CO-DEVELOP Vital Wave PATH Patrick J McGovern Foundation

Funded by Wellcome

Digital Public Infrastructure for Health

Charting a path to implementation in LMIC health systems

List of Acronyms

AI	Artificial Intelligence
CHIS	Community Health Information System
CHW	Community Health Worker
CR	Client Registry
DIAL	Digital Impact Alliance
DPG	Digital Public Good
DPI-H	Digital Public Infrastructure for Health
EA	Enterprise Architecture
EMR	Electronic Medical Record
FHIR	Fast Healthcare Interoperability Resources
GG	Global Good
HCW	Health Care Worker
HIE	Health Information Exchange
HIS	Health Information System
HIV	Human Immunodeficiency Virus
HMIS	Health Management Information System
ICT	Information and Communication Technology
LMIC	Low- and Middle-Income Country
LMIS	Logistics Management Information System
ML	Machine Learning
OpenHIE	Open Health Information Exchange
OpenHIM	Open Health Information Mediator
OpenMRS	Open Medical Record System
PEPFAR	The U.S. President's Emergency Plan for AIDS Relief
SHR	Shared Health Record
SMART	Standards-based, Machine-readable, Adaptive, Requirements-based, and Testable
WHO	World Health Organization

Table of Contents

Foreword

The global development community is entering a new phase in our collective conversation on implementing digital systems. After decades of investment in siloed, built-for-purpose digital development, countries have begun to realize the inherent potential of a “building block” approach: creating integrated services that can be combined to address multiple use cases. This digital public infrastructure (DPI), which includes governance structures and policies that prioritize safety and inclusion, can achieve development goals more efficiently, faster, and at a larger scale than the old, monolithic models. In country after country, digital public infrastructure has emerged as a critical priority, and we at Co-Develop are committed to responding to that demand with support to accelerate the adoption of DPI.

With this context in mind, we began asking how a “DPI approach” can be applied to digital health – not as a substitute for foundational DPI such as payment or identity systems, but as a domain-specific complement to it. What are the critical, underlying functionalities that are missing in current digital health systems that, if enabled at scale, could fill the gaps to support many of the priority use cases and health system goals of the next five to ten years? Are there architecture or implementation approaches that could unlock benefits to improve health outcomes?

This report is an early step in that process. It offers a framework for conceptualizing digital public infrastructure for health that is grounded in priority digital health interventions. It analyzes current challenges and gaps in the landscape to develop investment approaches that can generate conditions for implementations of DPI for health to succeed. While this work is necessarily technical in nature, it creates an anchor for further implementation research to determine how DPI for health can best support an ecosystem of interoperable digital health applications to achieve the inclusive, person-centric, equitable outcomes we strive for. I hope this work will prompt discussion in the global digital health community and hasten the next phase, not only of conversation but also of investments that bring to digital health what DPI is beginning to realize for broader digital transformation.

Tim Wood
Chief Partnerships Officer
Co-Develop

Executive Summary

The Potential of Digital Public Infrastructure for Health
The next decade of global digital health has the potential to greatly accelerate the achievement of sustainable, scaled, interoperable digital health systems in low- and middle-income countries (LMICs), while laying the foundation for emerging capabilities like predictive analytics and artificial intelligence (AI) and enabling their use in these contexts. This report maps out a pathway toward realizing these possibilities through an infrastructural approach to digital health investments. Rather than focusing on specific digital health interventions, an infrastructural investment approach prioritizes a subset of digital systems and services that act as the foundation for many priority use cases in digital health, today and in the future. This infrastructure-focused approach enables an ecosystem of innovators, implementers, and governments in gaining efficiency by leveraging a common set of open, widely scaled, relatively basic digital functions for health systems.

This approach builds on the emerging discourse around digital public infrastructure (DPI) by recognizing the transformative potential of inclusive, foundational, open, and transparently governed digital platforms. DPIs are intended to perform a basic function, at scale, in a way that others can build upon. For example, digital identity and data exchange systems based on an open, interoperable codebase enable an efficient pathway to scale for a multitude of applications and services. Digital public infrastructure for health (DPI-H) encapsulates health-specific digital functions that similarly enable a multitude of digital health interventions to effectively scale (Figure 1). Creating conditions for countries to adopt DPI-H opens a critical path to achieving the goals of the global digital health community.

Figure 1: Conceptual Model for DPI and Health Domain-Specific DPI

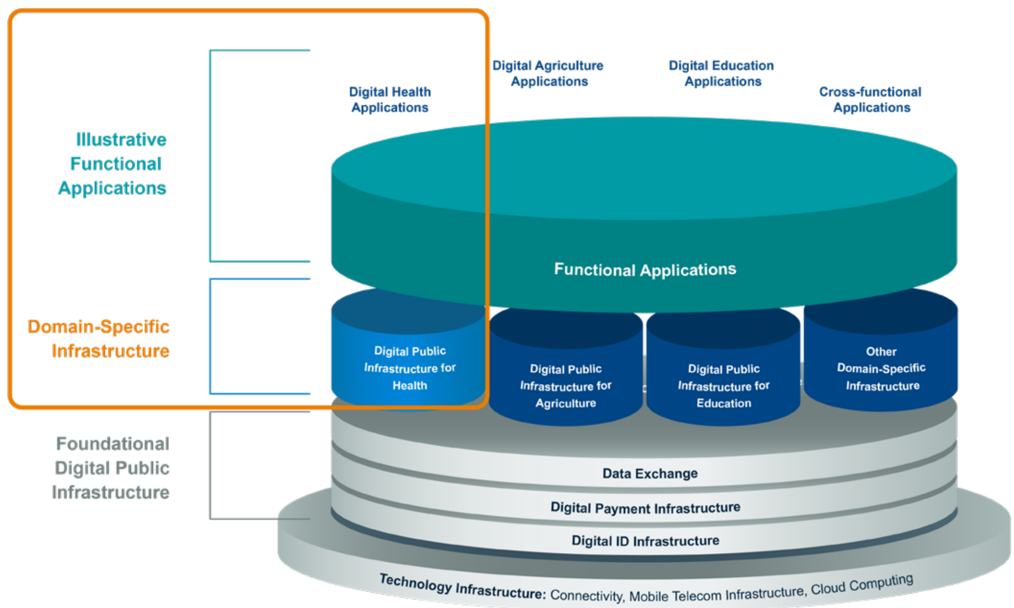

Existing landscape of health-related global goods and digital public goods is not ready to function as DPI-H

For countries to realize the benefit of DPI-H, there must be a reliable set of products that can provide core DPI-H functionalities at scale and that allow other solutions to leverage their underlying functionality. While there is a strong desire in the digital health donor community to support the use of open-source products and digital public goods, the existing set of open applications that correspond to DPI-H functions are often not, in practice, functional at scale or able to support other applications in leveraging their functionality. Limited investment in core product development, deployment tools, interoperability functionality, enterprise-standard testing, and routine security patching and maintenance hinder the ability of these products to serve as DPI-H in country implementations.

At the same time, open data in the form of representative training data for machine learning models and open models needed to utilize predictive analytic capabilities are only starting to be available as digital public goods. AI models are emerging with nascent governance frameworks to ensure safety, accountability, and effectiveness of their use. Open data and representative training data sets to develop such models are not yet widely available and easily discoverable. Active multi-disciplinary communities of data science and digital health experts are only starting to emerge. Strengthening the ecosystem for DPI-H will require strengthening both the market of supporting DPGs and governance structures that will ensure their suitability for use with DPI-H.

Overlapping challenges at global and country levels hinder the potential use of DPGs for DPI-H

At a country level, numerous persistent challenges in digital health enabling environments hinder the ability to implement DPI-H. Leadership and governance capacity to direct digital health investments following an approved architecture for health information exchange, as well as regulatory capacity to enforce compliance with existing architectures, remain critical capacities to strengthen. At the same time, funding patterns that prioritize investments in disease verticals and bespoke applications to support specific programmatic goals have created a fragmented landscape. Investment in foundational digital platforms that cut across disease domains to serve the health sector as a whole has been limited, particularly for infrastructural platforms such as the components of a health information exchange. Similarly, technical expertise on standards-based interoperability platforms is scarce. Further, efforts to develop and champion the adoption of open-source tools that could provide infrastructural functions have historically lacked continuity and the amount of support needed to ensure secure, interoperable, scalable implementation. Together, this has left most countries with insufficient resources to close the gap between fragmented implementations of individual components and comprehensive investments in DPI-H.

Complementary global and country-level investments are needed to advance implementation of DPI-H in LMICs

Three broad approaches can hasten the realization of DPI-H. First, the global digital health community needs to invest in strengthening the existing market to provide a set of DPGs better able to function as DPI-H. Further investments need to fill gaps in the product landscape related to health data security and consent, strengthen the ecosystem of Supporting DPGs and building blocks, and facilitate the operationalization of governance for future predictive analytics.

Second, countries need continued, consistent support to strengthen digital health enabling environments that support the implementation of DPI-H. Strengthening leadership capacity to support compliance with approved strategies and architectures, technical expertise to implement and maintain DPI-H components, and regulatory capacity to navigate evolving digital health landscapes and enforce legislation will create conditions in which countries are positioned to advance DPI-H implementation.

Third, as the global and country-level ecosystems strengthen for DPI-H, coordinating donor investments focused on the implementation of DPI-H will be positioned to succeed. Long-term support dedicated to the implementation of multiple DPI-H components can create the foundation for a thriving digital health ecosystem in which many priority use cases can scale.

The last decade of investment in global health has laid the groundwork for a new vision for coordinated, planned digital systems in which DPI-H provides the foundation. Now is the time to support the operationalization of that vision through coordinated investments at the global and country levels. Working together, the global digital health community can support digital health ecosystems to effectively leverage DPI-H and accelerate progress toward the health outcomes we all seek to achieve.

Introduction

Collectively, the digital health community is working towards a future in which health system performance is strengthened and health outcomes are improved through the use of sustainable, scaled digital health interventions across low and middle-income countries (LMICs). At the same time, digital health capabilities are expanding. New analytic approaches, including machine learning (ML) and other predictive analytics, have the potential to further improve outcomes and efficiency by enabling personalized, risk-based care; more accessible diagnostics; and earlier action to prevent health threats. The next decade of global digital health initiatives and activities will require investments that both accelerate progress towards the long-held vision of sustainable, scaled digital health tools and systems and also create a foundation for forward-looking capabilities such as predictive analytics. This report sets out to identify the underlying infrastructure and ecosystem enablers that will support both aims. Further, it maps out investment approaches that strengthen the realization of infrastructure-based implementation of digital health systems.

Rather than focusing on specific digital health interventions, an infrastructural investment approach prioritizes a subset of digital systems and services that act as the foundation for many priority use cases in digital health, today and in the future. This infrastructure-focused approach enables an ecosystem of innovators, implementers, and governments to gain efficiency by leveraging a common set of open, widely scaled, relatively basic digital functions for health systems. This report begins by identifying the underlying digital public infrastructure for health (DPI-H) that will support numerous priority use cases. From there, it explores the extent to which existing products could serve as DPI-H and the implementation experience that several countries have had to date. Analysis of these challenges informs a set of investment strategies that can facilitate the realization of successful DPI-H in LMICs.

