



**REGIONAL OFFICE FOR THE WESTERN PACIFIC
BUREAU RÉGIONAL DU PACIFIQUE OCCIDENTAL**

REGIONAL COMMITTEE

WPR/RC77/6

**Seventy-seventh session
Manila, Philippines
19–22 October 2026**

17 July 2026

ORIGINAL: ENGLISH

Provisional agenda item 9

ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI) has become a national political commitment for many Member States in the WHO Western Pacific Region. Prior to the proliferation of generative AI, machine learning and deep learning were already in use to address systemic health challenges, from managing population health to strengthening the continuum of care. Generative AI plays a much broader role. Trained on the vast archives of recorded information and knowledge, generative AI has moved beyond solving specific problems to providing more general-purpose support for health system functions.

Well-resourced settings are already operationalizing AI in health systems, while many resource-constrained settings in the Region have yet to start because of limitations in infrastructure, data, financing, governance and regulatory oversight. Concerns remain – revealed by early experiences and evidence from well-resourced settings – around the performance, ethical and governance implications of AI. As AI technologies evolve and transform societies, they will continue to transform health systems. The cost of not engaging AI will be high. Health systems architects that understand the abilities and limitations of AI will be uniquely positioned to unlock the full potential of AI for system transformation. The question is not whether ministries of health should adopt AI, but what actions should be taken today so that AI helps to deliver better health and well-being for people tomorrow.

This draft document provides Member States with three interconnected operational frameworks to shape the role of AI in public health: the first prioritizes use cases; the second helps determine the minimum necessary ecosystem; and the third helps build institutional and workforce capacity. These frameworks, which were developed through expert and Member State consultations, have been validated through country workshops. The Regional Committee for the Western Pacific is requested to consider for endorsement the draft *Guidance for Responsible AI in Health in the Western Pacific Region*.

DRAFT
GUIDANCE FOR RESPONSIBLE AI IN HEALTH IN THE
WESTERN PACIFIC REGION

1. BACKGROUND

Across the Western Pacific Region, ministries of health are confronting growing and interlinked system pressures. Demographic change is reshaping demands for care, with ageing populations and the rising prevalence of noncommunicable diseases straining health services. At the same time, many countries continue to face workforce shortages and maldistribution, constrained fiscal space, fragmented data systems and limited interoperability. Climate change and emergencies have increased the frequency and complexity of public health threats, while misinformation and broader social and commercial determinants of health further complicate decision-making and service delivery.

These challenges weigh heavier on the Region's lower-middle income countries. With limited resources, health ministries have increasingly recognized the importance of innovation to expand access in remote areas, strengthen the health workforce, improve system effectiveness and achieve better population health outcomes. Within this context, AI has emerged as a potential enabler of broader health system transformation.

AI refers to a series of machine-based technologies that analyse data to generate outputs, such as predictions, recommendations, content or decisions that may support or influence human action. AI encompasses a range of approaches, including statistical learning, machine learning, deep learning, natural language processing and, more recently, generative AI.

As a national political priority in some parts of the Region, the use of AI in health has expanded over the past decade. Prior to the emergence of generative AI, machine learning-based applications were already being used to augment epidemiological modelling for climate-sensitive diseases, support screening and diagnosis for high-burden conditions, digitize paper-based health records and improve specific aspects of service delivery. The emergence of generative AI-based applications has expanded the scope of AI from task-specific applications to more general-purpose tools that can influence clinical documentation and decision-making at scale.

While some well-resourced settings are beginning to operationalize AI within health systems, engagement with AI remains constrained in many countries because of limitations in digital infrastructure, availability and quality of data, sustainable financing beyond pilot projects, governance arrangements and regulatory oversight. Early evidence from those operationalizing AI emphasizes concerns related to performance and evidence gaps, safety and system effects, including risks of

hallucination, bias, over-reliance on AI and the potential deskilling of health workers. As AI technologies continue to evolve and reshape societies, however, the cost of inaction is likely to increase. Health systems that lack the expertise to understand the capabilities, limitations and risks of AI may be increasingly exposed to fragmented, vendor-driven or externally shaped adoption, which would likely limit the ability to safeguard public value and equity.

