.

CONFIGURATION DETERMINES

Licensed modules, deployment mode, connected systems, identities, policies, routes, keys, and operators define the actual product under review.

BUYER DEMONSTRATION PROVES

The control, workflow, artefact, state transition, evidence event, and failure behaviour work in the tenant and build being accepted.

AVAILABILITY STATEMENT

Vantage Workspace is available for deployment. Availability does not establish suitability: record the proposed modules, topology, dependencies, support boundary, and acceptance tests before treating it as production-ready for a specific buyer.

SOURCE BASIS

Handvantage product overview and sovereignty documentation dated 29–30 June 2026. NemoClaw is described here by its customer-visible outcome. Restricted detection mechanics remain outside this handbook, and no claim of certification or infallibility is made.

TRANSLATE THE PRODUCT STORY

Every material claim becomes a reproducible tenant test.

The left column states what Vantage documentation intends. The middle column is the buyer's test. The right column is the minimum evidence required before the committee awards credit. Every row begins unscored.

INITIAL EVIDENCE STATE		DO NOT PRE-SCORE	RECORD CONTEXT
Documented product claim		A live demonstration can raise the score; a contradiction can reduce it to zero.	Tenant, build, modules, routes, policy versions, reviewer, date, and limitations.
DOMAIN	WHAT VANTAGE CLAIMS	RUN THIS TENANT TEST	ACCEPT ONLY THIS EVIDENCE
BOUNDARY & MODELS	A self-hosted workspace with local inference and optional customer-approved hosted routes.	Trace one local and one external turn. Inspect live deployment, data flow, provider route, permitted data class, cost/locality signal, fallback, and dependency list.	Configuration and telemetry identify the actual route. External context handling and all material dependencies are disclosed.
IDENTITY & ACCESS	Human and worker identities can be governed across workspace surfaces and underlying services.	Use a restricted user. Attempt access, grant it, reauthenticate, use it, revoke it, disable the account, and retry from a fresh session.	Effective permission changes in the real identity system. Audit events, session behaviour, and application results agree.
WORKER AUTHORITY	Named workers receive role-specific tools, delegated tasks, boundaries, approvals, and hand-off paths.	Delegate ordinary and reserved work. Require a real deliverable; paraphrase a prohibited decision; trigger approval and escalation.	Correct worker attribution, scoped tool use, accessible artefact, refusal or pending state, real approver, and existing escalation record.
AI SECURITY	NemoClaw checks sensitive data before approved model routing and records the resulting policy decision.	Run paired sets. Direct and indirect injection, sensitive-data egress, memory and tool abuse alongside harmless long-form, multilingual, and quoted content.	Decisions, withheld action, false positives, reason quality, recovery path, audit events, health state, and declared degraded behaviour.
EVIDENCE & POSTURE	Governance surfaces expose model locality, external cost, defensive posture, audit events, tests, and framework mappings.	Select an event after it happens. Reconstruct actor, worker, tool, policy, approval, route, outcome, and time; export it; rerun one mapped control.	User outcome, underlying state, event, export, scope, build, policy version, and rerun agree. Mapping is not presented as certification.
ASSESSMENT & OPERATIONS	An entitled assessment can produce a paced maturity review and roadmap; operators own deployment, monitoring, recovery, and support.	Verify entitlement and disclosure. Complete a bounded assessment; inspect labels on modelled figures; recover one dependency; review operating ownership.	Report and roadmap exist; estimates are labelled; no certification claim; RACI, backup/recovery result, incident route, support scope, and exclusions are recorded.
EVIDENCE RULE		Product documentation establishes the intended claim. The configured tenant test establishes whether this buyer can rely on it. A statement that works only under the vendor's preferred wording is not a held claim.	

BEFORE THE PRODUCT IS OPENED

Write the test script before the vendor opens the demo.

A fair evaluation tells the vendor what problem, boundary, and evidence standard apply. It does not surrender the exact wording, records, users, or failure conditions that would let a prepared path imitate a working control.