Approach to Defining Digital Public Infrastructure for Health

For this report, DPI-H is generally understood as the health-specific complement to foundational digital public infrastructure (DPI). The Digital Public Goods Alliance defines Digital Public Infrastructure as "solutions and systems that enable the effective provision of essential society-wide functions and services in the public and private sectors." The Center for Digital Public Infrastructure further explains DPI as a set of technology building blocks that are powered by interoperable, open standards and that operate under a set of enabling rules and open, participatory governance. Relying on DPI as core digital infrastructure can drive innovation, inclusion, and consumer choice at scale. DPIs can reduce both physical and cost barriers to service access by enabling secure, inclusive, low-to-no cost digital alternatives. Digital identity, payments, and data exchange infrastructure are common examples of DPI that have shown how publicly-governed, interoperable digital platforms can spur improvements in health and economic welfare, particularly for women.

DPI-H, then, is conceptualized as the health-specific components of a country's digital infrastructure that enable an ecosystem of inclusive, scaled, user-driven digital applications in a health system. Figure 1 depicts a conceptual model for how foundational digital public infrastructure, together with technology infrastructure, forms a base on which domain-specific digital public infrastructure can operate. This enables actors within the broader ecosystem to leverage foundational DPI and technology infrastructure, domain-specific infrastructure, and other innovation-supporting DPGs to create cross-sector and domain-specific applications.

Figure 1: Conceptual Model for DPI and Domain-Specific DPI

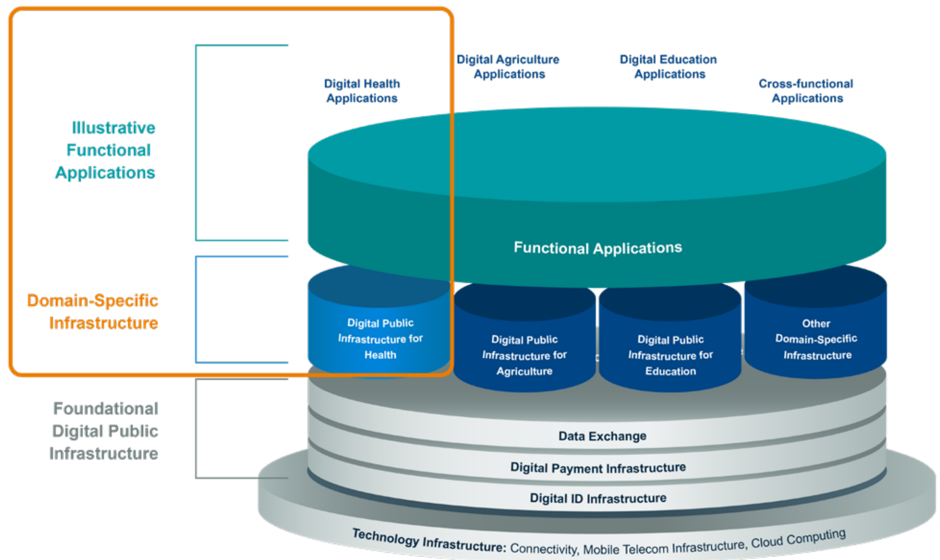

Conceptual model adapted from DIAL

To add more specificity to the domain-specific layer, the first section of this report unpacks the component parts of DPI-H conceptually – the combination of foundational digital systems and services that would enable an inclusive ecosystem of functional digital health applications to be created, scaled, and sustained within a country's digital health system.

Conceptually defining and then identifying specific components of an ideal DPI-H enables a gap analysis between the current state and potential future states of digital health systems in regard to the development of DPI-H. This analysis is developed throughout subsequent sections of the paper, as described in the Report Roadmap below.

Report Roadmap

Part 1: Identifying Digital Public Infrastructure for Health
This section articulates specific components of DPI-H based on an inductive exploration of three user journeys. User journeys allow infrastructure to be considered from different user perspectives and to more clearly distinguish the basic, common components of the functionalities required across use cases at different levels of the health system. Further, adding forward-looking use cases that include predictive analytics in the user journeys enables identification of potential gaps or needed additions to the common conception of infrastructure. This section breaks down each user journey to identify the basic and infrastructural components common across required functionalities.

Part 2: Current Product Landscape of DPGs, Global Goods, and Building Blocks
This section explores in greater detail the gap between existing DPGs, global goods (GGs), and building blocks that could play an infrastructural function, and the current reality of scalable, sustainable DPI-H in LMICs. It reviews the current market of DPGs, GGs, and building blocks against DPI-H infrastructure components and explores the potential of existing tools to function as DPI-H. It also includes three brief country case studies to understand how DPGs have been implemented in LMIC contexts as well as the challenges and limitations of their use as infrastructure across the health domain.

Part 3: Challenges Using DPGs as DPI-H
This section synthesizes challenges related to using DPGs, GGs, and building blocks for DPI-H within a country environment, and it highlights considerations around the broader enabling environment that are critical to DPI-H implementation.

Part 4: Recommended Approaches to Support DPI-H Implementation
This section describes multiple investment approaches that address challenges in using existing DPGs, GGs, and building blocks for DPI-H. It includes approaches for both the global digital health community and country-level investments. Recommended approaches include specific strategies and are accompanied by an order of magnitude estimate of costs associated with illustrative investments.

Part 1:

Identifying Digital Public Infrastructure for Health

23898: Laboratory technician logging information into a computer in the regional hospital laboratory in Thies, Senegal. PATH/Gabe Bienczycki

Part 1: Identifying Digital Public Infrastructure for Health

This section takes a user-journey-based approach to identify the components of DPI-H that would enable a range of health applications. It explores the functional needs of digital health systems by unpacking three user journeys, each representing a different perspective in the healthcare system: a client, a healthcare worker, and a health system manager.

For each perspective, the user journey articulates multiple digital health functionalities that are further broken down into potential applications, supporting content and platforms, and the underlying DPI-H components needed to enable them. The user journeys include digital health functionalities that are familiar today (yet often absent or not mature) as well as emerging functionalities that incorporate advanced data analytic approaches such as machine learning and natural language processing (see Appendix 1 for more information). The inclusion of both familiar and emerging functionalities ensures that the identified DPI-H components consider both current and future needs.

The user journeys provide a composite of important digital health interventions and use cases identified through a series of consultations with digital health and LMIC regional representatives. The user journeys are not intended to reflect a comprehensive mapping of all functional applications that will be useful in the future, nor do they include all of the pain points individuals may experience in the health system. Rather, the user journeys are a mechanism used to identify a range of functional needs that are valuable to end-users and within the realm of applications one would expect DPI-H to support. To further validate the relevance of the functionalities explored in each user journey, they are cross-referenced to the WHO Classification of Digital Health Interventions.

The following section details the user journeys, associated digital health interventions and functionalities, and the underlying DPI-H needed for scaled operation. At the end of the user journeys, a summary table maps the digital components required for the use cases across all three journeys, segmented into basic technology infrastructure, foundational DPI, functional applications, and candidates for DPI-H that provide a basic function underlying many digital health use cases.

1.1 User Journey 1: Client Engagement with Primary Health Services

Primary healthcare services are an essential component of country public health systems, particularly in the context of universal health coverage. Comprehensive primary health care enables equitable and accessible coverage to citizens, especially those in rural and marginalized communities, and improves health outcomes by providing preventative and curative care within these communities. Effective care at the primary health level relieves the burden on hospitals and tertiary services. This user journey focuses on the needs of clients engaging with primary health services.

"As a client, I want to be able to find where to get the health services I need. I would like to be able to access and share my health record to be sure I get the right treatment, even if I go to different places. I also want the drugs that I am prescribed to be safe, of good quality, and available when I go to collect them.

I'm thankful for the care that I get, but sometimes I feel like the health information and the treatment I'm given isn't really what I need. And I wish there was an easier way to pay my service fees and use the vouchers and insurance products that are starting to be available to me."

This user journey has multiple functional requirements. These include clients finding and connecting with care through **care service discovery**, tracking and sharing their health information to improve **continuity of care**, accessing safe and reliable **medication**, benefiting from **precision public healthcare**, and more easily using **digital payments and vouchers** that support service delivery. Each of the functional requirements relies on several underlying DPI-H components, as described below.

1.1.1 Desired Functionality: Care Service Discovery

WHO Classifications: 1.6 – On-Demand Information to Clients, 3.7 – Facility Management, 4.3 – Location Mapping, 4.4. – Data Exchange and Interoperability, 1.5 – Citizen-Based Reporting

Care services discovery involves an online platform in which clients can search for the location of a facility based on needed services. Clients in LMICs often are unaware of where the services they require are available, and as a result, can waste time traveling to and waiting to be served at facilities that do not offer what they need. For example, with care service discovery, mothers can go straight to the correct primary health care facility for free immunization of their children, rather than wasting time and money going to the wrong facility or going to a general practitioner who will charge them for the service. Similarly, individuals seeking care that might be stigmatized in their community, e.g., HIV testing or reproductive health services, can search and find care discreetly. Care service discovery relies on systems that can geospatially identify facilities, link providers with facilities, and link information on services offered with the providers and facilities.

Illustrative Functional Applications:

- Mobile or web applications that allow users to find facilities based on location and services offered. These could be further enriched by applications that support citizen-based reporting, capturing feedback and quality ratings, which can be linked to facilities and providers.

These functions are enabled by **canonical registries** of geographic information system (GIS) linked facilities and providers, enriched with information on the services provided at each facility. For the registries to share information with each other and with the search application, **interoperability platforms and standards** are necessary. **Health data security frameworks** should guide the presentation of care services information that may be politically sensitive or stigmatizing, such as reproductive health or HIV services. Care services discovery applications would also need to leverage **foundational DPI**, such as open mapping infrastructure to facilitate geo-spatially linked facilities and services as well as technology infrastructure to provide connectivity.

1.1.2 Desired Functionality: Continuity of Care

Digital Health Interventions: 1.4 – Client Health Tracking, 2.2 – Client Health Records, 4.2 – Data Coding, 4.3 – Data Location Mapping, 4.4 – Data Exchange and Interoperability

Continuity of care, in this context, refers to the availability of a longitudinal client health record and the ability to share this health information across public and private facilities and providers to ensure that the client is provided with appropriate care and services. As electronic medical record (EMR) systems in LMICs are often implemented as stand-alone, disconnected instances, and there may be multiple solutions used to capture clinical data, a central shared health record (SHR) service is necessary to enable the storage, linkage, and sharing of these clinical data for client care. Records from clinical information systems, including laboratory, pharmacy, radiology, and emergency service systems that may not yet be integrated with an EMR can be included in the SHR, which can provide access to the most important clinical information for a client in the form of a Summary of Care Record or International Patient Summary. Access to the SHR reduces duplicative lab testing and imaging and allows providers to continue treatment for a client across different services and, potentially, across international borders.

