

AGENT-AGNOSTIC EDITION

The Augmented Leader Playbook



Build governed AI workflows with source discipline, memory, and approval gates.

The canonical H-EXEC operating model for governed AI workflows. Vendor-neutral by design: it works with any approved assistant, and keeps the leader accountable for every decision.

CONTENTS

What's inside

The canonical, vendor-neutral H-EXEC operating model: the governance frame, the six techniques, and a realistic first thirty days for a healthcare executive.

Executive summary	03
-------------------------	----

01 The Operating Model

1.1	What this playbook is for	05
1.2	What you will learn to build	06
1.3	Healthcare executive use cases	07
1.4	The H-EXEC operating model	08
1.5	Canonical reference architecture	09
1.6	Healthcare AI data lanes	10
1.7	Governance roles & boundaries	11

02 The Six Techniques

01	Context File	13
02	Source Grounding	14
03	Persistent Memory	15
04	Workflow Library	16
05	Guardrails & Approval Gates	17
06	Staged Automation	18

03 In Practice

3.1	Worked example: daily operating brief	20
3.2	Executive checklists	21
3.3	Your first 30 days	22

EXECUTIVE SUMMARY

A Governed Operating Model

This playbook gives the healthcare executive a way to use approved AI tools inside governed workflows that preserve context, ground outputs in named sources, and stage work for human review. Throughout, the leader stays accountable for every final decision.

WHAT YOU GET

- **An operating profile.** Your Executive AI Operating Profile with starter CONTEXT, RULES, VOICE, and MEMORY, plus a SOURCE_INDEX of approved sources, owners, data classes, and currency.
- **Workflow SOPs.** Two or three repeatable procedures such as a daily operating brief, an executive SBAR, or a vendor decision analysis.
- **A review gate.** A control that keeps drafts staged until the accountable owner approves, revises, or rejects them.

YOUR FIRST 30 DAYS

- **Set the profile.** In week one, stand up CONTEXT and RULES and choose the data lane so your boundaries are explicit from the start.
- **Make boundaries clear.** Define what the assistant owns and what it must never touch before any live work begins.
- **Run one artifact.** In week two, draft a first daily brief or SBAR and review it, so meeting prep gets faster.

YOUR BOUNDARIES

- **Human approval gate.** Nothing is sent, filed, posted, or acted on without explicit approval from the accountable owner.
- **Approved sources only.** Every source must be approved for its tool, workspace, data class, use case, access, and retention.
- **No sensitive data.** No PHI, patient-level, or other sensitive data enters a workflow unless the organization explicitly authorizes it.

PART ONE

01

The Operating Model

Use approved AI tools inside governed workflows that preserve context, ground outputs in sources, stage work for human review, and keep the leader accountable for final decisions.

THE H-EXEC PROMISE

1.1 THE OPERATING MODEL

What this playbook is for

This playbook teaches the H-EXEC method without assuming a specific AI product. It is for healthcare executives who want to use approved AI tools more safely, consistently, and accountably across recurring leadership work.

It is vendor-neutral: it works with approved AI tools and internal assistants. The promise is simple:

THE PROMISE

Use approved AI tools inside governed workflows that preserve context, ground outputs in sources, stage work for human review, and keep the leader accountable for final decisions.

WHAT THIS IS NOT

This is not autonomous AI, clinical decision support, compliance or legal automation, procurement advice, or a replacement for approved enterprise systems.

1.2 THE OPERATING MODEL

What you will learn to build

Use the playbook to design the operating pieces behind governed AI work:

- An Executive AI Operating Profile.
- A starter **CONTEXT**, **RULES**, **VOICE**, and **MEMORY** set.
- A **SOURCE_INDEX** that identifies approved sources, source owners, data classes, and currency.
- Two or three workflow SOPs selected from Daily Operating Brief, Executive SBAR, Email Draft Queue, Vendor Decision Analysis, Meeting Prep Brief, or Project Monitor.
- A review gate that keeps drafts staged until the accountable human owner approves, revises, or rejects them.