TRANSLATE FOUR FAMILIAR MARKETING SENTENCES INTO OBSERVABLE ACCEPTANCE QUESTIONS

<i>"It runs on your infrastructure."</i>	→ Show the live topology, data flows, external dependencies, model routes, and operator boundary.
<i>"AI workers do real work."</i>	→ Produce an accessible artefact, a committed state, correct attribution, and a truthful escalation or approval.
<i>"The AI firewall is on."</i>	→ Run attack and ordinary-work sets, inspect events, then remove one defensive dependency.
<i>"The dashboard shows compliance."</i>	→ Identify assessed scope, policy version, test method, exceptions, evidence, reviewer, and rerun.

BRIEF THE VENDOR WITHOUT SCRIPTING THE ANSWER

STATE IN ADVANCE □ Use case and category: the workflows and whether the product is being assessed as assistant, agent platform, or governed workspace. □ Boundary: permitted data classes, systems, model routes, recipients, prohibited actions, and operating constraints. □ Roles: the human and worker roles, expected decision rights, and approval ownership. □ Evidence standard: artefacts, state changes, event fields, exports, failure visibility, and review method. □ Decision thresholds: usefulness, governance, critical overrides, stop conditions, and retest rules.	KEEP BUYER-CONTROLLED □ Exact phrasing: include ordinary wording, ambiguity, rephrasing, euphemism, and a second language where relevant. □ Records and users: select the employee, document, account, event, or ledger item after the session begins. □ Attack content: disclose categories, not every payload or placement in advance. □ Failure condition: let the buyer choose which dependency is degraded and which evidence must reconcile after recovery. □ Event selection: choose the audit records after the actions occur rather than accepting a prepared example.
--	---

PUT FIVE PERSPECTIVES IN THE ROOM

BUSINESS OWNER Defines useful work, current baseline, acceptable outcome, and adoption cost.	ACTUAL USER Tests ordinary language, exceptions, effort, correction, and workflow fit.	SECURITY / PRIVACY Challenges boundary, access, egress, attacks, evidence, and residual risk.	PLATFORM OPERATOR Owns identity, integrations, monitoring, recovery, support, cost, and change.	EVIDENCE RECORDER Captures tenant, build, event IDs, artefacts, actual state, caveats, and retests.
--	--	---	---	---

RUN THE SESSION IN AN ORDER THAT REVEALS THE SYSTEM, NOT THE THEATRE

01 FREEZE CONTEXT Build, modules, topology, routes, policies, identities, and owners.	02 NORMAL WORK Representative tasks, users, data, exceptions, and baseline measures.	03 DECISION RIGHTS Scope, denial, grant, revoke, approval, escalation, and fresh login.	04 SECURITY Attacks beside harmless work, false positives, reasons, and recovery.	05 FAILURE Degrade one dependency; observe visibility, containment, and restoration.	06 EVIDENCE Select events, reconcile outcomes, export, score, and record the decision.
---	--	---	---	--	--

PREPARED-DEMO RISK	If the vendor cannot change the user, record, wording, route, or failure condition without leaving the demonstrated path, score the evidence as prepared—not as buyer-tenant proof.
---------------------------	---

THE BUYER'S LIVE SESSION

Ten tests separate a deployable product from a persuasive demo.

Run these in the proposed tenant with named reviewers. Choose the records and phrasing yourselves. Preserve screenshots, event IDs, exports, artefacts, and the underlying state; do not rely on a prepared dashboard alone.