Person-centered monitoring using SHRs reduces the need to capture service indicators specifically for reporting, lightening the burden on health workers and improving client experience and care. Further, electronic records that are accessible to clients empower them to take ownership of their health data and give consent to others to view and use their data for self-care. With client consent, these centralized records can be used for research purposes and as training data to develop ML and AI models. In the absence of a fully functional and transactional SHR, records from clinical systems can be contributed to central data repositories, where they can be linked and aggregated to enable visualization and analytics for program management and research.

Illustrative Functional Applications:

- Provider-facing client health record with a view of client interactions with the health system, medication history, and insurance details (e.g., EMR connected to an SHR).
- Client-facing applications with access to shared records of the clients' interactions with the health system, enabling clients to manage consent to share their data.

A client-centric data repository that functions as an SHR service requires digitization and scale of point-of-service data across all health programs, including medical records, laboratory information, and details of prescribed and dispensed medication. Provider-facing applications interface with the SHR through standards-based **interoperability platforms**. Data need to be linked using a unique client identifier issued and maintained by a **canonical client registry**, following **health data security frameworks** to ensure that client health information is stored securely and shared with consent. Individual providers, facilities, and medications need to be listed and managed in respective **canonical registrie**s. Clinical observations, diagnoses, and laboratory tests and results need to be coded using common terminology, and data need to be securely and reliably exchanged with central services using **health messaging standards** such as Health Level 7 Fast Health Interoperability Resources (HL7 FHIR). Mature national governance policies, frameworks, and strategies, including an enterprise architecture, are critical to ensure conformance to this approach. Finally, technology infrastructure requirements include reliable connectivity and secure cloud hosting for SHR.

1.1.3 Desired Functionality: Access to Medication

WHO Classifications: 1.4 – Personal Health Tracking, 2.9 – Prescription and Medication Management, 3.2 – Supply Chain Management, 4.4 – Data Exchange and Interoperability

Access to medication includes ensuring that medications are safe, of good quality, and consistently in stock and available to clients. Stock-outs of essential medicines and other drugs at all levels of the health system are common in LMICs, and a substantial proportion of available medication is falsified or substandard. Recent estimates suggest nearly two billion people do not have access to the life-saving medication that they need.

Illustrative Functional Applications:

- Data visualization application for monitoring stock in near real-time.
- Product verification application that checks for product authenticity and expiration.

To provide visibility into stock levels at warehouses and facilities, ensure that medication is authentic and not expired, and enable supply planning and pre-positioning of stock based on demand, data need to be linked to **canonical facility and product registries** and shared between **scaled health services** such as warehouse management systems (WMS), electronic logistics management systems (eLMIS), and health management information systems (HMIS) via **interoperability platforms** using **health messaging standards**. Analyzing historical data to estimate future demand and plan forward distribution would be facilitated by a **central data repository** inclusive of data from WMS, eLMIS, and HMIS. This capability can be further strengthened by introducing machine-learning-based demand forecasting, which leverages open, representative data sets (potentially from scaled digital health services) to train ML models and open AI models that may be available to retrain on locally representative data. New health data governance may need to be considered to guide the sharing and use of data for predictive analytics.

1.1.4 Desired Functionality: Precision Public Healthcare

WHO Classifications: 1.1 – Targeted Client Communication, 4.1 – Data Collection, Management, and Use

Precision health has gained significant momentum in high-income countries, enabling medical decisions, interventions, and treatments to be tailored to the individual characteristics of each client. Leveraging predictive modeling capabilities, healthcare providers can deliver more accurate diagnoses and personalized treatment plans, and health interventions can be directed based on a more precise understanding of client risks and behavioral preferences. For example, health promotion and behavior change messaging can be tailored to content that will resonate most with a client, or, alternatively, be directed to a particular client based on the likelihood of that client responding. Segmentation algorithms can quickly identify clients as one of several distinct archetypes and provide each client with the information most relevant for them. Similarly, machine learning models may help identify the interventions that work best for clients with specific vulnerabilities and can inform tailored care plans based on those specific needs. This moves towards more personalized care, increasing the relevance of information and care given to individuals, creating more value for them, and strengthening their engagement with the healthcare system.

Illustrative Functional Applications:

- Community health worker (CHW) application that runs a “segmentation module” on top of community health information systems (CHIS) to advise on which of several possible interventions may work best for clients given their specific characteristics.

The development of machine learning models for market segmentation will require additional DPGs and building blocks that can leverage and integrate with underlying **scaled health services** such as a CHIS. The services can be a source of representative training data and give rise to **open AI models**. Storing such data in a **central data repository** separate from operational data will facilitate its use for analytics. Developing predictive models for precision health will require additional open data to develop training data sets constructed to be representative of local context. These data would need to be managed through **health data security frameworks** to ensure the implementation of appropriate safeguards when working with, sharing, and analyzing individual-level training data. To operationalize predictive models within the digital health systems, open AI models will need to leverage underlying infrastructure to work in conjunction with **client registries** and **shared health record services** connected to the specific point of service EMRs or CHIS that providers use. Training of models may require underlying infrastructure for cloud computing and connectivity.

Use of applications that integrate predictive AI models in service delivery will also introduce a need for new governance policies, frameworks, and strategies to ensure that models meet agreed-upon performance standards, documentation requirements, and evidence of effectiveness as well as having licensing structures that facilitate reuse when appropriate.

1.1.5 Desired Functionality: Digital Payments and Vouchers for Health Services

WHO Classifications: 1.7 – Client Financial Transactions, 3.5 – Health Financing, 4.2 – Data Coding

In LMICs, many individuals lack access to formal financial services and rely on cash-based transactions. These individuals are unable to access electronic payments and digital vouchers for health services such as COVID-19 relief payments, remote payment for healthcare interventions, online purchasing of medication, and access to health insurance programs and targeted subsidy schemes. In many countries, especially those moving towards middle-income status and taking on an increasing portion of health financing domestically, ministries of health require copayments for certain treatments and medications. Additionally, some clients may prefer and have means to access private-sector services, which require payment, even where public services are free. Access to digital payment mechanisms enables vulnerable populations to participate in the digital health ecosystem and reduces administrative burdens and the risk of fraud. Digital platforms can also be leveraged for conditional cash transfer programs, where individuals are provided with vouchers or cash transfers to incentivize health-seeking behaviors such as antenatal clinic visits and immunization. In many countries, vouchers are provided to individuals in order to access ambulance services and other interventions.

Illustrative Functional Applications:

- Digital financial services integrated with an insurance management information system that enables clients to digitally co-pay for services or access entitlements.
- Health insurance application that verifies service eligibility and provides electronic voucher for service.

Digital payments in the health sector depend on the foundational, non-health-specific DPI capabilities of digital identity, digital payment, and secure digital data exchange. Governance frameworks such as national digital enterprise architectures can guide how DPI-H interfaces with these foundational DPI. DPI can engage with DPI-H through a **services gateway** to access **canonical client and provider registries** to verify identities and enable bi-directional payments, following **health data security frameworks**. **Scaled digital health services** such as SHRs and insurance information management systems need to exchange data with the digital payment systems to coordinate and record payments, ideally through standards-based **interoperability platforms**. Linkage to national telecoms infrastructure will enable payers and beneficiaries of the financial transactions to receive direct messages, which is particularly important for the provision of vouchers.

1.2 User Journey 2: Improving Healthcare Provider Performance and Efficiency

In LMICs, healthcare workers (HCWs) perform multiple roles in service delivery in resource-constrained settings. Insufficient support, training, and capacity of these frontline workers lead to mental and physical health issues that impact the health system and its clients. This user journey is motivated by critical challenges with healthcare worker capacity as well as the need to improve efficiency and quality of care provided.

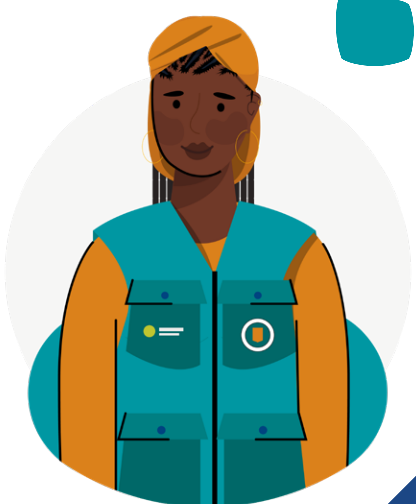

> “As a healthcare worker, I want digital apps and tools that I can trust to save me time – reducing the time I spend entering data and following up with clients to provide basic information. I’ve noticed there are many occasions where clients do not need to physically be at the clinic for me to assist them. Having access to a telemedicine service would enable clients to access a wider network of care providers and lessen demands on my time and on already-congested clinics.
>
> When I do care for clients in person, I want to make sure I’m spending time where it will really make a difference. I’d like to know which clients will benefit the most from my care, so I can prioritize my time.
>
> Lastly, as I am increasingly asked to take on more complex tasks. I want access to engaging, relevant training materials, supportive supervision, and clinical decision support to help me provide better quality care.”

This user journey includes several functional requirements that will improve healthcare worker performance and, potentially, sustainability in their roles. These include **streamlined data capture and messaging**, **remote consultations**, the ability to use **risk-based client prioritization**, and access to **online HCW training, support, and guidance**. Each of these functional requirements relies on several underlying DPI-H components, as described below.

1.2.1 Desired Functionality: Streamlined Data Capture and Messaging

WHO Classifications: 2.1 – Identification and Registration, 2.2. – Client Health Records, 2.7 – Health Worker Activity Planning and Scheduling, 4.2 – Data Coding, 4.4 – Data Exchange and Interoperability

HCW are consistently overburdened with high client loads, time-consuming administrative tasks, and frequent need for redundant data collection. Removing the need to re-enter data that have already been collected could be achieved by linking client information through a shared health record. The ability to directly message clients from a client records-management system can also reduce time by allowing HCW or CHW to quickly notify clients of results and send reminders and other supportive care messages directly to clients without having to make separate calls.

Illustrative Functional Applications:

- CHW application connected to an SHR, embedded messaging capability to message clients directly from the app.

Reducing redundant data entry requires applications that leverage a unique identifier for each client and a **canonical client registry, scaled digital health services** such as a shared health record in which client data is recorded, and robust, secure, functional point-of-service applications. To ensure that data can move from a shared record to an approved point-of care application, **interoperability platforms and standards** are needed. **Health data security frameworks** will ensure that data are managed safely and securely.

Direct messaging to clients and health workers relies on unique identification managed through **canonical client and health worker registries**, and **interoperability platforms and standards** allow components to participate in a health information exchange. In addition, access to underlying telecommunications infrastructure and messaging DPGs though a **service gateway** would enable direct messaging and/or calls without having to manage individual telecom channels. Curated messaging content can be leveraged by messaging services to provide relevant information to recipients.

