



NORTHERN LOGIC AI

MANAGED AI OPERATIONS FIELD GUIDE

Human-Led, AI-Assisted: An Operating Model for Specialized AI Assistants

A practical operating model for adding preparation, coordination, monitoring, and knowledge capacity through specialized AI assistants under clear human ownership.

ASSESS

Find the value

BUILD

Create the system

OPERATE

Keep it useful

Prepared by Logan Abell · Northern Logic
Published July 28, 2026 · northernlogic.ai

In this field guide

Use this paper as a working reference for business decisions. The architecture is intentionally nontechnical: it names the goals, knowledge, responsibilities, controls, and operating practices that should exist around the technology.

01**The operating model at a glance****02****What current research suggests about work design****03****Capacity should be defined before the role****04****The task allocation matrix****05****Design every role on one page****06****Shared knowledge is the team's foundation****07****A portfolio of specialized roles****08****Example: a human-led client delivery system****09****Human roles and decision rights****10****Five levels of action****11****The Managed AI Operations hub****12****Operating cadence****13****Measurement without invented precision****14****Introduce the operating model with the team****15****Common operating failures****16****The Assess, Build, Operate sequence****17****Owner's readiness checklist**

The central idea: company knowledge gives the system context, business goals give it direction, people keep consequential decisions, and managed operations keeps the whole system visible and improving.

This publication provides general business and design guidance. It is not legal, financial, clinical, employment, or security advice for a specific organization.

A specialized AI assistant adds capacity when it owns a defined preparation, coordination, monitoring, or knowledge responsibility inside a human-led process. People continue to own business goals, judgment, relationships, exceptions, and consequential decisions. The operating model connects those responsibilities so the system helps the team without creating a new management burden.

For this paper, an **AI agent** means software that can use an AI model to plan steps and work with connected tools. From here on, we use **specialized AI assistant** because it describes the practical business design: a bounded role with approved knowledge, tools, review, measurement, and ongoing support.

The model has four commitments:

1. **Human-led:** A named person owns the result and decision rights.
2. **Role-specific:** Each assistant has a narrow, reviewable responsibility.
3. **System-managed:** Goals, work, knowledge, permissions, cost, and changes are visible in one operating layer.
4. **Evidence-improved:** Expansion follows reviewed value and quality, not feature availability.

01 The operating model at a glance

Element	Purpose	Owner question
Business goal	Gives the system direction	What result matters now?
Role portfolio	Divides preparation and coordination into bounded responsibilities	Which role should handle each part?
Shared knowledge	Gives roles approved company context	What may the system know and trust?
Tools and flow	Moves eligible work through existing systems	Where does work begin, happen, and end?
Human control	Keeps judgment and consequence with qualified people	Which actions require review or approval?
Operating hub	Makes goals, tasks, cost, issues, and changes visible	Can the owner see and manage the work?
Measurement	Connects system activity to business value	Is this creating useful capacity?

The flow is:

Business goal → eligible task → specialized AI assistant → approved knowledge and tools → human review → useful result → measurement and correction

Several assistants may contribute to one result, but they should not be added merely to make the system appear sophisticated. Separate roles only when the responsibility, knowledge, permission, or review standard is meaningfully different.

02 What current research suggests about work design

The [International Labour Organization's 2025 occupational exposure index](#) analyzed tasks within occupations and concluded that job transformation is the most likely overall effect because most occupations contain work requiring human input. It is a global research finding, not a prediction for any one company. Its practical lesson is to redesign tasks with the people who understand the work.

The U.S. Census Bureau's [2026 AI diffusion working paper](#) reported AI use across business functions and worker tasks, with 57% of adopting firms using AI in three or fewer business functions during the supplement reference period. That supports a focused operating approach, although each company still needs its own assessment.

The [June 2026 Anthropic Economic Index](#) introduced methods for studying longer-running AI tasks and distinguished several collaboration patterns, including directive work, feedback loops, task iteration, learning, and validation. This is data from one provider's products, not a general workforce benchmark. It illustrates a useful design point: the same underlying AI can support different levels of delegation and collaboration depending on the role, interface, and controls around it.

03 Capacity should be defined before the role

Do not begin by naming an assistant. Begin by naming the capacity constraint.

Preparation capacity

The team spends time gathering records, comparing sources, checking required fields, and formatting a result before judgment can begin.

Potential role: prepare a decision-ready account, project, document, or operations brief.

Coordination capacity

Work waits because inputs, handoffs, owners, and exceptions are spread across tools or conversations.

Potential role: check readiness, prepare handoffs, route missing information, and keep the next action visible.

Monitoring capacity

A person repeatedly checks the same conditions across records to find what needs attention.