Hence, the question is not whether to engage with AI, but rather how to ensure that the adoption of AI in health is in the right direction and aligned with the broader health system transformation. For resource-constrained settings, this means starting with use cases that are feasible today, while deliberately building the institutional capacity, safeguards and evidence base needed for more transformative AI use over time.

2. PLAN OF WORK

The starting point for responsible adoption of AI in health is not the selection of a technology; it is the clear definition of the issue to be solved, such as improving access to screening for high-burden diseases in underserved areas, supporting frontline and village health workers in care delivery, or enabling more timely decisions informed by higher-quality data. Clearly defining the challenge determines the appropriate AI use case and, in turn, the choice of technology.

The end goal is not deployment of an AI solution; it is resolving the issue. Between the starting and ending points, an enabling ecosystem governed by the public sector is the critical success factor. This includes aligned changes to care pathways, workforce roles, incentives, accountability arrangements, information systems and service delivery models. Technology is only as effective as the systems and governance within which it is embedded.

At the 2025 session of the Regional Committee for the Western Pacific, Member States identified three action priorities for ministries of health during the technical discussion on AI in health care: (1) prioritization: identify high-impact use cases that can advance health equity and address health system constraints; (2) ecosystem strengthening: build fit-for-purpose governance, regulatory frameworks and enabling conditions for safe and effective implementation; and (3) capacity-building: strengthen public sector capacity to govern, adopt and use AI responsibly.

Those priorities translate into three questions that Member States may use to pragmatically guide responsible engagement of AI:

1. Which AI use cases should be prioritized?

2. What minimum safeguards and enabling conditions are required?
3. And what institutional and workforce capacities are required to govern AI effectively over time?

The frameworks provided herein are designed to support ministries in answering these questions in ways that are feasible today, while building towards more transformative use of AI as system capacities grow.

2.1 Prioritization of use cases

Prioritization helps ministries of health decide which health system problems are appropriate candidates for AI and when non-AI interventions would deliver greater value. Not every problem warrants an AI solution, and not every AI use case warrants public investment. Prioritization should begin with a problem-driven assessment of public value, before considering specific technologies. Use cases that emerge as priorities should then be subject to tool-specific assessment, including evidence of validity and regulatory status, before implementation decisions are made.

Benefits must be balanced against risks. Does addressing this problem with AI generate meaningful benefit relative to the actual standard of care available in the intended setting, particularly for underserved populations? What level of risk is inherent in the use case itself, given the severity of harm if the AI-influenced decision is wrong, the degree of human oversight between AI output and impact, and the reversibility of the decision?

Use-case risk helps ministries determine whether AI is appropriate for addressing the problem at all, and whether the expected benefits justify proceeding. Higher-risk use cases may still be justified when they address critical unmet needs, but they require stronger safeguards at the implementation stage.

Other questions revolve around the health system absorptive capacity, cost and sustainability. Can the health system act on AI outputs through available services, referral pathways, workforce capacity, workflows and decision processes? A use case that generates information and needs the system cannot act upon may compromise the real-world benefit. Can the use case be financed across its full life cycle, including implementation, infrastructure, training, governance, monitoring and any additional service demands generated, at a cost proportionate to expected value? Is there a pathway to sustainability if it proves beneficial?

Risk-proportionate approach

Across global AI governance frameworks, a low-, medium- and high-risk typology is widely used for practical classification. These frameworks typically assess risk at the level of the use case, based on the type of decision involved and the potential impact on patients or populations. Such classifications provide a useful starting point for identifying applications that warrant closer scrutiny.

From the health systems as AI implementers perspective, however, risk is not solely determined by use case. Implementation risk – the risk that arises from how, where and by whom an AI system is deployed within a specific health system – deserves greater concern. The same AI technology can present very different levels of risk depending on the deployment context. For example, an AI system for diabetic retinopathy screening may function as a productivity tool in a hospital setting with specialist oversight. But that same system may become the primary decision-maker in a rural clinic without access to an ophthalmologist. The implementation risk in the latter setting is higher, even if the underlying technology is the same. At the same time, the potential benefits in the higher-risk setting may also be greater, as the alternative may be no screening.