1.2.2. Desired Functionality: Remote Consultation

WHO Classifications: 2.4 – Telemedicine, 4.1 – Data Collection, Management, and Use (for use of chat interface on CHW app)

Remote consultation refers to a client's access to a healthcare provider or nurse via a telephone or video conferencing platform. Remote care can reduce congestion and waiting times, lessen the burden on clinic-based health workers, and benefit clients. Tools that enable client referrals between primary health workers and specialist physicians and surgeons and provide secure channels for consultation between health workers and specialists save time and support less-skilled HCWs in providing appropriate care to clients. Digital tools can also support asynchronous care, easing the scheduling burden on HCWs, through video-observed tuberculosis therapy, for example. Remote care may also be provided through a virtual chatbot interface that can reply to simple questions as well as direct clients to nurses or clinicians when needed, another way to optimize HCW time.

Illustrative Functional Applications:

- Voice, video, and chat-based consultation apps.
- Service directory from which to locate telehealth providers.
- Client rating applications to show client feedback in the service directory.

Remote care applications require **canonical provider registries** to locate qualified providers as well as **scaled health services** such as SHRs and laboratory information systems (LISs) to provide and access health records and lab services. **Interoperability platforms and standards** are necessary to exchange data with a shared health record, supported with **health data security frameworks** to ensure that information shared virtually is protected. Remote consultation requires secure **underlying technology infrastructure** including video and audio exchange, likely with direct client messaging capability (as described above). Governance also plays an important role in setting standards for telecare and ensuring providers are approved and certified.

If remote consultation is provided through a chat interface, it may be augmented using natural-language processing (NLP) capabilities to develop more intuitive chat capabilities in local languages.[1] This relies on many **Supporting DPGs** to develop context-appropriate NLP (e.g., text corpora in local languages[1] to train models, existing chat platforms, no-code platforms to train bots, etc.). As with other introductions of AI and predictive analytic tools, it introduces new dimensions to governance guidelines to evaluate AI model performance and ensure that content is contextually appropriate and factually accurate.

1.2.3 Desired Functionality: Risk-Based Client Prioritization

WHO Classification: 2.3 – Healthcare Provider Decision Support, 2.7 – Healthcare Worker Planning and Scheduling, 4.1 – Data Collection, Management, and Use

HCW and CHW time may be further optimized by using machine learning models to identify clients with the highest need for care. For example, a CHW can prioritize clients for adherence support intervention based on their risk of experiencing interruption in treatment. This functionality is currently emerging for HIV and TB clients in multiple LMICs, including Nigeria, Mozambique, Kenya, and India. In these cases, client-level data from routine service delivery, and occasional additional open data, are used to train machine learning models to produce a treatment interruption risk score. This model can be run on top of either a scaled health service such as a SHR or on an individual instance of an EMR such as OpenMRS.

Illustrative Functional Applications:

- CHW apps with an embedded risk-scoring model to prioritize client care plans based on risk.

[1] *For chatbots to work with local languages, significant amounts of digital text (text corpora) in local languages will be needed for natural language processing models. This is a generalized need and not necessarily health-specific, though text corpora to support NLP for local languages would be made stronger by including language generated in health contexts. Several initiatives to strengthen NLP for local languages are on-going, including Lacuna Fund and Mozilla Common Voice project.*

To develop such machine-learning based models, there is a need for sufficient quality and quantity of client-level training data. If available, **open AI models**, and documentation to support tailoring and reuse of existing models to local contexts, may eliminate the need to repeatedly create new models. Training data may come from scaled point-of-service applications, though this is greatly facilitated by a **central data repository** collating data from multiple applications and **scaled health service** data streams, as well as additional **open data**. Health data governance considerations should guide the sharing and use of this data for the development of predictive models as well as use of pre-trained models to ensure they meet agreed-upon standards of performance and demonstrated benefit.

1.2.4 Desired Functionality: Electronic HCW Guidance and Training

WHO Classification: 2.8 – Healthcare Provider Training, 3.1 – Human Resource Management

The continuous and increasing exodus of health professionals from LMICs poses a particular threat to health systems. As such, remaining frontline workers are required to perform more and more complex tasks. Access to global, national, and regional curated health content, as well as supportive supervision from healthcare supervisors, is vital to ensure health workers are referencing the correct literature and clinical guides. Providing this content will rely on decision support tools to aid HCWs and CHWs in practice, content management platforms to organize appropriate guidance, and online training courses to support workers in taking on new tasks.

Content management platforms streamline content creation, management, and indexing and enable economies of scale by allowing additional content creators to use the same platform. Online training platforms can then go beyond content management to provide structured interaction with providers. Training platforms often consume and enrich data from content management platforms and, similarly, enable economies of scale by training more professionals with minimal incremental investment.

Illustrative Functional Applications:

- Decision support tools and job aid applications to give providers structured support with care plans and diagnoses that follow approved guidance and best practices.
- Dynamic data dashboards to analyze and visualize data from HCW apps and information systems to support management and supervision.
- Online content management platforms to collate and facilitate searchable, up to date guidance.
- eLearning courses that offer economies of scale for training, upskilling, and maintaining skills.

Decision support tools developed in an ecosystem with DPI-H would leverage **scaled health services**, such as an SHR, and add additional workflows based on **open content**, such as clinical guidance or global best practices. Content creation for such tools needs to follow standards set through governance policies, frameworks, and strategies to ensure content is appropriate and that decision-support workflows are suitable for **scaled health services** and supporting point-of-service functional applications, for example using the SMART Guideline approach (Standards-based, Machine-readable, Adaptive, Requirements-based, and Testable). **Open AI models** and analytics platforms leveraging **central data repositories** could enable feedback and data-driven management and supervision for quality improvement.

Content management applications and eLearning courses will depend on domain-agnostic content management or training platforms such as a wiki page or Moodle training platform.[2] These platforms could contribute to HCW performance support by sharing data with health-specific infrastructure, namely, **provider registries via an interoperability platform**, to produce individual training profiles and records of training outcomes.

1.3 User Journey 3: Strengthening Surveillance and Outbreak Management

Outbreaks of Ebola and other severe infectious diseases, including the recent global COVID-19 pandemic, have highlighted the need for effective public health systems with near real-time data capabilities that can detect, respond to, and contain outbreaks to reduce mortality and morbidity. The ability to see data in near real-time and, eventually, predict outbreaks can enable earlier preventive action, inform timely communication to health officials and the public, and enable pre-positioning of critical supplies and resource shifting throughout healthcare facilities. The third user journey focuses on health system managers and a series of functional needs involving population-level data and system-level management and response.

"As a health system manager, I want to be able to see all the data that might help me understand emerging disease outbreaks, including surveillance data, global and local routine reporting data, and supply chain data. Given the increasing importance of One Health approaches in our changing climate, I want to see animal and weather data, too. Ideally, we could overcome delays in availability of routine data by using models to get real-time estimates of cases and alert relevant health facility managers with timely updates.

In the future, I would like to leverage new kinds of data to predict outbreaks even before we see cases rise through routine health reporting and surveillance, and those predictions could even inform logistics planning and forecasting. But I would need to have confidence in the predictions and have a way to disseminate the information."

This user journey includes three primary functionalities: **outbreak analytics**, **direct communication between health officials and healthcare workers,** and **outbreak forecasting and alert capabilities.**

[2] This analysis considers domain-agnostic training and content management platforms part of "innovation-supporting DPGs and building blocks" (explained more in 1.4.3), though they might also be considered part of the domain-specific infrastructure for education, as a core platform for sharing and exchanging knowledge.

1.3.1 Desired Functionality: Outbreak Analytics

WHO Classifications: 3.2 – Supply Chain Management, 4.1 – Data Collection, Management, and Use, 4.2 – Data Coding, 4.3 – Location Mapping, 4.4 – Interoperability and Data Exchange

To proactively manage disease outbreaks, health system managers need a holistic view of current case counts across a wide range of diseases affecting population health. However, many reporting systems and surveillance systems are disease-specific and siloed, making it difficult for health system managers to have a comprehensive view of what is happening in communities and at facilities. Further, supply chain data is often siloed from health information, making it difficult to rapidly adjust supplies to meet demand. A useful outbreak analytics platform requires integration of surveillance data, routine reporting data, and supply chain data in one environment for analysis and visualization. As climate change continues to contribute to the emergence of new pathogens and shifts endemic zones for known pathogens, being able to bring in non-health information streams, such as data from animal surveillance and meteorological systems, will be increasingly important.

Illustrative Functional Applications:

- Integrated disease surveillance response (IDSR) specific data repositories and analytics platforms to collect and visualize epidemiological data.
- Case-based surveillance and management tools to enable collection, analysis, and linkage of case information including contact tracing.

Pulling in disease information by location requires GIS-linked **canonical facility registries**. **Scaled digital health services**, such as global, integrated disease surveillance systems and disease-specific global reporting systems provide core indicators.

A **central data repository** can facilitate collating aggregate data from multiple sources into one environment and enable visualization of the data, and **interoperability platforms and standards** are needed to enable the harmonization of these data. Incorporating climate or animal disease data following a One Health approach requires access to other databases outside the health system, which can be facilitated through a **services gateway** to an underlying data exchange platform. **Health data security frameworks** are needed to ensure the responsible use of aggregate population-level data and demographically identifiable data (such as GIS-linked case counts).

1.3.2 Desired Functionality: Direct Communication Between Health Officials and HCWs

WHO Classification: 2.5 – Healthcare Provider Communication, 4.4 – Data Exchange and Interoperability

When suspected or confirmed outbreaks are identified through surveillance and reporting dashboards, health system managers need to alert key stakeholders in the system, including facility managers, logistics supply managers, public health communications managers, and frontline health workers. Official communications need to be rapidly conveyed to the workforce to manage emergency disease response protocols, flag changes to guidelines, and distribute curated informational content.

Illustrative Functional Applications:

- Surveillance dashboard with embedded messaging capability to enable real-time notification and communication with health officials.

A communications system between health officials and staff would require a **canonical provider registry** of health system managers, administrative staff, and health workers that includes reporting hierarchies and up-to-date contact information. The communications system also needs to leverage **interoperability platforms and standards** to share information across **canonical registries**, and access to **foundational technology infrastructure** such as a secure telecommunications platform for messaging both for one-to-one and group communications. This system would be strengthened by supporting DPGs, such as **open content** and guidelines that follow global recommendations on communications for outbreak response.

1.3.3 Desired Functionality: Outbreak Forecasting and Alert Capabilities

WHO Classification: 4.1 – Data Collection, Management, and Use, 4.2 – Data Coding, 4.3 – Location Mapping, 4.4 – Data Exchange and Interoperability, 2.5 – Healthcare Provider Communication

Advanced outbreak forecasting is necessary to identify emergent epidemics before they are measurable through routine reporting or surveillance systems. Machine learning and predictive modeling approaches are increasingly applied to develop models that support these early warning systems so that outbreaks can be predicted and contained quickly.

Illustrative Functional Applications:

- Forecasting application leveraging AI-enabled and other predictive analytics to enable real-time case estimates and advanced prediction of emergent outbreaks.

Developing outbreak forecasting models requires large quantities of training data generated through underlying infrastructure such as **scaled digital health services** including logistics, health management information, and disease surveillance systems. Case-based surveillance may also draw upon SHR, if available. These services require **interoperability platforms and standards** to exchange data collected through approved point-of-service health information system (HIS) applications.