1.3 THE OPERATING MODEL

Healthcare executive use cases

ROLE	MONDAY-MORNING USE CASE	OUTCOME
Service-line leader	Prepare a daily operating brief across access, denials, staffing, quality, budget, vendor, and project issues	Fewer stale decisions; meeting prep arrives current and sourced.
COO / VP Operations	Convert messy operational notes into an executive SBAR or escalation memo	Messy notes become a clear executive SBAR that speeds decision communication.
Ancillary / nursing / procedural leader	Summarize throughput, staffing constraints, patient-flow risks, and follow-ups from sanitized sources	Sharper throughput and staffing priorities, no patient-level exposure.
Finance / operations partner	Compare vendor options, assumptions, TCO, implementation lift, and decision risks	Vendor and capital choices rest on compared, source-grounded assumptions.
Quality / regulatory leader	Track survey-readiness workstreams, evidence gaps, owners, and stop conditions	Survey-readiness workstreams carry named owners, evidence gaps, and stop conditions.

1.4 THE OPERATING MODEL

The H-EXEC operating model

H-EXEC turns recurring executive work into a governed workflow system. The leader defines the work, the sources, the rules, the output standard, and the approval gate. The AI assistant helps synthesize, draft, compare, and refresh memory inside those boundaries.

THE REPEATABLE LOOP

- 1 **Frame the work.** Name the decision, deliverable, meeting, risk, or follow-up.
- 2 **Name the sources.** Identify approved files, notes, metrics, emails, or public references.
- 3 **Set the rules.** Define audience, tone, format, sensitive-data boundary, and approval requirements.
- 4 **Run the workflow.** Ask the assistant to synthesize, compare, draft, or prepare.
- 5 **Inspect the output.** Check facts, source fit, assumptions, tone, and missing context.
- 6 **Approve or revise.** The leader decides what is sent, filed, shared, or killed.
- 7 **Refresh memory.** Update current state so the next run starts from truth.

1.5 THE OPERATING MODEL

Canonical reference architecture

COMPONENT	PURPOSE	REQUIRED	UPDATED
CONTEXT	Stable role, scope, stakeholders, work style, operating boundaries	Yes	Monthly
RULES	Source, privacy, approval, tool-lane, and outbound communication rules	Yes	When rules change
VOICE	Writing style, tone rules, phrases to use or avoid	Yes	Quarterly
MEMORY	Current active threads, open decisions, blockers, recent changes	Yes	Every session
LOG	Historical status notes and completed decisions	Yes	As memory rolls forward
SOURCE_INDEX	Approved source list, owner, date, data class, and currency	Yes	Before each workflow
WORKFLOWS/	One SOP per recurring executive workflow	Yes	When a workflow changes
OUTPUTS/	Draft artifacts staged for review	Yes	Every workflow
REVIEW/	Approvals, open questions, blocked items, escalation notes	Yes	Every workflow

The exact workspace may vary by organization and approved tool, but these components should remain recognizable in any H-EXEC implementation.

1.6 THE OPERATING MODEL

Healthcare AI data lanes

LANE	EXAMPLES	DEFAULT H-EXEC RULE
Public / sanitized	Sanitized practice packet, public policy, anonymized examples	Allowed for approved training and shared examples.
Internal non-PHI	Sanitized project status, de-identified metrics, non-sensitive meeting themes	Allowed only in approved tools and workspaces.
Confidential business	Vendor pricing, strategy, contracts, budget assumptions, nonpublic financials	Requires explicit tool approval, access control, retention rules, and owner approval.
Regulated / PHI / patient-level	Patient identifiers, clinical details, patient-level schedules, protected health information	Prohibited unless the organization explicitly authorizes that exact tool, data class, use case, access model, retention model, and BAA/contractual controls where required.
Privileged / prohibited	Legal advice, labor relations, investigations, credentials, security incidents, passwords, internal URLs, high-risk clinical decisions	Do not use in shared examples; escalate to the appropriate owner before any AI workflow.

SHORT VERSION

Short version: a user-provided file does not override policy. A source counts as approved only when its tool, workspace, data class, use case, access, and retention are all cleared. The handling rules for each lane are stated once under Guardrails & Approval Gates.

