

THE AI INSTITUTE / OPERATING INTELLIGENCE

EXECUTIVE CASE NOTE · 12 SEP / 2026



Approve the AI and Its Operating Environment Together

A UK health sandbox shows why a product decision is incomplete without the people, workflow and controls that must keep it safe.

Boards, C-suite and senior management

Research period: September 2026 – September 2026

First published: 12 September 2026

Current to 12 September 2026

PUBLICATION RECORD

An Institute thesis and decision framework

For high-stakes AI, approve the product and the operating environment as one decision because performance and safety depend on both.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

- 01
- Treat product capability and operating readiness as one investment case.
- 02
- Name the conditions the user environment must maintain before deployment.
- 03
- Fund monitoring, incident reporting and withdrawal alongside the licence.
- 04
- Reapprove material model, data, workflow or population changes.
- 05
- Keep local law and professional accountability inside every regional rollout.

Research method

This official-source case note compares the UK National Commission's 10 September 2026 recommendations with MHRA AI Airlock programme records, its independent evaluation, the planned London real-world sandbox, public deliberative research and official comparative materials. The 11 September arXiv batch was scanned as a development radar; new submissions and revisions were distinguished where stated, and preprints were not treated as deployment evidence. The four-part joint approval gate is Institute analysis.

First published 12 September 2026 · Updated 24 September 2026 · Research period September 2026 – September 2026 · Version 1.1 · Suggested citation: The AI Institute, *Approve the AI and Its Operating Environment Together* (2026).

CONTENTS

Inside this paper

01

Eleven products changed the unit of approval
The case

02

The host system changes the result
The condition

03

Approve four things together
The operating gate

04

The next sandbox moves into live care
What happens next

05

The control is global; the conditions are local
What travels

06

Rewrite the next approval paper
What leaders should do

M

Method and limitations

R

References

How to use this brief

Start with the decision and use the framework to challenge the current plan. Figures retain their original populations and definitions. Institute analysis is labelled, and limitations are recorded before the references.

Eleven products changed the unit of approval

The UK's Medicines and Healthcare products Regulatory Agency created AI Airlock to test regulatory questions against real AI-enabled medical products. Across a pilot and a second phase, the programme worked with 11 innovators on seven challenge areas, including clinical note-taking, cancer diagnostics, eye-disease detection and support for obesity treatment.²³

That work now sits underneath recommendations published on 10 September by the National Commission into the Regulation of AI in Healthcare. The commission's central conclusion is plain: oversight must become proportionate, lifecycle-based and system-wide because AI products can change, behave differently by setting and depend on the organisations around them.¹

This is a recommendation to government, not enacted law. It is also a useful implementation case. Testing real products revealed that the approval object is larger than the product.

The host system changes the result

A high-stakes AI system does not arrive alone. It enters local data, professional practice, user interfaces, access controls, staffing, escalation paths and reporting systems. Each can change performance or determine whether a weak result is caught before it matters.

The commission therefore recommends that manufacturers specify the operating conditions needed for safe use, including cybersecurity, user training and institutional readiness. It also calls for clearer responsibility, stronger implementation governance and monitoring throughout the product lifecycle.¹

The transferable point is not medical-device regulation. It is that product risk and operating risk cannot be approved in separate rooms. If procurement approves the capability while operations, risk and the workforce discover the conditions later, the organisation has approved only part of the system.

Approve four things together

First, approve the product boundary: the intended job, population, prohibited use and material limitations. Second, approve the host conditions: data quality, user roles, integration points, training and cyber controls. Third, approve the learning loop: what is monitored, who can report a concern and which changes trigger reassessment. Fourth, approve the exit: who can pause the system, how service continues and how affected people obtain review or redress.

These four decisions belong in one paper because they fail together. A technically capable model with an untrained user is not the approved system. A validated model connected to changed data is not the same system. A useful pilot without monitoring or withdrawal is not ready to become an operating dependency.

Use the gate at initial deployment and whenever the model, data, workflow, user group, population or level of authority changes.