Forecasting models may also benefit from additional **open data** (including historical data, climate data, social and demographic data) and **open AI models** that can be retrained on context-appropriate data. Once models are developed, **central data repositories** can facilitate the collation and management of source datasets that will be used to train and process data using approved predictive models. The forecasted outputs can be utilized as inputs to other models, such as for supply chain demand forecasts. Using data in this way requires an underlying central data repository to be optimized for analytical workloads, separate from operational data stores. Together, the data used for training, open AI models, and health analytics platforms are supporting global goods, DPGs, or building blocks that leverage the underlying scaled digital health services, interoperability layer, and central data repository.

To facilitate trust in these predictive models, their outputs should be subject to a review process, creating new governance needs within a health system. Countries will need to adopt and operationalize emerging AI governance to establish norms for documenting how models work, the context in which they are originally developed, and known limitations. Governance policies and frameworks should also guide model review processes to ensure they meet agreed-upon standards for performance and lay out thresholds of model performance to be met before triggering broader alerts. A health system manager would be able to plan response protocols earlier, communicating with other officials and HCW using **direct messaging platforms** that support the communication functionality noted above.

Table 1: **Summary of Infrastructure Components Required by Desired Functionalities Across User Journeys**

Layer	Category	Component	Care Service Discovery	Continuity of Care	Access to Medication	Precision Public Health	Digital Payment & Vouchers	Streamlined Data Capture and Messaging	Remote Consultation	Risk-Based Client Prioritization	Online HCW Training, Support and Guidance	Integrated Outbreak Analytics	Direct Communication: Health Officials and HCW	Outbreak Forecasting and Alert Capabilities
Supporting DPGs and Functional Applications	Illustrative Supporting DPGs and Global Goods	Open Curated Data										X		X
		Open AI Models				X				X				X
		Open Curated Content				X								
DPI-H	Health Data Security Framework	Consent, Security, and Authorization Protocols	X	X	X	X	X	X	X	X	X	X	X	X
	Canonical Registries	Facility Registry	X	X	X		X	X	X	X		X	X	X
		Product Registry	X		X			X		X				X
		Client Registry		X		X		X	X			X		
		Health Worker Registry		X			X	X	X		X		X	X
		Terminology Service		X		X	X	X	X			X		X
	Health Domain Interoperability	Interoperability Layer, Messaging Standards & Orchestration	X	X	X	X	X	X	X	X	X	X	X	X
	Scaled Health Services	SHR		X	X	X		X	X	X		X		X
		EIR		X	X			X	X	X				X
		LMIS		X	X				X	X				X
		LIS		X	X			X	X	X		X		X
		HRIS		X				X	X	X	X	X		X
		IDSR										X		X
		IMIS		X	X		X	X	X					
	Central Data Repositories	Data Warehouses and Analytics				X				X	X	X		X
Foundational DPI and Technology Infrastructure	Foundational Digital Public Infrastructure	Mapping Infrastructure	X									X	X	X
		Payments Infrastructure					X							
		Cross-Sector Data Exchange Infrastructure												X
		Identity Infrastructure					X	X	X		X		X	
	Technology Infrastructure	Connectivity	X	X	X	X	X	X	X	X	X	X	X	X
		GSM Messaging Networks				X	X	X	X			X	X	X
		Cloud Computing				X			X					X

1.4. A Generalized Framework for DPI-H

The user-journeys explored above show how diverse functional applications are supported by common underlying components (Table 1). This section codifies the identified DPI-H needs into a general framework. Following the layers of the conceptual model introduced in Figure 1, the detailed DPI-H framework (Figure 2) shows that foundational DPI and technology infrastructure components can be leveraged through a services gateway. These foundational DPI are separate from DPI-H but essential to support DPI-H and an ecosystem of applications. DPI-H components include the basic functionalities of health data security along with health information exchange: registries, health domain interoperability, scaled health services and central data repositories. These components are required by a large majority of the functional applications explored through the user journeys. Finally, above core DPI-H, this analysis identifies several Supporting DPGs, GGs, and building blocks that feature in multiple applications, yet have content-specificity. They enrich the basic functionalities of DPI-H to give rise to multiple functional applications but are not themselves providing infrastructural functionality.

Figure 2: Conceptual Model for DPI-H

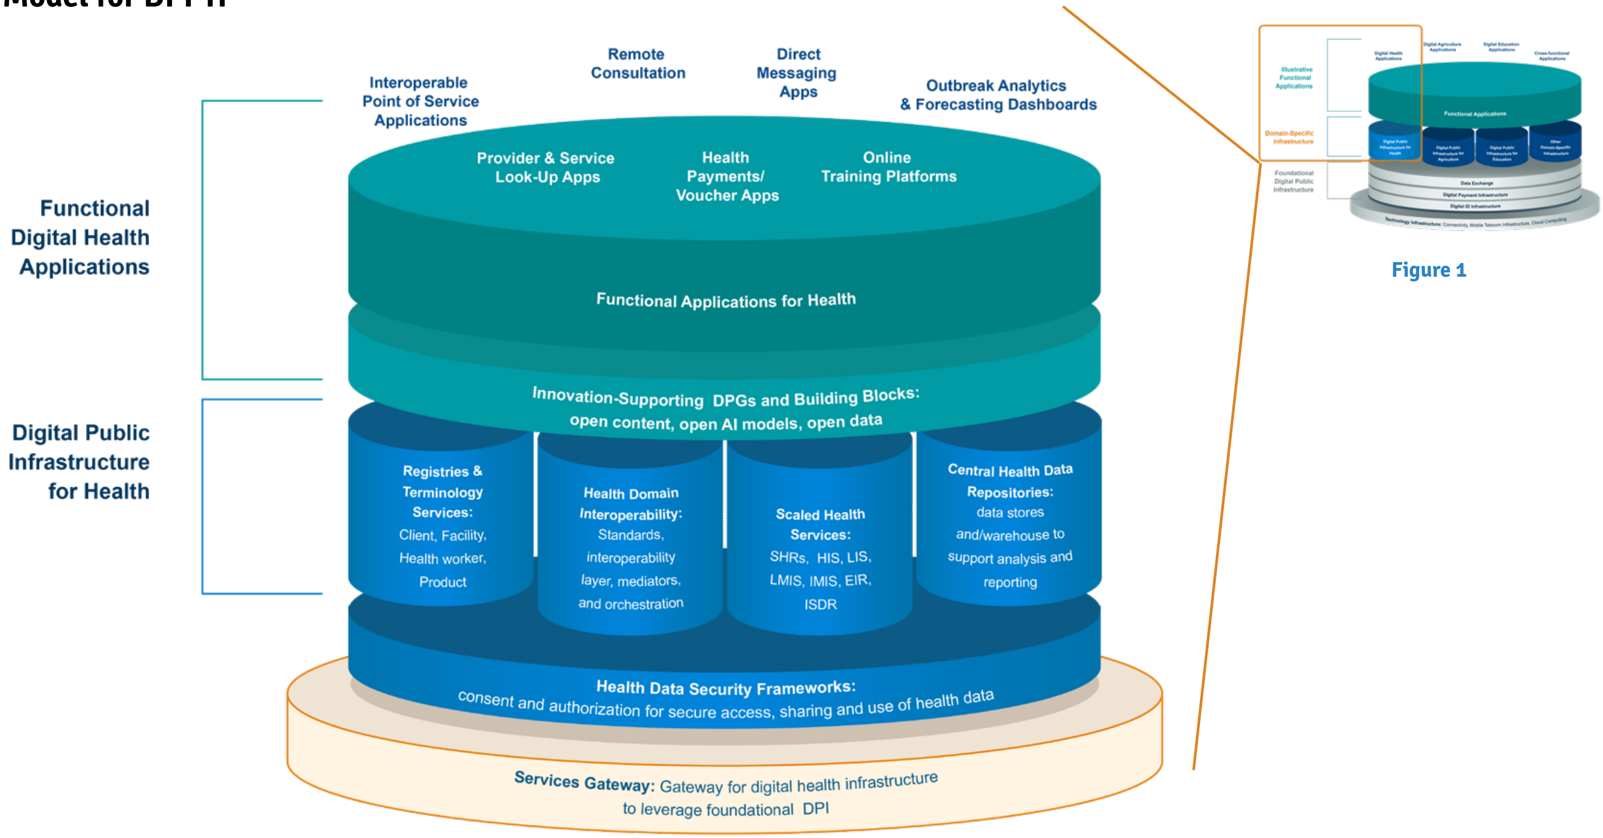

1.4.1 Cross-Domain Base Layers: Foundational DPI and Technology Infrastructure

In this conceptualization of DPI-H, the health-specific infrastructure elements leverage foundational DPI like payments, identity systems, and data exchange with other domains. For example, health insurance payments and voucher transactions should be able to leverage underlying digital payments infrastructure, care services discovery applications may need to access geospatial data and mapping infrastructure to layer health facility coordinates onto a base map layer, and surveillance systems may benefit from exchanging data with systems outside the health sector, such as climate and agricultural monitoring systems.

Utilization of these infrastructure components is represented through the Digital Services Gateway layer, pictured at the bottom of Figure 2.

While not health-specific, most domain-specific infrastructure will need a gateway through which to access and leverage other domain-specific infrastructure as well as underlying DPI and technology infrastructure.

Similarly, DPI and domain-specific infrastructures like DPI-H rely on some level of technology infrastructure. Connectivity, access to mobile network infrastructure for basic mobile messaging capability, and, for analytics use cases, cloud computing infrastructure similarly will need to be part of national DPI to support DPI-H at scale.

1.4.2 Domain-Specific Layers: Digital Public Infrastructure for Health

Domain-specific infrastructure for the health sector includes the multiple digital components required by the identified use cases and those which can be combined with other components or DPGs to produce functional applications. These core components are described below with several illustrative products to provide concrete examples of the function, with the caveat that individual products may serve more than one functional category, and any individual product, by itself, is generally not infrastructure – it is an example of a product that performs an infrastructural function but may not always be deployed in a way that enables it to serve as infrastructure.

Canonical Registries refer to services that uniquely identify key entities, facilities, people, and products in a health information exchange. These include client, facility, product, health worker, and provider registries (e.g., OpenCR, GOFR, PCMT and iHRIS, respectively) as well as terminology services (e.g., OpenConceptLab). Registries list canonical, uniquely identified data that are relatively static but can also be enriched with ancillary data that changes more rapidly, such as contact information for clients, providers, and facilities to enable communications with these entities, information about the services that a facility offers, and qualifications and skills of health workers. Terminology services allow real-time mapping of entities as part of health information exchange transactions.

Interoperability Platforms and Standards refer to technology and standards that allow health information to be electronically shared across information systems. Messaging standards (e.g., HL7 FHIR) and interoperability layers (e.g., OpenHIM) are central components of a health information exchange used, for example, to ensure data elements can be exchanged and accurately interpreted across systems.