A risk-proportionate approach in the context of health system implementation asks what the inherent risk is and how it will be used in practice, and applies safeguards that are proportionate to the level of implementation risk. High risk does not necessarily mean “do not deploy”. It does mean that stronger governance, validation, monitoring and accountability mechanisms are required.

2.2 Minimum ecosystem for implementation

Whether an AI technology delivers real benefits depends both on the technology and the conditions under which it is introduced and used. The minimum ecosystem refers to the essential conditions that must be in place for a specific technology to be implemented safely, effectively and sustainably within a health system. Its requirements are determined by the use case, but assessed against the specific AI solution and deployment context.

Countries do not need to wait for a fully developed national AI ecosystem before acting. They can begin with feasible, lower-risk use cases, while strengthening institutions and systems over time. The level of readiness required should be proportionate to the implementation risk of the use case.

Five areas help ministries of health determine whether the minimum ecosystem is in place: governance, data and digital infrastructure, validation, financing and capacity.

Governance

There should be a clear public authority or function responsible for approving, overseeing and, when necessary, suspending the use of AI in health. Governance arrangements should be proportionate to risk: lighter oversight may be sufficient for low-risk AI implementation, while higher-risk implementation requires clearer accountability, stronger review and ethics or regulatory engagement, as appropriate.

Data and digital infrastructure

Health data used by AI should be protected, secure and fit for the intended purpose. Countries should be able to control where data are stored, who can access them, how they are shared and whether they are used to train or improve external systems. Data and infrastructure arrangements should support interoperability, privacy and data sovereignty.

Validation

There should be evidence that the AI performs as intended in the setting where it will be used, or in comparable contexts. Higher-risk implementation requires stronger prospective evidence and monitoring of performance to detect errors, bias, unintended effects or loss of effectiveness over time.

Financing

Funding should cover more than the initial procurement. Ministries should consider the full life cycle of implementation, including system integration, training, monitoring, cybersecurity and additional service demands generated by the AI application. Each implementation should have a clear pathway to sustainability, with decision points for continuation, scale-up or discontinuation based on performance and value.

Capacity

Institutions and individuals involved in governing, procuring, using and overseeing AI should have the skills and authority to fulfil their roles. Higher-risk applications require stronger technical, clinical and managerial capacity, as well as mechanisms for public and user engagement throughout the design, deployment and monitoring.

2.3 Capacity-building

The purpose of capacity-building is to ensure that institutions and actors involved are able to do what they need to do, in line with national priorities and ecosystem readiness. Countries should

identify the minimum capacities for different stakeholder groups required to begin safely, while building deeper and more durable capabilities over time through implementation, supervision, evaluation and learning.

For policy-makers and institutional leaders, the core capacity is stewardship. This encompasses understanding opportunities and challenges, setting priorities, establishing proportionate governance and regulatory pathways, defining boundaries for data use, clarifying accountability and making informed decisions about whether to adopt, scale, modify or discontinue AI applications.

For planners, evaluators and implementers, the priority is implementation readiness and evidence generation. This includes the ability to assess technologies, estimate full life cycle costs, design and implement evaluations.

For practitioners and providers, capacity-building should focus on safe and appropriate use in routine practice: understanding the intended use and limitations of AI tools, knowing when human judgment should override AI outputs, communicating AI-supported decisions, and identifying and reporting errors or unintended effects.

For patients, communities and civil society, the core capacity is meaningful participation. This includes the ability to engage in decisions about whether and where AI is introduced, what safeguards are required, how benefits and risks are communicated, and how effects on access, equity, trust and rights are monitored over time.

Across all groups, certain competencies are foundational: people-centredness, ethical responsibility, critical appraisal, evidence-informed decision-making, communication, collaboration and accountability.

2.4 Recommendations for Member States

Lead with ministry-owned prioritization of AI use cases

Ministries of health should lead the identification and prioritization of AI use cases that address national health priorities, advance equity and are likely to deliver public value in real-world settings. This process should inform a ministry-led AI road map or an initial list of priority use cases, rather than allowing adoption to be driven primarily by vendor offerings, donor projects or ad hoc institutional initiatives.