1.7 THE OPERATING MODEL

Governance roles & boundaries

ROLE	OWNS
Executive owner	Final judgment, approval, decision, and accountability
Operational owner	Source accuracy, workflow usefulness, and owner/action validation
IT / security	Tool approval, access controls, connectors, storage, retention, and auditability
Compliance / privacy	PHI, privacy rules, minimum-necessary standard, and policy fit
Legal / procurement	Privilege, contract terms, vendor obligations, BAAs, and procurement constraints
Finance	Financial assumptions, budget impact, and TCO logic
Implementation guide / coach	Workflow education, sanitized examples, and stop-condition discipline

WHAT H-EXEC DOES NOT DO

Make clinical decisions, or provide legal, compliance, privacy, cybersecurity, procurement, HR, clinical, or financial advice.

Send patient, employee, vendor, or executive communications without approval.

Update systems of record without human review and proper authorization.

Determine compliance status, or replace employer policy, approved systems, access controls, or data-handling rules.

PART TWO

02

The Six Techniques

Each technique adds one form of discipline: context, sources, memory, repeatable procedure, guardrails, and, only last, automation.

HOW THE SYSTEM COMPOUNDS

Technique 01 / Six



The Context File

Front-load who the leader is, so the assistant never produces generic output.

A Context File tells the assistant who the leader is, what they own, what systems they operate in, which stakeholders matter, and what rules govern the work.

■ FILLED STARTER · CONTEXT

WHO I AM: Service-line executive responsible for [service lines/functions].

WHAT I OWN: Access, staffing, quality, budget, vendor follow-up, and project execution for [scope].

KEY STAKEHOLDERS: COO/VP Ops, finance partner, quality partner, department managers, vendor contacts.

ACTIVE PRIORITIES: [Project] · [current status] · [owner] · [next decision].

APPROVED SOURCES: Sanitized project notes, de-identified dashboard exports, approved meeting summaries.

HARD BOUNDARIES: No PHI, patient-level details, employee relations details, privileged legal content, credentials, confidential vendor terms, or nonpublic financials unless explicitly approved for this tool and use case.

WATCH OUT FOR

Vague ownership. A context file that does not state who the leader is, what they own, and which systems and stakeholders matter produces generic output; name them explicitly.

Stale priorities. Active priorities and key stakeholders drift over time, so a context file left unrevised will steer work toward old problems; refresh it on a set cadence.

Missing boundaries. Without explicit hard boundaries (no PHI, no patient or employee-relations detail, no privileged matter, no credentials, no vendor terms, no nonpublic financials), the assistant has no line it knows not to cross; state every prohibited data class up front.

Technique 02 / Six



Source Grounding

The assistant should not guess; it should be pointed to named sources.

Point the assistant to named sources and ask it to distinguish facts, assumptions, gaps, and recommendations. Source grounding is what keeps an executive artifact defensible.

■ USE THIS PATTERN · SOURCE GROUNDING

Use only the sources I provide or explicitly approve for this workflow.

For each material claim, name the source, date, owner if known, and whether it appears current.

Separate source-backed facts from assumptions.

Flag missing information.

Do not invent metrics, dates, owners, decisions, or commitments.

WATCH OUT FOR

Silent invention. When a metric, date, owner, or commitment is absent, the assistant may fill the gap with a plausible figure; require it to flag missing information rather than supply it.

Unapproved inputs. A user-provided file does not override policy, so treating any uploaded document as a valid source bypasses your controls; accept only sources approved for tool, workspace, data class, use case, access, and retention.

Blurred facts and assumptions. Mixing confirmed facts with inference hides where the reasoning rests; separate the two and name source, date, owner, and currency for each material claim.

Technique 03 / Six



Persistent Memory

Keep the workflow from starting over every time, without storing sensitive records.

The **MEMORY** file stays short, redacted, and current, while older notes move into **LOG**.

- Store minimum-necessary operational summaries, not raw records.
- Use source IDs or redacted references for sensitive topics.
- The data classes that may never be stored unless explicitly authorized are the same restricted classes defined under **Guardrails & Approval Gates**.
- Proposed memory changes should be reviewed before persistence.