Scaled Digital Health Services refer to national or central-level services with centralized operational databases and standards-based open APIs that support point-of-service applications and enable sharing and combining of data across different point-of-service applications. This category includes shared health record (SHR) services, immunization registry services, logistics management information services, health management information services, health financing services, and disease surveillance systems. National instances of DHIS2, iHRIS, and OpenLMIS, and FHIR servers that form part of a health information exchange can be classified as scaled digital health services.

Central Data Repositories contain health data from multiple source systems. The category includes a very broad range of technologies, distinct from a shared health record service due to their predominantly offline functionality and focus on analytics. Many LMICs are implementing these technologies to consolidate data from multiple source systems, transforming data into a consistent format to enable analytical queries and visualizations to inform public health decision-makers.

These repositories commonly function as offline data warehouses, housing copies of data from other systems for reporting and analytical purposes e.g., OpenHexa and Harmony, or custom-developed national scale data warehouses such as Zambia National Data Warehouse or South Africa's Western Cape Provincial Health Data Centre. Software may be health-specific (Harmony), or generic (Microsoft SQL Server), and can provide built-in analytics or utilize generic visualization tools such as PowerBI and Tableau. Even where an operational data repository for a scaled health service such as a shared health record exists, data must be replicated from the transactional, operational databases into an analytical data store or data warehouse for reporting e.g., PowerBI FHIR connectors for FHIR servers, and Google's FHIR Data Pipes, to ensure that reporting operations do not interfere with real-time data exchange.

Health Data Security Frameworks reference the technology and related protocols that operationalize data protection requirements, security measures, and safeguards for the use of individual and population-level data. These include services for managing consent, such as a centralized consent framework to enable fiduciary control of health data and services for auditing and securing these data, as well as tools for linking, de-identifying, and re-identifying client records for research purposes.

1.4.3 Cross-Domain Layer: Innovation-Supporting DPGs and Building Blocks

The basic functionalities provided by core DPI-H in isolation are necessary but not always sufficient to support the use cases identified through the user journeys. In multiple cases, the basic functionality of DPI-H needs to leverage supporting data, content, analytic platforms, and content management platforms to create a diverse range of applications, particularly those that rely on predictive analytics. This framework focuses on three categories of Supporting DPGs and building blocks that often complement core DPI-H: open curated data, open content, and AI models. These kinds of DPGs are often, but not always, health-specific, e.g., open curated data for outbreak analytics and One Health will often rely on climate, demographic, and other data sets that are cross-domain. Similarly, open health content may be paired with domain-agnostic content management platforms or learning platforms to create health-specific solutions.

Curated Health Content refers to curated informational content that would be applicable across multiple applications but, necessarily, has content specificity. In the use cases above, approved clinical guidelines support the upskilling of providers; sets of revised guidelines in specific but recurring situations such as outbreak response, commodity shortages, etc. are also important. WHO clinical guidelines, SMART guidelines, and approved best practices for public health intervention are examples of curated content sources. Content may be made more discoverable and usable through content management platforms such as wikis or when curated on learning platforms such as Moodle.

Curated Open Data can be used for business intelligence and predictive analytic use cases. With enabling data governance guidelines, open data should be posted with consideration for relevant metadata, clear terms of use, and information about the context in which the data was gathered along with known biases or limitations. Open data may be accessed through specific dataset management platforms, such as OCHA's Humanitarian Data Exchange, or as part of other content platforms, such as DIAL's Digital Impact Exchange.

Open AI Models include trained models that use machine learning algorithms and are available for retraining and reuse. The growth of open AI models introduces new considerations for the governance of their use in digital health systems. In addition to the models themselves, Open AI models considered as a Supporting DPG may need to have documentation describing how the model was trained, appropriate performance metrics, and descriptions of known limitations and biases. Ideally, open models are available in a standard such as PMML, enabling predictive models to be represented in a standard way, independent of which programming language is used to build them. In some cases, it may be appropriate to make key features for a given predictive problem openly available, even if a model is commercially licensed. Open AI models may be available on platforms such as GitHub or or other software repositories.

1.5 Ecosystem Enablers

The concept of DPI-H requires intentional planning and design. While it is possible for each component part of DPI or DPI-H to exist in isolation, it is only when they are implemented together as part of an architecture, following a governance model, to support other solutions to leverage their functionality or data that they begin to function as public infrastructure at scale. In order to move from isolated implementations to a reality in which the functionalities could be efficiently scaled at population level, there are numerous capacities needed in the broader ecosystem. These include leadership, governance, and regulatory capacity, health data governance capacity, and technical workforce capacity.

Leadership, Governance, and Regulatory Capacity. Realizing the concept of DPI depends upon committed leadership and capacity to develop national digital health strategies, guiding policies, health enterprise architectures, and costed plans to direct digital health investments. Further, there is a need for governance capacity to ensure that digital systems and goods comply with the standards, requirements, and architecture design put forward. As new capabilities like predictive analytics and AI enter the digital health space, additional leadership and governance capacity will need to be developed to evaluate and govern the use of predictive models in public health and health care.

Health Data Governance Capacity. Given the central role of data exchange and data use in the user journeys above, health data governance policies and capacities are critical to creating an ecosystem in which actors operationalize the protocols and safeguards that are increasingly stipulated in data protection laws. In the absence of legislation, organizational policies and best practices may also guide use around consent, authorization, and decisions regarding which data can be shared according to specific conditions. This is also reflected in the DPI-H component of Health Data Security Frameworks, which refers to the software tools needed to operationalize consent for data sharing and use and appropriate security controls on access and use of sensitive data.

Technical Workforce Capacity. Lastly, the ability to design and implement DPI-H, and benefit from it, relies on sufficient numbers of adequately skilled human resources in the ecosystem (Figure 3). This applies to the technical workforce capacity needed to draft relevant strategies, architectures, and frameworks as well as create and maintain relevant software used in a given country›s DPI-H. It also includes a broader workforce capacity in digital health to support an ecosystem of innovation and functional application development across public and private sectors that can leverage underlying DPI-H. These enabling factors are as important as the infrastructure components themselves in the process of bringing multiple DPGs and building blocks together to function as DPI-H in a country.

Figure 3: Ecosystem Enablers for Digital Public Infrastructure for Health

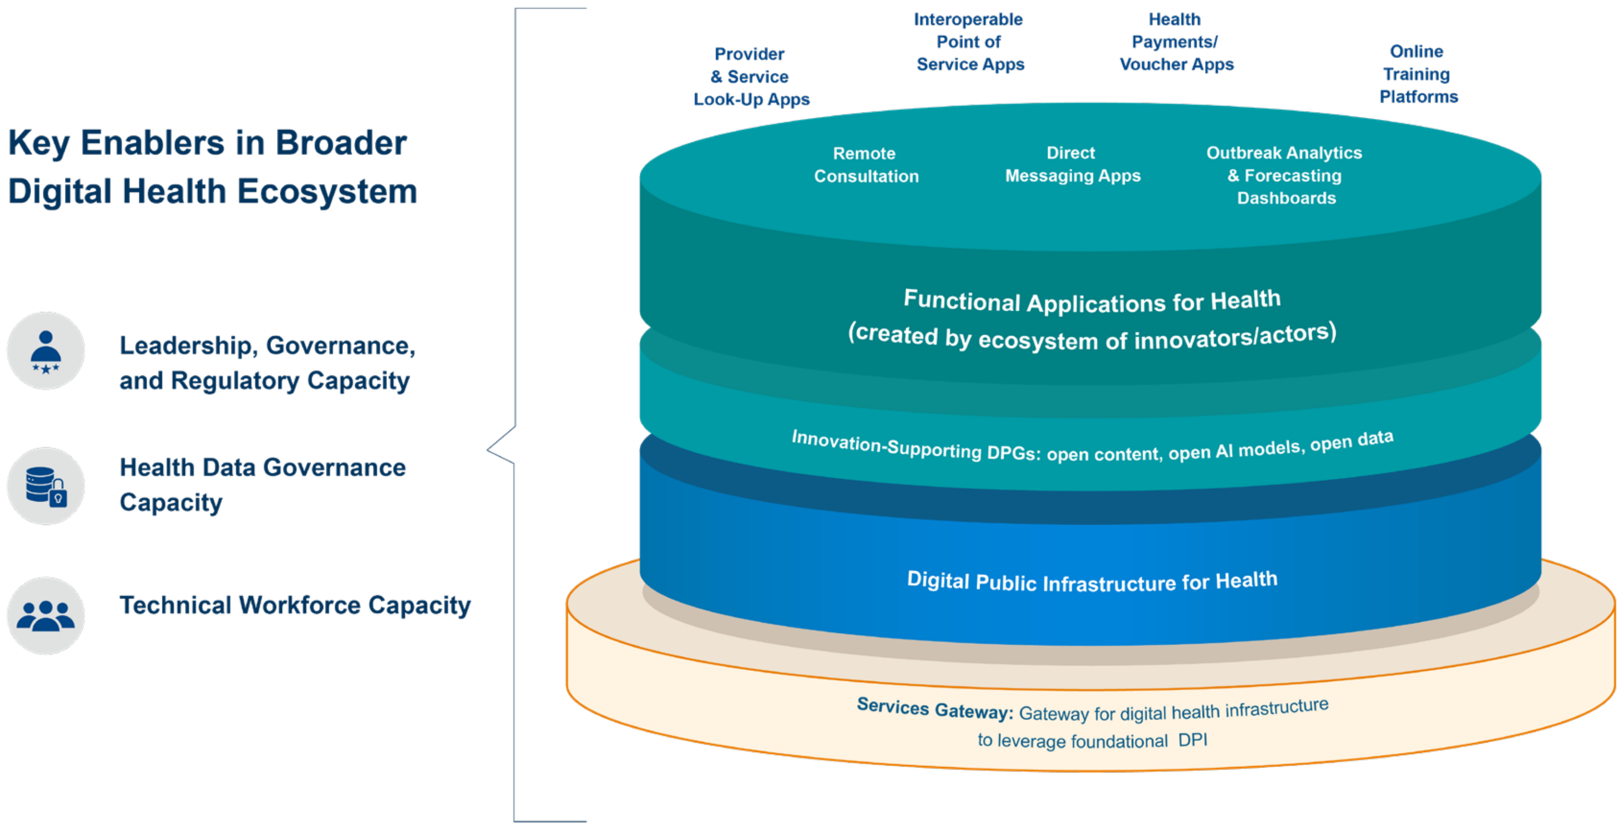

DPI-H in Practice

The National Health Service (NHS) in the United Kingdom offers one model of what DPI-H could look like in practice. It supports canonical registries, scaled health services, and consent protocols through an interoperability platform and utilizes dedicated data repositories to provide many of the desired functionalities envisioned in the user journeys. To accomplish this, the NHS is embracing the use of Fast Healthcare Interoperability Resources (FHIR) as a key standard for health data exchange. FHIR allows the NHS to facilitate interoperability between various healthcare systems, enabling the secure and efficient sharing of client information across different care settings and organizations. Alongside its FHIR implementation, the NHS is developing a robust data governance framework to ensure responsible and ethical handling of health data. This framework includes strict policies and procedures for data collection, storage, access, and sharing, emphasizing client privacy and confidentiality. Through the combination of FHIR and its data governance framework, the NHS will be able to harness the power of health data to deliver improved client care, research, and healthcare planning while maintaining a high level of data security and protection. Further, it supports robust, open documentation to enable many external organizations, public and private, to develop applications that can be approved to work with the NHS.

