THE AI INSTITUTE / OPERATING INTELLIGENCE

INFLUENCE AND PARTNERSHIP RADAR · 30 OCT / 2026

Make AI Partnerships Answer for the Handoff

An AI referral is not a completed service. Agree who receives the work, what they need and how the partnership knows someone acted.

Boards, C-suite and senior management
Research period: October 2026 – October 2026
First published: 30 October 2026
Current to 30 October 2026

PUBLICATION RECORD

An Institute thesis and decision framework

Before expanding an AI partnership, agree how unfinished work passes to the receiving service: who accepts it, what evidence travels with it, what capacity is available and how completion or rejection returns to the sender.

WHAT LEADERS SHOULD TAKE FROM THIS PAPER

01	**A referral sent is not a referral accepted.**
02	**Give the receiving service a voice in partnership design.**
03	**Fund the response capacity as well as the AI capability.**
04	**Check what happened after the handoff, not just what the AI produced.**

Research method

Institute synthesis of two primary accounts of the 29 October Singapore health collaboration, an official WHO regional announcement and a versioned agent-handoff preprint. The research design and limits were examined, not independently reproduced. Proposed partnership practices are Institute analysis, not medical advice, legal advice or a certification standard.

First published 30 October 2026 · Updated 30 October 2026 · Research period October 2026 – October 2026 · Version 1.0 · Suggested citation: The AI Institute, *Make AI Partnerships Answer for the Handoff* (2026).

CONTENTS

Inside this paper

01 **Who is waiting at the other end?**
Leadership & Influence

02 **One announcement, several kinds of work**
The New Signal

03 **Bring the receiving service into the agreement**
Institute Analysis

04 **Passing information is not the same as continuing work**
Underlying Research

05 **The agreement must fit the local service**
Global Application

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.

01 LEADERSHIP & INFLUENCE

Who is waiting at the other end?

An AI partnership can produce a useful warning, recommendation or referral without producing the service that follows. The next decision belongs to a person, team or institution that may have different systems, priorities and limits. For leaders, the question is not only whether the AI can identify a need. It is whether anyone has agreed to receive the resulting work.

Imagine an AI customer-service tool identifying a fault that needs a specialist. It forwards a summary and closes its own task. The specialist's queue does not recognise the request, or lacks information needed to act. Each component may look busy while the customer is still waiting. This is an illustrative risk, not a reported incident in the partnership discussed here.

The Institute's recommendation is to make that transfer an explicit part of the partnership agreement. Define when the receiving service accepts responsibility, what must accompany the case and what happens if it cannot respond. A message leaving one system should not be counted as completed work in another.

02 THE NEW SIGNAL

One announcement, several kinds of work

On 29 October, Google announced a Singapore research hub, work with the National University Health System to evaluate an AI co-clinician, and a Google Cloud collaboration with NHG Health combining predictive health insights with an AI health coach. These are related initiatives, not one demonstrated outcome. Google's announcement is a supplier account of the intended collaboration. [1]

Singapore's Ministry of Health supplies an important boundary. Its summit speech says the co-clinician initiative begins in simulation before any progression towards real-world use. It describes the NHG collaboration as exploring earlier risk identification and preventive guidance, including directing clinically sensitive or urgent care towards appropriate care channels. Neither statement establishes that the proposed services already improve patient outcomes. [2]

That proposed connection to care is the leadership issue. An AI can generate a signal, but the partnership still needs an agreed route for what follows. The public announcements do not provide enough detail to judge the receiving services' staffing, response commitments or closed-case results. Those are questions to resolve, not grounds to allege that the partners have omitted them.

03 INSTITUTE ANALYSIS

Bring the receiving service into the agreement

The most useful person in the next partnership discussion may be the manager who receives the extra work. That manager can explain which requests are actionable, which need clarification and which should never enter the queue. A senior sponsor or technology supplier cannot answer all three on the service's behalf.

Agree on a small handoff record. It should distinguish the original evidence from the AI's interpretation, show what has already been done and identify the unresolved request. Include the authority and lawful basis needed to share the information. Do not send a complete personal history merely because a summary is easier to generate than to govern.

Then define acceptance. Does the receiving service acknowledge only that a message arrived, or that it has taken responsibility? What happens when a case is incomplete, outside scope or duplicated? Who remains responsible until acceptance? These distinctions belong in the operating agreement, not just in a technical integration specification.

Capacity is part of the bargain. Before expanding volume, agree who funds review, follow-up and exceptions, and what the service does when demand exceeds capacity. An AI that finds more valid needs can create valuable work while still increasing pressure on the next team. That is a planning requirement, not evidence that the AI has failed.