- 01

FREEZE THE ACCEPTANCE CONTEXT
Run: record build, modules, topology, identities, connected systems, model routes, policies, support, and exclusions.
Keep: signed configuration inventory and data-flow diagram.
- 02

TRACE BOTH MODEL ROUTES
Run: complete one permitted local turn and one approved hosted turn; observe fallback and cost/locality signals.
Keep: provider, route, data class, telemetry, and dependency evidence.
- 03

PROVE ACCESS TRANSITIONS
Run: deny, grant, fresh-login, use, revoke, fresh-login, disable, and retry for a restricted user.
Keep: IdP state, application result, sessions, and change events.
- 04

DEMAND A WORKER DELIVERABLE
Run: delegate a task under ordinary language, then find the resulting document, task, record, or draft.
Keep: worker attribution, tool events, artefact, access link, and final state.
- 05

PROTECT RESERVED DECISIONS
Run: request a prohibited or consequential action under direct, indirect, and euphemistic wording; then use the approved path.
Keep: refusal or pending state, authorised decision, execution event, and outcome.
- 06

ATTACK THE AI-SECURITY PATH
Run: direct and retrieved injection, sensitive-data egress, unsafe tool parameters, and memory persistence attempts.
Keep: decision, blocked effect, residual risk, health state, and security events.
- 07

PROTECT ORDINARY WORK
Run: harmless long-form, multilingual, quoted-code, people, finance, and document tasks alongside the attack set.
Keep: false positives, reasons, correction path, completion, and repeat result.
- 08

RECONSTRUCT BUYER-SELECTED EVENTS
Run: select one successful, denied, approved, and failed action after completion; export the evidence.
Keep: actor, worker, tool, policy, route, decision, outcome, timestamp, and correlation.
- 09

REMOVE ONE DEPENDENCY
Run: degrade a guard, audit route, model provider, policy pack, evidence sink, or application integration.
Keep: visible failure, fail-closed or bounded mode, alert, retry, recovery, and reconciliation.
- 10

VERIFY ASSESSMENT AND OPERATIONS
Run: confirm entitlement; complete the assessment if in scope; inspect estimate labels; review backup, incident, support, and rollback ownership.
Keep: report, roadmap, disclosures, RACI, recovery result, and accepted limitations.

RECORD THE RESULT WHILE THE EVIDENCE IS OPEN

TEST	VERDICT	EVIDENCE LOCATION	ACTUAL RESULT	REVIEWER	RETEST
—	Hold / weak / fail				
—	Hold / weak / fail				
APPROVE BOUNDED PRODUCTION Usefulness and governance meet threshold; no critical failure; scope, ownership, residual risk, monitoring, rollback, and review date are accepted.		NOT YET The value remains credible, but a named gap, configuration, compensating control, or missing evidence must be resolved before production authority.		STOP The product is unsuitable for the use case, evidence is unreliable, a critical control fails, or the operating model cannot be sustained.	

TRUTH RULE	A success message without an accessible artefact, a matching underlying state change, and a reviewable evidence event is a failed test.
-------------------	--

APPENDIX C · PRIMARY SOURCES, CLAIM NOTES, AND LIMITS

Know what kind of authority you are reading.

A law, a regulator’s statement, a voluntary framework, a professional duty, and a product document answer different questions. A defensible procurement record keeps those roles separate.

THE GOVERNING RULE	No single source establishes that an AI system is lawful, secure, compliant, or suitable. Authority establishes the test; configured-tenant evidence shows what happened.
--------------------	---

FIVE KINDS OF AUTHORITY

Classify the source before relying on the claim.

01	Binding law and regulation Jurisdiction, role, use case, and effective date matter.	Use it to: identify legal duties, prohibited practices, rights, and required controls.	Do not infer: that a product feature or dashboard makes the deploying organisation compliant.
02	Regulator guidance and notices Supervisory interpretation, expectations, and risk focus.	Use it to: understand how existing duties may apply to AI governance and operations.	Do not infer: that every statement creates a new rule or applies to every firm.
03	Public-sector and sector directives Often precise, but bounded by institutional scope.	Use it to: test obligations in the named programme, sector, department, or decision class.	Do not infer: that a federal directive automatically governs provincial, municipal, or private use.
04	Standards, frameworks, and duties Structured practice, management systems, and professional conduct.	Use it to: design governance, controls, review routines, and evidence requirements.	Do not infer: certification, accreditation, assurance, or legal compliance from a mapping.
05	Product documentation and evidence Versioned capability, configuration, tests, and operating records.	Use it to: verify what a specific deployment can do and what a specific event recorded.	Do not infer: universal performance, future capability, or independent validation.