EXHIBIT / THE JOINT APPROVAL GATE

1 Product Job, population, limits and prohibited use Decision owner: sponsor	2 Environment Data, people, workflow, integration and cyber conditions Decision owner: operator	3 Learning loop Monitoring, reporting, review and reapproval triggers Decision owner: assurance	4 Exit Pause authority, service continuity and redress Decision owner: accountable executive
--	---	---	--

The next sandbox moves into live care

The UK is extending the approach. London Region I is designed to deploy up to 10 AI-enabled medical devices in controlled, real-world NHS settings under regulatory oversight. Its stated purpose is to generate additional information about benefit, safety and effectiveness before wider adoption. Results are not yet available.⁴

That distinction matters. A sandbox is not a badge of success and a recommendation is not a deployment outcome. Their value is the learning architecture: limited scope, named partners, live observation and a path to adjust the rules as the system meets reality.

A board can borrow that architecture without creating a formal regulatory sandbox. Limit the first population, define the permitted workflow, observe the complete service and expand only when the product and environment pass together.

The control is global; the conditions are local

The joint approval principle travels across markets and sectors. The exact conditions do not. In health, professional duties, clinical governance, patient communication and redress vary. In finance, credit, conduct and model-risk rules take their place. In public services, administrative law, accessibility and citizen challenge matter.

The European Union adds AI Act and sector-specific obligations. An approval copied from one country may therefore preserve the product description while missing the local operating system.

Trust also depends on context. UK public research published this month found support for healthcare AI was conditional on strong safeguards, accountability and human oversight. Those findings should shape the UK case; they should not be silently converted into a global percentage.⁵

Rewrite the next approval paper

Choose the next AI proposal that can affect a customer, employee, patient, citizen or material financial decision. Add one page titled “Operating environment.” Name the required data, workflow, skills, controls, monitoring, incident route and exit. Put an accountable owner beside every condition.

Then ask what would make this a different system. A model upgrade, a new population, a new data source or added authority should not inherit approval automatically. Set the reapproval trigger before the pilot begins.

The goal is not to make every deployment slow. It is to stop speed being purchased by leaving half the system unowned. Approve the product and its operating environment together, then let both earn the right to scale.

RESEARCH RECORD

Method and limitations

Method

This official-source case note compares the UK National Commission's 10 September 2026 recommendations with MHRA AI Airlock programme records, its independent evaluation, the planned London real-world sandbox, public deliberative research and official comparative materials. The 11 September arXiv batch was scanned as a development radar; new submissions and revisions were distinguished where stated, and preprints were not treated as deployment evidence. The four-part joint approval gate is Institute analysis.

Limitations

The National Commission recommendations are not enacted law and await a government response. AI Airlock involved a small, selected set of UK health products; innovator case reports are self-authored and do not establish population-level benefit. London Region I has not reported outcomes. Health-system law, infrastructure, workforce capacity and public expectations vary materially by market. The case supports an approval practice, not a claim that every AI product carries the same risk.



Sources and limitations

Source facts link to the original publication. Institute analysis reconciles definitions and turns the material into a management proposition. Frameworks are decision aids, not claims of universal causality. Read every important statistic with its population, date, definition and caveat.

Research cutoff: 12 September 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 UK Government, National Commission into the Regulation of AI in Healthcare

Independent advisory recommendations published 10 September 2026; government response pending.

02 MHRA, AI Airlock regulatory sandbox collection

Official programme record covering two completed phases, 11 innovators and seven regulatory challenges.

03 MHRA, AI Airlock Sandbox Phase 2 Programme Report

Official programme report and independent evaluation; updated 27 July 2026 and not formal guidance.

04 MHRA and NHS London, London Region I regulatory sandbox

Official launch, updated 30 July 2026; controlled real-world programme with outcomes pending.

05 The Health Foundation, public views on healthcare AI regulation

UK deliberative research and public polling published September 2026.

06 European Commission, Apply AI public-sector sectoral deep dive

Official programme page updated 11 September 2026; initiatives and future event, not measured outcomes.

07 US GAO, Artificial Intelligence Acquisitions

Independent audit released 13 April 2026; 13 acquisitions across four federal agencies.

08 arXiv, recent cs.AI submissions for 11 September 2026

Development radar listing 171 entries; mixed review status and mostly preprints.

10 US FDA, lifecycle management for AI-enabled device software functions

US draft guidance published January 2025; not final and not UK law.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Research · Readiness · Capability

theai.institute