Potential role: apply approved thresholds or checklists and prepare an exception queue for an owner.

Knowledge capacity

People repeatedly search for current company guidance or ask the same experts for context.

Potential role: find approved sources, return a concise answer with references, and route conflicts or sensitive questions.

Coverage capacity

Routine preparation should happen on a schedule or when a defined event occurs, including times when the responsible person is focused elsewhere.

Potential role: prepare recurring work within approved hours, volume, cost, and action boundaries.

Capacity only creates value when the business can use it. The team should be able to say where returned attention goes, such as serving clients, reducing a backlog, improving review quality, or removing dependence on a scarce owner.

The article [How Specialized AI Assistants Add Capacity](#) provides a shorter readiness scorecard.

04 The task allocation matrix

Assign tasks according to evidence, rules, and consequence, not according to a whole job title.

Work category	Specialized AI assistant can lead	Person should lead
Gathering	Collect approved records and source links	Decide which sources are authoritative
Preparation	Summarize, compare, classify, and draft	Judge usefulness and material accuracy
Coordination	Check status, prepare handoffs, route known exceptions	Resolve ownership conflicts and unusual cases
Monitoring	Apply defined rules and surface changes	Interpret impact and choose the response
Communication	Prepare messages from approved context	Own tone, commitments, and sensitive conversations
Decisions	Organize options, evidence, and unresolved questions	Make financial, legal, clinical, access, employment, safety, and relationship decisions
Improvement	Group corrections and recurring issue patterns	Change policy, process, priorities, or boundaries

The dividing line can move after evidence, but every movement should be a new operating decision with testing and review.

05 Design every role on one page

A role card prevents vague responsibility and permission sprawl.

Mission

State the useful result in one sentence:

Prepare a weekly client delivery readiness brief from approved project, service, and communication records for the delivery leader to review.

Eligible work

Define the standard cases the role accepts and the cases it must refuse or escalate.

Trigger

Name the schedule, business event, queue entry, status change, or approved request that begins the work.

Approved inputs

List the records, procedures, examples, decisions, and templates the role may use. Link each durable source to an owner and current version.

Tools and permissions

Separate search, read, draft, create, change, send, approve, and delete capabilities. Grant only the functions required by the role.

Output

Specify the format, required fields, source references, delivery location, and definition of complete.

Human review

Name the reviewer, what they must verify, and what happens after approval or rejection.

Stop conditions

Define missing inputs, conflicting sources, sensitive content, high-impact requests, tool failures, or confidence limits that should pause the task.

Measures

Choose business value, quality, adoption, review burden, cost, and control signals.

Owner

Name the person who can change the role, accept risk, and decide whether it continues.

The role card is not a prompt. It is the business agreement that guides the technical implementation and ongoing operation.

06 Shared knowledge is the team's foundation

Specialized AI assistants should use approved company knowledge rather than separate copies of procedures and examples. A shared foundation includes:

- Current definitions and business language.
- Procedures and checklists.
- Policies and constraints.
- Approved examples and templates.

- Decisions and known exceptions.
- References to current operational records.
- Owners, effective dates, access classes, and review dates.

The assistant's temporary working context should not become durable company guidance without approval. Likewise, a model-generated answer should not be treated as policy because it sounded correct.

The whitepaper [From Scattered Information to Shared Business Memory](#) provides the full architecture. Northern Logic's [Business Knowledge Base service](#) builds and maintains this foundation around the work that needs it.

07 A portfolio of specialized roles

Start with one role. As the system proves useful, adjacent roles can share knowledge and coordinate under the same operating model.

Context preparation role

Gathers the records and approved knowledge needed for a decision or review.

Typical output: a source-linked context package.

Boundary: reads approved information but does not interpret consequential meaning.

Work preparation role

Turns context into a structured draft, brief, checklist result, or proposed next action.

Typical output: a review-ready work product.

Boundary: prepares but does not make the final commitment.

Quality check role

Compares the prepared result with approved requirements, examples, and source evidence.

Typical output: completeness notes and exceptions.

Boundary: identifies possible gaps but does not substitute for qualified professional judgment.

Coordination role

Tracks status, required inputs, owners, approvals, and handoffs.

Typical output: visible task status and routed exceptions.

Boundary: coordinates known rules but does not resolve business conflicts.

Knowledge maintenance role

Groups unanswered questions, source conflicts, correction patterns, and review dates for knowledge owners.

Typical output: a maintenance queue with evidence.

Boundary: proposes updates but does not approve company guidance.

These are patterns, not required job titles. A simple workflow may need only one role. The goal is clarity, not role count.

08 Example: a human-led client delivery system

Business goal

Give the delivery leader a complete, current view of project readiness before the weekly operating review.