1.6 Insights from DPI-H Conceptual Framework

DPI-H Components are Interconnected

The user-journey approach identified clear patterns in combinations of digital systems that need to be implemented together to function as DPI-H. Across the various functionalities explored, canonical registries, health data exchange through an interoperability layer, scaled health services, central data repositories, and authorization and consent tools were repeatedly required. Central data repositories are included as core infrastructure due to the growing interest in using data collected from multiple and often non-interoperable point-of-service applications, scaled health services, and open or otherwise siloed data for the predictive analytics functionalities noted in the user journeys. While not every functionality will require all these components, almost every functionality in the user journeys requires more than one. The implication is that implementation, even when scaled, of one of these DPI-H components in isolation has limited value. When implemented together, at scale, as part of a functioning data exchange, a broader ecosystem of functional applications in the health domain can be supported. Implementing these components together is critical to realizing the potential of DPI-H.

Supporting DPGs and Building Blocks Complement DPI-H to Enable Desired Functionalities

Further, the user-journeys and DPI-H framework show there are multiple types of Supporting DPGs and building blocks that are important in achieving prioritized functionalities but that are dependent on one or more components in the DPI-H layer. For example, open content and content management for HCW training are high-priority categories for global goods and building blocks that will be important to support a stretched workforce. DPGs and building blocks in the form of open data and open AI models are needed to realize the emerging functionalities of personalized care, risk-based prioritization, and reliable outbreak forecasting. Accordingly, advances in the use of data and AI models for predictive analytics require new dimensions to governance policies, frameworks, and strategies to ensure that predictive models are trustworthy at scale. This may include requirements around performance reporting for predictive models and evidence generation to show their effectiveness and the validation of models.

DPI-H Enables Health Information Exchange

As a whole, there is a strong overlap between the DPI-H components identified through this process and the components of the OpenHIE framework. Indeed, a functioning health information exchange (HIE) would be a good example of DPI-H. As in the DPI-H framework outlined above, an HIE can connect a multitude of point-of-service applications (as well as other functional applications) to a data exchange layer to enable many of the envisioned applications in the user journeys.

Further, the DPI-H components identified through this user-journey approach generally align with the categories of DPI described by the Center for Digital Public Infrastructure (CPDI), which are envisioned to apply across sectors. Table 2 shows the alignment between the three frameworks. While there is not complete overlap with CPDI categories, the general alignment supports the notion that domain-specific DPI will be variations on similarly underlying categories rather than wholly new. Domain-specific considerations will necessitate some additional functionalities, in some cases creating solutions on top of foundational DPI and in other cases directly leveraging foundational DPI.

Table 2: **Alignment between DPI-H, OpenHIE, and CDPI DPI Categories**

Digital Public Infrastructure for Health	OpenHIE Element	Related CDPI DPI Category
Health Data Security Protocols	N/A **authentication included in Interoperability layer*	Signatures and Consent
Canonical Registries and Terminology Services	Registries	Identifiers and Registries
Health Domain Interoperability	Interoperability Layer	Data Sharing and Models **interoperability is recognized as a principle rather than a product in the CDPI framework, and data sharing and models refers to sharing of individual level data with consent*
Scaled Health Service	Business Services	Discovery and Fulfillment (of Business Services) **scaled health services could be conceptualized as the critical services to discover and fulfill in the health domain*
Central Data Repository	N/A	N/A
(Via Services Gateway)	N/A	Payments

Enterprise Architecture Approaches Support DPI-H Implementation

The concept of utilizing shared infrastructure and applications that can be leveraged to support many use cases is also highly consistent with enterprise architecture (EA) approaches, in which health enterprise architectures outline the business needs of an enterprise to guide the needed applications in the ecosystem. These applications are then mapped to information and technology layers to meet needs in an efficient, flexible way. This often entails shared use of common assets including those of a health information exchange, a central data warehouse, and underlying technology infrastructure. The overlap between DPI-H and EA approaches reinforces the value of health enterprise architectures as key governance assets that can guide the adoption of DPI-H in a specific country context.

Part 2:

Exploring Current Product Landscape

Part 2: Exploring Current Product Landscape

2.1 Product readiness to function at DPI-H

DPI-H will necessarily be implemented through applications – specific products that provide the infrastructural functionality. While DPIs ideally stem from digital public goods (DPGs), not all DPGs will be infrastructure. Rather, DPI-H will manifest as a subset of DPGs that have infrastructural properties: they perform a basic function, at scale, and can be leveraged by other products to create many different functional applications. This subset of DPGs becomes infrastructural when they are part of a transparent governance framework, in wide use, and enable other solutions to leverage their functionality or the data they host.

What are Digital Public Goods?

The Digital Public Goods Alliance defines DPGs as open software code, platforms, applications, open data, open AI models, open standards, and open content used to advance sustainable development outcomes. Health-related DPGs are also referred to as “global goods” and enable a wide range of digital health use cases and interventions in the service of health-related sustainable development goals.

Two bodies currently assess solutions for recognition as a digital public good or, in the health sector, a global good (GG). The Digital Public Goods Alliance (DPGA) and Digital Square assess prospective products against a set of standards to be awarded the status of DPG or GG by each body, respectively. While DPGs are referenced as examples of global goods that could serve as DPI-H if widely adopted and leveraged within a digital health system, this report defines components of DPI-H in terms of functions rather than specific software products.

“Building blocks” in this report refers to digital assets that are interoperable, provide a basic digital service at scale, and can be reused for multiple use cases and contexts.

Across LMIC health systems, there are several global goods and building blocks already in use with the potential to contribute to DPI-H. This section explores the existing landscape of products and provides insights into how they can serve as DPI-H.

This analysis focuses on a subset of products on the market that have been recognized as a DPG, global good, or recognized building block in the DIAL Digital Impact Exchange. Not all products in use have sought certification, and new products and approaches to achieving these functionalities will emerge over time. Thus, this review is not comprehensive, but it identifies a landscape of product candidates for countries interested in developing DPI-H in the next five to ten years.

To produce this analysis, the identified products were categorized according to the DPI-H framework components described above (Figure 4). Some products are present in all three lists, while others are only listed in one or two. The "total" in the charts below reflects the number of unique products in the category. Many products perform more than one function, which introduces some subjectivity into the categorization of products. This is intended to be a high-level overview of the DPG and building block product landscape and acknowledges some products may justifiably fall into more than one category.

At a high level, each of the components in the DPI-H layer has at least one associated global good or DPG-certified product, though these are not evenly distributed. Products that provide data governance protocols like consent and authorization are limited to only one in the global goods and DPG designation (OpenAttestation), though there are other products on the market that provide this service, either alone or as part of a more complex product. On the opposite end, there are many products that fall into the scaled health service and registry categories, which reflect the breadth of services in a typical health system.

Figures 5 and 6 break down the *scaled health services* and *canonical registries and terminology services components*, as there are multiple services and types of registries detailed within those categories.