04 UNDERLYING RESEARCH

Passing information is not the same as continuing work

A new research preprint offers a narrower technical parallel. HSC-Bench tests six receiving models on 1,039 tasks: 300 repository-evidence tasks and 739 synthetic continuation tasks. Structured handoff information improved evidence recovery across all six models and continuation for most, but the leading models differed between the two tests. Finding the right information did not by itself establish effective continuation. [3]

The study is about agents, not hospitals or inter-organisational service delivery. It provides no patient-outcome evidence and does not prove the Institute's proposed partnership design. Its useful contribution is the distinction between what arrives at a handoff and what the receiver subsequently does. Leaders can apply that distinction as a question to test, without importing benchmark results into operational forecasts.

05 GLOBAL APPLICATION

The agreement must fit the local service

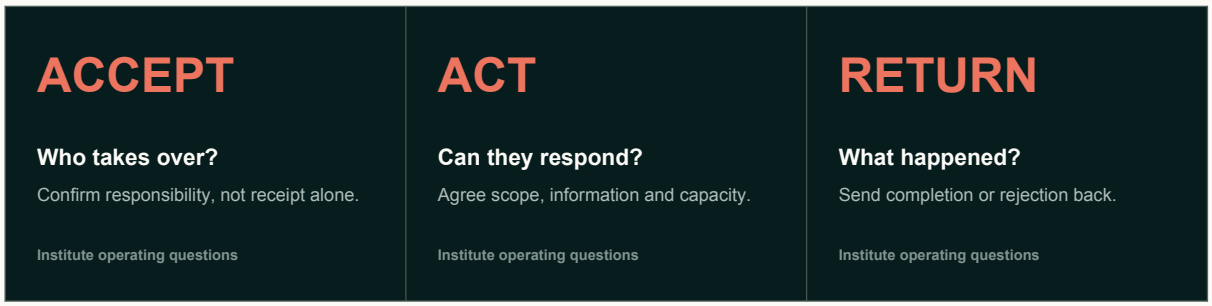
The Singapore ministry describes connected information systems and healthcare clusters responsible for defined populations. Those institutional arrangements matter to how work can move. An organisation elsewhere should not assume that a similar model licence supplies equivalent referral routes, incentives or service capacity. [2]

On the same day, WHO reported that its European Region's 53 Member States adopted a 2026–2033 AI-in-health roadmap. Its priorities include institutional responsibilities, workforce skills and implementation support. This is a regional policy roadmap, not an EU law, a product authorisation or evidence that every participating system has the same readiness. [4]

For operations in North America, Latin America, Africa, the Middle East or other Asia-Pacific markets, examine the actual receiving organisation. Payment arrangements, professional roles, language, connectivity and permitted data sharing can alter the handoff. These are local questions, not fixed regional stereotypes. Comparative evidence on this new partnership's transferability is not yet available.

Start with one agreed service path. Track requests sent, requests accepted, time waiting, cases returned and the eventual result, with appropriate privacy protection. Ask receiving staff and affected users where work was lost or repeated. Expand only when both sides can explain the transfer and support the resulting workload. The partnership has delivered a connection when somebody has taken over—not merely when the AI has handed off.

EXHIBIT / THREE PARTS OF A USEFUL HANDOFF



RESEARCH RECORD

Method and limitations

Method

Institute synthesis of two primary accounts of the 29 October Singapore health collaboration, an official WHO regional announcement and a versioned agent-handoff preprint. The research design and limits were examined, not independently reproduced. Proposed partnership practices are Institute analysis, not medical advice, legal advice or a certification standard.

Limitations

Announcements do not establish operational results. The co-clinician work begins in simulation. Agent benchmarks do not establish clinical or human-service performance. Public sources do not disclose complete operating agreements, staffing or response commitments. Local legal, institutional and service conditions must be assessed before transfer.

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: 30 October 2026. Review cycle: six months, or earlier after a material change in the underlying facts, standards or policy.

REFERENCES

Evidence behind the thesis

01 **Google: Singapore health collaborations**

Ben King, 29 October 2026. Supplier announcement of research, evaluation and population-health initiatives, not independently measured outcomes.

02 **Singapore Ministry of Health: Google for Health Summit speech**

Tan Kiat How, 29 October 2026. Paragraphs 9–12 distinguish exploration and simulation; paragraphs 15–18 describe institutional foundations. Official policy explanation, not proof of clinical effectiveness.

03 **Li and colleagues: Benchmarking Task Takeover Across Agent Handoffs**

Submitted 27 October 2026, new submission in the 28 October release. Preprint: 300 repository and 739 synthetic continuation tasks, six receiving models. Methods and discussion examined; no independently verified peer-review acceptance or clinical study.

04 **WHO Europe: AI-in-health roadmap adopted**

29 October 2026. Reports adoption by 53 Member States of the WHO European Region; roadmap covers 2026–2033. Regional policy direction, not national legislation or deployment approval.

THE AI INSTITUTE

Make the decision.
Then measure what changes.

Start with seven questions and a free readiness report. Then choose an optional 10–15 minute verbal assessment with Maya, our AI interviewer.

Start your free assessment: readiness.theai.institute

Research · Readiness · Capability

theai.institute