Role sequence

Step	Specialized role	Work	Human responsibility
1	Context preparation	Gather scope, milestones, owners, recent updates, open issues	Define approved sources and eligible projects
2	Knowledge guidance	Retrieve current delivery standards and known exceptions	Approve standards and resolve conflicts
3	Readiness preparation	Compare current context with requirements and draft the brief	Judge risk, priority, staffing, and tradeoffs
4	Coordination	Route missing inputs and unresolved decisions	Assign owners and decide the next action

Tools and access

The roles may need read access to the project system, approved document area, and issue records. They may create a draft brief and internal follow-up tasks. Customer messages, scope changes, staffing decisions, financial commitments, and access changes remain with people.

Shared knowledge

- Delivery status definitions.
- Readiness checklist.
- Escalation rules.
- Approved project template.
- Known exception decisions.
- Examples of strong readiness briefs.

Useful result

One source-linked review that identifies what is ready, what is missing, which decisions are open, and who owns each next step.

This is more than a single summary feature. It is an operated capacity system that combines several tools, shared company knowledge, specialized preparation, and human decisions.

09 Human roles and decision rights

Human role	Decision rights
Business sponsor	Sets the goal, priority, investment, and continuation decision
Process owner	Defines eligible work, standards, exceptions, and fallback
Reviewer	Accepts, corrects, or rejects prepared work
Knowledge owner	Approves sources, meaning, and company guidance
Managed operator	Monitors tasks, issues, costs, tests, and changes
Technical or security adviser	Reviews access, vendors, data flow, and response needs

In a smaller company, one person may hold several roles. The decisions should still be explicit.

Avoid review theater. A human checkpoint is useful only when the reviewer has the skill, time, source context, and authority to judge the work. The assistant should present evidence and uncertainty in a format that supports that review.

10 Five levels of action

Level	Assistant activity	Human control
1. Observe	Read approved information	Source and access review
2. Prepare	Summarize, compare, classify, draft	Reviewer checks the result
3. Propose	Recommend a next action	Named owner approves
4. Execute reversible work	Create a task or update a low-impact field	Logging, limits, and review
5. Consequential action	Send, pay, commit, grant access, approve, or delete	Qualified approval outside the model

The [OWASP guidance on excessive agency](#) recommends limiting tool functionality, permission, and autonomy and requiring human approval for high-impact actions. A first implementation should usually concentrate on levels one through three, where useful preparation can be proven before broader action is considered.

The guide to [human-in-the-loop AI decisions](#) helps place approval based on consequence.

11 The Managed AI Operations hub

Several assistants should not create several disconnected management systems. The operating hub should provide one view of:

Direction

- Current business goals.
- Prioritized workflows and roles.

- Owners and review responsibilities.

Work

- Queued, active, completed, failed, and paused tasks.
- Schedules and event triggers.
- Exceptions and approval status.

Knowledge and tools

- Approved source collections.
- Tool identities and permission scopes.
- Changed, stale, or unavailable dependencies.

Control

- Volume, cost, and runtime limits.
- Approval rules and pause conditions.
- Issue, incident, and change history.

Value

- Accepted work.
- Handling and wait time.
- Quality and correction patterns.
- Adoption and review burden.
- Operating cost and current improvement.

This operating layer is the connection between the business and the technical components. The [Managed AI Operations Blueprint](#) shows how goals, knowledge, roles, tools, controls, and measurement fit together. The [managed AI operations service](#) provides the ongoing responsibility to keep that system useful.

12 Operating cadence

Per task or daily

- Failed and blocked work.
- Urgent exceptions.
- Pending approvals.
- Volume or spending limits.
- Access or tool errors.

Weekly

- Completed and accepted work.

- Material corrections.
- Unresolved exceptions.
- Review time and adoption.
- Open issues and the next improvement.

Monthly

- Business value against the baseline.
- Cost per accepted unit.
- Knowledge and tool changes.
- Role fit and task eligibility.
- Permission and approval review.
- Test updates and recurring failure patterns.

Quarterly or at a material change

- Continue the role as designed.
- Expand one adjacent responsibility.
- Restrict access, cases, or actions.
- Redesign the process or knowledge.
- Retire the capability.

The cadence should match consequence and volume. The important point is that someone reviews the live system after launch.

13 Measurement without invented precision

Measure one business result, then the operating signals required to interpret it.

Measure group	Questions
Business result	Did the system produce the review-ready unit the process needs?
Capacity	Did handling time, wait time, backlog, or owner attention change?
Quality	What was accepted, corrected, rejected, or escalated?
Adoption	Did eligible work use the system, and did the team find it useful?
Cost	What did building, usage, review, maintenance, and support cost?
Control	Were permissions, approvals, limits, and fallback used as designed?
Improvement	Which recurring issue was corrected, and did the retest pass?