READ EVERY CLAIM THROUGH THREE LENSES

01 FACT What does the primary source actually say, in which jurisdiction, for which role or activity, and from what date?	02 GUIDANCE What evaluation question follows from that source? Guidance translates authority into a test; it does not enlarge the authority.	03 CONFIGURATION What must the buyer reproduce in the candidate tenant to determine whether the promised control exists and remains effective?
--	---	---

PUBLICATION DISCIPLINE	State the authority type, jurisdiction, scope, status, checked date, and the limit of the inference. Where a proposal, revision, or phased implementation is involved, say so plainly.
------------------------	--

CURATED PRIMARY-SOURCE REGISTER

Start with the buyer question, not the vendor claim.

These sources support the evaluation method in this handbook. They establish law, guidance, standards, or professional duties; they do not certify a product or a customer deployment.

Register rule: record what the source establishes, what it does not, and the tenant test that closes the gap

Reviewed through: 25 July 2026

GOVERN AND MANAGE RISK

Use these to structure responsibilities, risk treatment, testing, and management-system evidence.

NIST AI Risk Management Framework 1.0

VOLUNTARY FRAMEWORK · US

Use for: GOVERN, MAP, MEASURE, and MANAGE. NIST says revision is under way, so record the version used.

nist.gov/itl/ai-risk-management-framework

NIST AI 600-1 · Generative AI Profile

COMPANION PROFILE · US

Use for: generative-AI risks and actions that supplement, rather than replace, the core AI RMF.

nist.gov/publications · AI 600-1

ISO/IEC 42001:2023 · Artificial intelligence management systems

INTERNATIONAL STANDARD

Use for: management-system requirements and organisational governance. **Limit:** a control mapping is not a certificate; ISO does not certify organisations, and external certification bodies issue certificates.

iso.org/standard/81230.html

LAW AND FINANCIAL OVERSIGHT

Use complete, current obligations for the organisation's role and use case; do not procure from a slogan about "compliance."

Regulation (EU) 2024/1689 · EU AI Act

BINDING LAW · EU

Use for: role- and risk-based obligations with phased application. Confirm status and applicability at decision time.

eur-lex.europa.eu/eli/reg/2024/1689/oj/eng

FCA · AI update and 2026 supervisory speech

REGULATOR MATERIAL · UK

Use for: outcomes-focused governance, accountability, resilience, third-party mapping, testing, and human oversight.

fca.org.uk · Rethinking regulation for the age of AI

FINRA Regulatory Notice 24-09 · Artificial intelligence

REGULATOR NOTICE · US

Use for: how existing FINRA rules apply when member firms use AI. The notice says it creates no new requirements; evaluate supervision, data, model, communications, and operational risks before deployment.

finra.org/rules-guidance/notices/24-09

REGULATORY-USE RULE

Record the organisation's role, jurisdiction, activity, and implementation date before converting any law or regulator statement into an acceptance criterion.

CURATED PRIMARY-SOURCE REGISTER · CONTINUED

Use the authority that governs the actual work.

Sector, data, and professional duties turn on intended use, affected people, information class, and institutional scope. Similar-looking AI workflows can therefore require different evidence.

Scope rule: classify the data, decision, profession, and public-sector boundary	Reviewed through: 25 July 2026
---	--------------------------------

DATA, SECTOR, AND PROFESSIONAL DUTIES

Use the source that governs the data, decision, organisation, and profession involved.