Establish a risk-proportionate governance pathway

Ministries of health should designate or strengthen governance mechanisms with the mandate to approve, oversee and, when necessary, suspend the use of AI in health. Use cases should be classified according to risk, with higher implementation risk cases subject to stronger governance and oversight. Where national regulatory capacity is limited, Member States may draw on reliance mechanisms and trusted external assessments while retaining national decision-making authority.

Define and assess the minimum ecosystem required for each priority use case

For each priority use case, ministries of health should define what minimum conditions must be in place for safe and effective implementation of a solution. This assessment should be used to decide whether implementation can proceed now, should be staged with parallel system strengthening or should be deferred until key gaps are addressed.

Fund the ecosystem, not only the solution

Public investment decisions should cover the full implementation life cycle and any downstream service demands generated by AI use. Ministries of health should require sustainability plans to avoid short-term pilots without long-term system integration plans.

Build institutional and workforce capacity for AI

Ministries of health should define role-specific competencies for priority AI use cases to be pursued, starting with the minimum capacities needed. Capacity-building should be embedded into existing education and professional development systems.

3. ROLE OF THE WHO SECRETARIAT

WHO will assist Member States to move together towards responsible adoption of AI in health by turning global principles into practical national actions, facilitating shared learning across countries and areas, and supporting coordinated approaches to reduce duplication and strengthen decision-making.

Develop operational tools for implementation: WHO will adapt, operationalize and complement global guidance through practical regional tools, implementation support packages and methodologies that respond to the diverse health system contexts of the Western Pacific Region. This may include prioritization tools, minimum ecosystem assessment tools, monitoring approaches and implementation guidance aligned with health technology assessment and quality assurance principles.

Convene policy dialogue and intergovernmental exchange: WHO will convene Member States and technical partners to exchange experience in AI implementation, governance and regulation, and identify areas for collaboration. A regional technical advisory mechanism will be established to support this work by providing expert input on priority issues of shared interest.

Catalyse country implementation: WHO will provide technical assistance to country efforts in prioritization of use cases, risk-proportionate safeguards and ecosystem strengthening, and catalyse country implementation of AI use cases for which the public health value is high but market incentives or local capacity are insufficient.

Provide regional learning platform: WHO will support a regional catalogue or learning platform documenting AI use cases, implementation risk classifications, evidence status and lessons from implementation, enabling countries to learn from comparable settings and avoid duplicating efforts. Existing WHO learning platforms, including the Pacific Open Learning Health Net and the Global Health Learning Centre, can be used to support structured training pathways for national focal points and relevant institutions.

Build partnerships for public goods: WHO will facilitate partnerships to strengthen collaboration among ministries, academic institutions, implementers, civil society and appropriate technology developers and open-source communities to develop regional public goods, such as common AI models adapted to settings where they are needed in the Region.

4. CONCLUSION

AI is no longer an optional consideration for health systems in the Western Pacific Region. It will reshape how health services are delivered, how decisions are made and how scarce resources are used. Used strategically and responsibly, AI can help countries expand access, strengthen public health functions and improve system performance. Using AI poorly or not at all, countries run the risk of deepening inequities. The central policy task for ministries of health is not whether to adopt AI in health, but how to shape its use so that it serves the public, creating value and promoting equity, safety and long-term health system transformation that prioritizes population health and well-being for all.

This guidance provides Member States with practical pathways for shaping AI in health through three linked approaches: prioritizing the right use cases, applying a risk-proportionate minimum ecosystem, and building institutional and workforce capacity. Together, these translate into five concrete actions for ministries of health: lead use-case prioritization; establish risk-proportionate governance for AI implementation; define and assess minimum conditions before deployment; fund the

ecosystem, not just individual solutions; and build stewardship and implementation capacity as core health system functions.

The time to begin is not when every condition is perfect, but when countries have a clear public-interest purpose, adequate safeguards and a commitment to monitor, learn and adapt over time. WHO stands ready to support Member States in translating this regional guidance into national action through technical assistance, regional collaboration and shared learning.