Do not count every minute returned as cash. Identify how the company will use the capacity. Do not assign revenue to a prepared brief unless the business can establish the connection. Do not turn a high number of outputs into a success measure if people do not use them.

The [AI ROI guide](#) provides the full value model and a seven-factor opportunity scorecard.

14 Introduce the operating model with the team

The people closest to the work should help shape the role.

Before build

- Walk through real examples.
- Ask where time and attention are lost.
- Separate preparation from judgment.
- Identify trusted sources and known exceptions.
- Agree on what good output looks like.

During shadow mode

- Run the assistant beside the current process.
- Let reviewers explain correction reasons.
- Capture missing context and edge cases.
- Show the team what the system did and which sources it used.

At controlled launch

- Explain eligible work and exclusions.
- Show where results and approvals appear.
- Provide a correction and concern path.
- Name the process, knowledge, and operating owners.
- Keep the fallback process available.

After launch

- Share improvements and unresolved limitations.
- Review whether the capacity is reaching the intended work.
- Update the role when business goals or processes change.

Use language that reflects support and accountability. The assistant prepares defined work. The team owns the business.

15 Common operating failures

Broad role, unclear result

"Help operations" cannot be tested or managed. Rewrite the role around one reviewable output.

Assistant sprawl

Several tools and roles appear without shared goals, knowledge, or ownership. Use one portfolio and operating hub.

Permission sprawl

Technical connectivity is mistaken for business authorization. Separate capabilities and grant the minimum required for the role.

Hidden work

Owners cannot see what is queued, completed, changed, or failed. Make status, sources, actions, exceptions, and cost visible.

Knowledge drift

Each role has a different copy of company instructions. Build shared maintained sources with owners and review dates.

Approval without evidence

Reviewers receive a polished result but no sources, uncertainty, or exception context. Design the output for judgment.

No fallback

The process depends on a connection or model with no manual route. Define pause, recovery, and continuity before launch.

No operator after build

The system has no owner for monitoring, correction, tool changes, access, or value review. Assign ongoing responsibility before it enters live work.

16 The Assess, Build, Operate sequence

Assess

- Name the business result and capacity constraint.
- Map the current workflow, sources, tools, and decisions.
- Baseline handling time, wait, quality, exceptions, and owner attention.
- Compare process improvement, product configuration, automation, specialized assistance, and deferral.
- Select one bounded first role and accountable owner.

Build

- Write the role card and eligible scope.
- Prepare approved company knowledge.
- Connect only required tool functions.
- Define human review, permissions, limits, and fallback.
- Build representative tests.
- Run offline and in shadow mode before controlled live use.

Operate

- Manage goals, tasks, schedules, cost, and approvals.
- Review accepted work, corrections, exceptions, and adoption.
- Maintain knowledge, tools, tests, and permissions.
- Keep issue and change history.
- Measure business value and full operating cost.
- Expand, restrict, redesign, or retire from evidence.

The [90-day AI implementation roadmap](#) turns this sequence into practical gates and deliverables.

17 Owner's readiness checklist

Before launching a specialized AI assistant, confirm:

- One business result and owner are named.
- The returned capacity has a planned use.
- The role fits on one page.
- Eligible and excluded work are clear.
- Approved knowledge sources and owners exist.
- Search, read, draft, change, send, approve, and delete are separated.
- Qualified people own consequential decisions.
- The output includes enough evidence to review.
- Normal, unusual, and hostile cases were tested.
- Tasks, exceptions, costs, and changes are visible.
- The team knows how to correct or challenge the work.
- A fallback and pause path exist.
- Someone is accountable for ongoing operation.

Northern Logic helps owners identify the right capacity constraint, build the specialized role and shared system, and manage it after launch. Explore the [specialized AI assistant service](#) or [book a free 15-minute AI assessment](#) with one workflow, bottleneck, or business goal.

18 Sources

Primary and authoritative references used throughout this field guide.

- [International Labour Organization, Generative AI and Jobs: A Refined Global Index of Occupational Exposure](#)
- [U.S. Census Bureau, The Microstructure of AI Diffusion](#)
- [Anthropic Economic Index, June 2026 report](#)
- [NIST AI Risk Management Framework](#)
- [NIST AI RMF Core](#)
- [NIST Generative AI Profile](#)
- [ISO/IEC 42001:2023 AI Management Systems](#)
- [OWASP LLM06:2025 Excessive Agency](#)
- [CISA and NCSC Guidelines for Secure AI System Development](#)