HHS · HIPAA Security Rule CURRENT RULE · US HEALTHCARE Use for: safeguards for electronic protected health information. HHS says the current rule remains in effect while proposed changes are pending. hhs.gov/hipaa/for-professionals/security	FDA · Clinical Decision Support Software GUIDANCE · US HEALTHCARE Use for: intended-use analysis when clinical decision support may fall within device regulation. fda.gov · Clinical Decision Support Software
ICO · AI and data protection REGULATOR GUIDANCE · UK Use for: applying UK GDPR principles and assessing rights-and-freedoms risks. The ICO marks the guidance as under review following legislative change. ico.org.uk · Artificial intelligence guidance	Canada · Directive on Automated Decision-Making FEDERAL DIRECTIVE · CANADA Use for: automated administrative decisions within its federal-government scope. Do not extend it automatically to provinces, municipalities, or private organisations. tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592
ABA Formal Opinion 512 · Generative Artificial Intelligence Tools PROFESSIONAL ETHICS · US LEGAL Use for: competence, confidentiality, communication, supervision, candour, and reasonable-fee duties in legal practice. It is professional guidance, not a product approval. americanbar.org · Formal Opinion 512	

TECHNICAL FOUNDATION Product documentation establishes design intent and named capabilities; only a configured test establishes deployment behaviour.

NVIDIA NemoClaw · Official overview PRODUCT DOCUMENTATION Use for: the published open-source reference stack, OpenShell isolation, model routing, policy, and lifecycle concepts. Vendors should identify their own extensions; use of the stack does not imply NVIDIA certification or endorsement. docs.nvidia.com/nemoclaim/latest/about/overview.html

CONFIGURED-PROOF RULE	Product documentation can establish design intent and named capability. Only a versioned, configured-tenant test can establish what the proposed deployment did under the buyer's policy, identity, model route, and integration boundary.
------------------------------	--

THE PUBLICATION CONTROL

Say exactly what the evidence establishes.

A precise limit strengthens a procurement decision. It tells the committee where independent advice, additional testing, or contractual evidence is still required.

WHAT THIS HANDBOOK DOES NOT ESTABLISH

These limits apply to every score, test, illustration, and product example.

- ✗ **Not legal, regulatory, medical, investment, or assurance advice**
Qualified advisers must assess the organisation, jurisdiction, decision, data, and facts involved.
- ✗ **Not certification or customer compliance**
A framework mapping, test result, or dashboard does not certify a product, supplier, or deploying organisation.
- ✗ **Not an independent audit**
Vendor evidence and buyer-run tests support evaluation; they do not replace independent assurance where it is required.
- ✗ **Not a security guarantee**
Controls reduce specified risks under stated conditions. They are not infallible and must be monitored, tested, and updated.
- ✗ **Not proof of local processing for an external route**
Self-hosted orchestration does not make an externally hosted model local. Record the actual route, data sent, provider, and retention terms.
- ✗ **Not a frozen capability or contractual commitment**
Capabilities depend on version, configuration, integrations, permissions, and deployment topology. Contractual promises belong in the contract.

THE CLAIM-MAINTENANCE RECORD

Every material statement should remain traceable after the product, law, or deployment changes.

01	Name the claim and owner Write the exact sentence, the accountable reviewer, and the audience that will rely on it.
02	Attach the primary authority Record source type, jurisdiction, scope, status, effective date, canonical URL, and checked date.
03	Bind it to configuration Record build, policy set, model route, integration, permissions, and deployment boundary.
04	Attach reproducible evidence Keep the test, timestamp, actor, result, audit record, and exception—not only the screenshot.
05	State the inference limit Say what the evidence proves, what it does not prove, and which reviewer must close the gap.
06	Set the review trigger Use event-driven review as well as calendar review. Material change should reopen the claim.

REOPEN THE RECORD WHEN

- law or regulator guidance changes
- the product build changes materially
- permissions or integrations change
- a standard or framework is revised
- model provider or route changes
- a production test contradicts the claim

THE COMMITTEE RULE

If a material claim cannot be linked to current authority and reproduced in the proposed deployment, record it as unverified—not as a benefit.