■ FILLED STARTER · MEMORY

LAST UPDATED: 2026-08-11.

ACTIVE THREADS: Access dashboard refresh pending from finance partner; vendor implementation timeline needs source confirmation.

OPEN DECISIONS: Whether to escalate staffing coverage issue to VP Ops after sanitized source packet is reviewed.

WAITING ON: Updated metric definition for turnaround-time report.

DO NOT REPEAT: Do not reuse old timeline assumptions from a prior packet.

WATCH OUT FOR

Sensitive extracts. Persisting PHI, patient-level detail, employee-relations notes, legal advice, credentials, internal URLs, or vendor terms into memory turns a convenience into a liability; store only minimum-necessary operational summaries, using source IDs or redacted references for sensitive topics.

Memory bloat. Memory that keeps growing goes stale and crowds out current state; keep it short and move older notes to the log.

Unreviewed writes. Memory updated automatically can quietly persist something it should not; route proposed memory changes through review before they are saved.

Technique 04 / Six



Workflow Library

Any task that should run the same way twice deserves an SOP.

Good first workflows: Daily Operating Brief, Executive SBAR, Email Draft Queue, Vendor Decision Analysis, Project Monitor, and Meeting Prep Brief.

■ WORKFLOW SOP STARTER · WORKFLOWS/

PURPOSE: Prepare a source-grounded executive artifact for [workflow].

TRIGGER: Use when [event/cadence].

INPUTS: Approved sources, source dates, source owners, current MEMORY, applicable RULES.

OUTPUT: Draft artifact staged in OUTPUTS with approval status.

STOP CONDITIONS: Missing source, suspected PHI, patient safety concern, employee or labor relations, legal privilege or legal exposure, contracting terms, cybersecurity incident, regulatory finding, clinical decision support, media/reputation risk, or conflicting data.

VERIFICATION: Check source names, dates, owners, metric definitions, assumptions, and approval gate.

APPROVAL: Human owner approves, revises, or kills output.

WATCH OUT FOR

SOPs without stop conditions or verification. A workflow that names no point at which to halt (missing source, suspected PHI, patient safety, legal privilege) and no verification step will produce confident, unchecked output, so every SOP must specify both before it runs.

Over-documenting one-offs. Building a formal SOP for work that runs only once wastes effort and clutters the library, so reserve SOPs for tasks that recur the same way twice or more.

Missing approval ownership. An SOP that omits the accountable human approver lets drafts move forward without a clear owner, so every workflow must name who approves, revises, or rejects the output.

Technique 05 / Six



Guardrails & Approval Gates

This is the safety center of H-EXEC.

CORE RULES

No PHI, patient data, employee records, confidential vendor terms, privileged material, credentials, internal URLs, or nonpublic financial details in shared examples.

No PHI or patient-level data in any AI workflow unless explicitly authorized by the organization for that exact tool, data class, use case, access model, retention model, and contractual controls.

Nothing is sent, filed, posted, or acted on without explicit human approval.

Drafts are staged for review.

Unsupported claims are labeled or removed.

Sensitive work stays in approved systems only.

WATCH OUT FOR

Soft or implicit approval gates. Treating silence or a quick glance as approval lets unreviewed drafts get sent, filed, or acted on, so require explicit human sign-off before anything leaves the staging area.

Sensitive data in the wrong system. Pulling confidential business or regulated content into an unapproved tool or workspace breaks the data-lane controls, so keep sensitive work inside systems approved for that exact tool, data class, use case, access model, and retention.

Unsupported claims passed through. Letting assertions with no named source reach a final draft erodes trust in the output, so label or remove any claim that is not grounded in an approved source.

Technique 06 / Six



Staged Automation

Automation comes last, and only after a workflow has been run and verified by hand.

A workflow should only be scheduled after it has been run manually, verified, and constrained by rules.

SAFE AUTOMATION

- Prepare tomorrow's meeting brief from sanitized or approved notes.
- Refresh a project tracker using approved source IDs.
- Summarize open decisions from approved notes, and stage draft replies for review.
- Produce a weekly status packet.

UNSAFE AUTOMATION

Sending messages without approval, or updating records without review.

Pulling from unapproved systems, or inferring sensitive facts from incomplete data.

Persisting sensitive extracts into memory, logs, outputs, or review files.

PART THREE

03

In Practice

Build a system that helps the leader prepare, decide, communicate, and follow through with more discipline.

THE H-EXEC RULE OF THUMB

3.1 IN PRACTICE

Worked example: daily operating brief

Scenario. A service-line leader has a Monday morning huddle, a vendor follow-up, a staffing coverage question, and a quality metric review. The packet contains sanitized meeting notes, a de-identified metric summary, and a project status excerpt.

INPUT PACKET

- **Source A.** Sanitized meeting notes, dated 2026-08-07, owner: operations manager.
- **Source B.** De-identified throughput metric summary, dated 2026-08-08, owner: quality analyst.
- **Source C.** Vendor implementation status excerpt, dated 2026-08-06, owner: project manager.

■ EXPECTED OUTPUT · STAGED FOR REVIEW

Executive snapshot: Three prep-worthy items need attention before Monday noon: staffing coverage decision, throughput metric definition, and vendor timeline confirmation.

Prep-worthy meetings:

1. Morning huddle, clarify staffing assumption and owner.
2. Vendor follow-up, confirm timeline source and implementation dependency.
3. Quality review, ask analyst to confirm metric definition and reporting period.

Open questions: Is the throughput metric mean, median, or percentile? Is the vendor date confirmed or tentative? Who owns the staffing coverage decision by end of day?

Approval status: Draft for leader review. Do not send or file.

3.2 IN PRACTICE

Executive checklists

BEFORE I RUN A WORKFLOW

- I know the data lane, and the tool/workspace is approved for that source type.
- Sources are named, dated, and current enough for the task.
- Stop conditions are clear.

BEFORE I TRUST AN OUTPUT

- Facts are tied to sources, and assumptions are labeled.
- Missing information is visible, and metrics and definitions are checked.
- The output uses the right audience, tone, and format.

BEFORE I SHARE AN OUTPUT

- Sensitive details are removed or authorized.
- The executive owner approved it.
- Legal, compliance, IT, privacy, finance, or procurement reviewed it if needed.
- The final artifact has an approval status and a source note.

3.3 IN PRACTICE

Your first 30 days

WEEK	BUILD	ENTRY & EXIT CRITERIA	OUTCOME
Week 1	Executive AI Operating Profile	Enter with one low-risk workflow selected. Exit with CONTEXT, RULES, and data lane chosen.	Operating profile set; data boundaries written and explicit.
Week 2	First live workflow artifact	Enter with sanitized or approved sources ready. Exit with a daily brief or SBAR drafted in under 10 minutes and reviewed.	First live workflow artifact drafted and cleared through human review.
Week 3	Workflow library expansion	Enter after the first workflow is reviewed. Exit with two or three SOPs drafted with stop conditions.	Repeatable work begins; two to three SOPs carry stop conditions.
Week 4	Staged automation / follow-up cadence	Enter after a workflow is manually verified. Exit with prep or status refresh running on cadence with a review gate.	Follow-up runs on cadence, never without a review gate.

THE H-EXEC RULE OF THUMB

Do not treat AI as a substitute for leadership judgment. Build a system that helps the leader prepare, decide, communicate, and follow through with more discipline.

BEGIN

Governed AI workflows, accountable to the leader.

This Agent-Agnostic Edition is the canonical H-EXEC operating model. Start with CONTEXT and RULES, run one low-risk workflow on sanitized sources, and keep the approval gate on every output.



Educational; not legal, compliance, privacy, cybersecurity, procurement, HR, clinical, or financial advice. H-EXEC does not make clinical decisions, determine compliance status, or replace approved enterprise systems, access controls, or data-handling policy. No PHI, patient data, employee records, confidential vendor terms, privileged material, credentials, internal URLs, or nonpublic financial details should be placed in any AI workflow unless the organization explicitly authorizes that exact tool, data class, use case, access model, retention model, and contractual controls. Examples are anonymized composites and do not represent any specific organization, individual, or matter.