Figure 4: DPG and Building Block Product Landscape by DPI-H Category

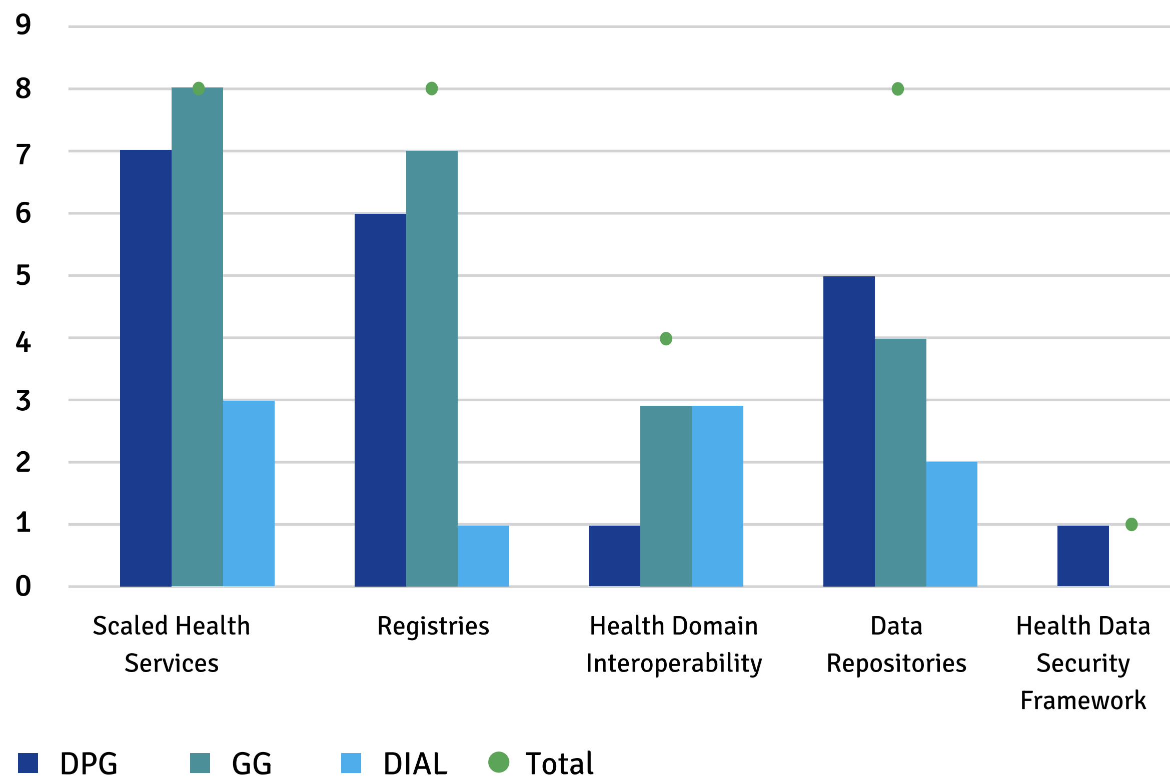

Figure 5: Product Landscape Breakdown for Scaled Health Services

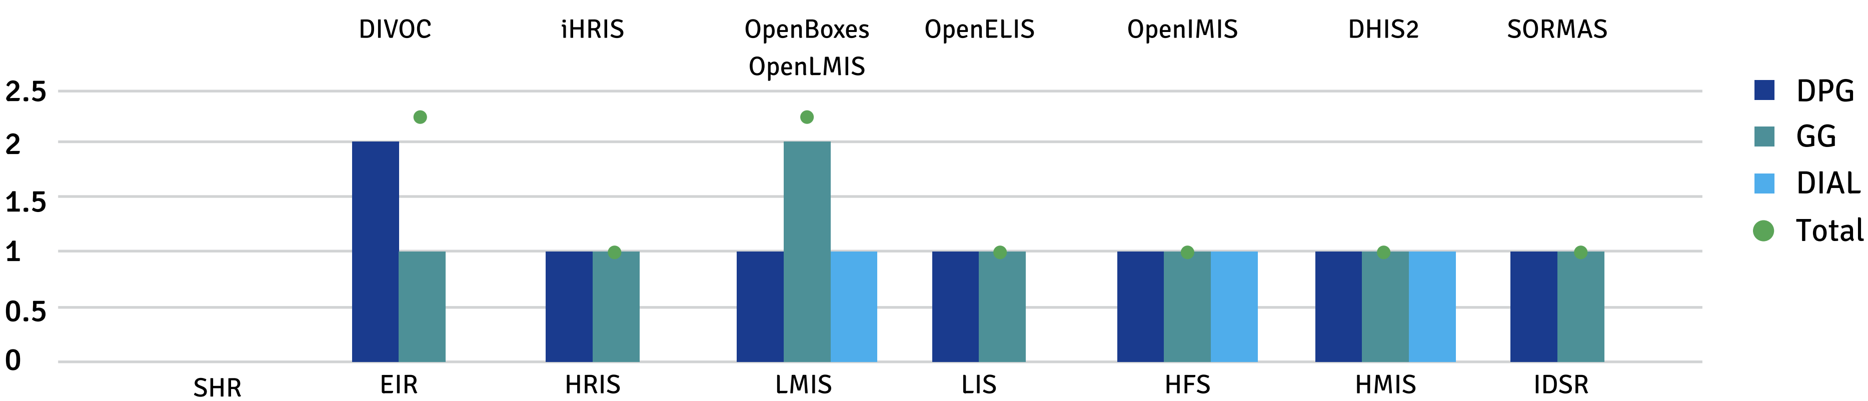

Among the products that may have the potential to function as scaled health services, there is a notable gap in shared health record services, which enable storage and exchange of person-centric records between point-of-service applications. While there are many client-centric point-of-service applications deployed at scale, most are implemented as stand-alone instances that cannot themselves serve as a SHR.

For other categories of shared health services (services for electronic immunization registries and human resources, logistics, insurance, and lab information), products that can function as central services, which can be leveraged by multiple point-of-service applications, do exist.

Figure 6: Product Landscape Breakdown for Registries and Terminology Service

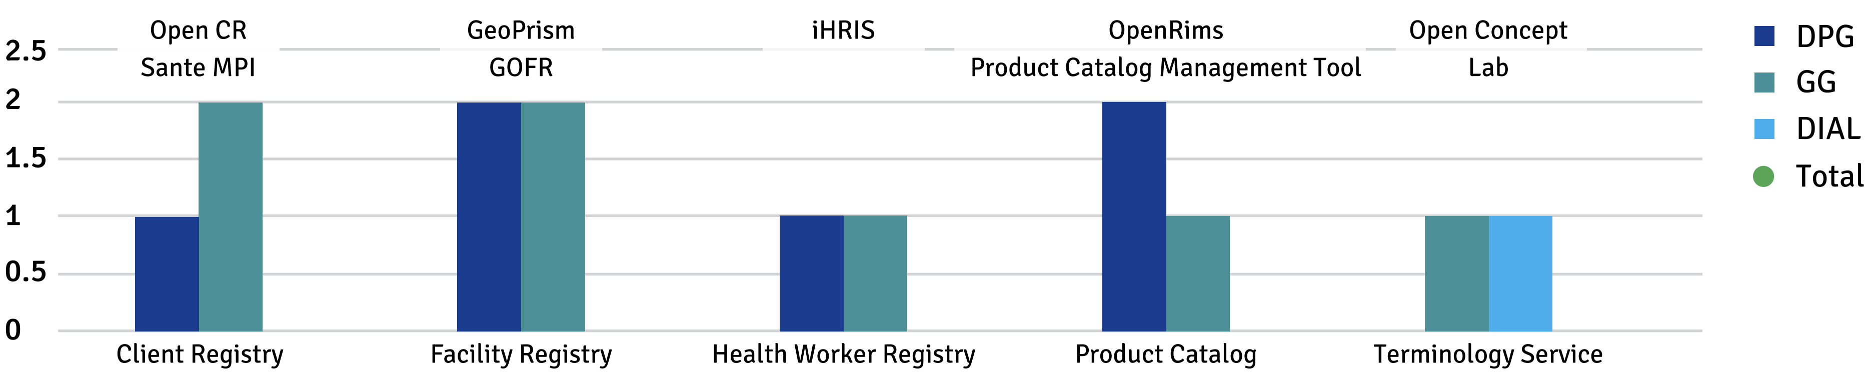

For registries and terminology services, there is also at least one recognized product in each category on the market today. While not a diverse market, one or two products with the right functionalities should make it possible for countries to utilize those tools as part of a DPI-H approach. To that end, the next section considers the readiness of a sample of existing products to function as DPI-H.

2.2 Quality and Readiness of Existing GG and DPG Products for DPI-H Implementation

Significant investment has been made in the development of DPGs and global goods over the last decade. The establishment of DIAL in 2015 and Digital Square in 2016 has strengthened the maturity and conformity of tools with global standards. The OpenHIE framework and associated community have been instrumental in the development of public health information exchange architecture, implementation of standards, and productization of various reference applications. However, the suites of products available are often not ready to implement at scale without significant technical configuration and additional development.

This review looks at a subset of exemplary DPG products in each of the infrastructural components identified in the framework, which strongly overlap with the components of the OpenHIE architecture. Registry services enable the unique identification and listing of various entities; the interoperability layer provides a central point of access to the health information exchange as well as functionality to enable message routing, logging, auditing, and authentication. Scaled health services are centralized health exchange components designed to support specific health system business domains and can combine data from multiple point-of-care systems. While not specified in the OpenHIE Framework, data repositories are highlighted as an vital component of DPI-H, given the increasing role of data analytics in priority health use cases (in some cases, HMIS may act as a data repository). Similarly, products operationalizing health data security frameworks are also considered an important component of DPI-H and are included below.

The review highlights the extent to which products are configurable, in use at scale, interoperable with other systems, maintain necessary security features and documentation, and enable other solutions to leverage their functionality or the data they store. This understanding of readiness to function as DPI-H aligns closely with Digital Square's concept of "shelf-readiness," but also includes the element of configurability.

Health Data Security Frameworks

Underrepresented DPI-H component among existing product landscape

There is a dearth of DPGs that support the operationalization of health data security frameworks and tools that support consent from clients for their data to be used for various purposes: auditing, tracking, and secure transfer of data, and de-linkage and re-linkage of anonymized client records. DPGs such as X-Road and OpenAttestation, which facilitate these functions, are not health-specific. They can be leveraged in the health sector but are currently not in use in LMICs.

Tools to facilitate informed consent from clients to allow their data to be used for research purposes do exist, but primarily in the academic space, and they are not commonly integrated into point-of-service applications. A consent service that acts as a fiduciary framework is absent, although there is active research and discussion in this space, especially in the context of India's National Digital Health Mission. Continued discourse around context-appropriate tools that balance the need for quality care with client data rights is needed to ensure that neither care nor privacy are compromised.

Registry Services – Client Registry

Unique client identification and linkage are necessary for continuity of care

Several client registry products exist and are implemented at a national scale. The Digital Square Global Goods Guidebook indicates that SanteMPI is operational in Fiji, Myanmar, Solomon Islands, and Tanzania, and OpenCR is used in Botswana, Haiti, and Uganda. Both tools implement OpenHIE workflows for client registries and are compliant with HL7v2/v3 and HL7 FHIR standards. However, many countries have developed their own client registries, e.g., South Africa's Health Population Registration System and bespoke platforms in Malawi and Kenya.

While client registries are one of the backbones of a client-centered health information exchange, they are seldom implemented at scale as their use is dependent on real-time transactions, requiring stable internet connectivity and power. In addition, the technology is complex, requiring a deep understanding of client-matching algorithms, local context, and significant engineering capacity to configure and maintain the services. Ongoing curation by a dedicated and skilled team is necessary to maintain the hygiene of the master patient index and ensure that duplicates are resolved and that, more importantly, incorrect linkages are discovered and removed. To ensure that unauthorized users and systems cannot access sensitive patient-level information, strict security measures must be in place. Linkage of client registries to existing digital identity DPI can provide additional levels of confidence in the unique identification of an individual, especially where biometric identification and/or verification are available. However, there is ongoing discussion on the risks and benefits of linking health identifiers to national identifiers, and indeed, in many high-income countries, including the United States, United Kingdom, and Australia, a national identifier does not exist.

Interoperability Layer – OpenHIM

Connective product that enables data exchange between other infrastructure components

OpenHIM acts as an interoperability layer in an HIE and enables secure communication, routing, orchestration, and translation of requests between disparate health information systems. The product was first released in 2012 as part of the Rwanda Health Information Exchange project and was redeveloped in 2014 using more modern technologies (node.js and MongoDB). OpenHIM is implemented in several countries, including Rwanda, Malawi, Tanzania, and South Africa, at a global level as part of the DATIM platform, and as part of iNGO technology stacks (e.g., PSI). However, metrics on usage are not shared, and the maturity of the implementations is not well documented.

Lack of core funding hampers the maintenance and evolution of OpenHIM. Upgrades to the JavaScript frameworks are needed, with support for the front-end framework, AngularJS Version 1, having officially ended in January 2022. Patching emerging security vulnerabilities is challenging, and deployment of new versions of the product is not operationalized. While the product is packaged as an easy-to-deploy Docker instance, the range of versions of OpenHIM deployed in production is not known. The curating organization does not currently have sufficient resources to maintain a robust open-source community.

Scaled Health Services – Shared Health Record

A missing component needed to fully realize the potential of EHRs as infrastructure

A shared health record is a collection of person-centric records for clients and a necessary part of an HIE to enable real-time exchange of clinical documents between point-of-care applications to ensure continuity of care for a client as described in User Journey 1. There are only a few pilot efforts to integrate OpenMRS with a shared health record (e.g., Rwanda, see case study below). Point-of-services applications, including EMRs, while scaled in terms of deployment, are generally unable to exchange client data between instances.

OpenHIE identifies HAPI FHIR as a reference SHR application, and FHIR servers are emerging as platforms that can function as shared health record services. These are implemented at scale in primarily high-income country HIEs, including in the United Kingdom (NHS), Canada, Australia, and Japan. There are few examples of scaled use of FHIR-servers in LMICs.

Where an SHR service is implemented, clinical information can be provided in the form of Summary Care Records or International Patient Summaries, where the most important clinical events for a client can be packaged and provided to health workers or to the clients themselves.

Implementation is complex and is dependent on other components of the HIE to be present (interoperability layer and registries as well as entity mapping functionality to align terminology across the enterprise). In the absence of a transactional SHR service, many countries rely on offline analytics databases and data warehouses, which are examples of the central data repositories component of DPI-H. Consent frameworks enabling clients to opt-in to sharing data across the health system, in alignment with national data-sharing legislation, is necessary but lacking.

Multi-function Applications

Given that DPIs are implemented through products and platforms, some serve more than one function. One example is **District Health Information System 2** (DHIS2), which is traditionally considered a health management information system (HMIS). While DHIS2 performs many functions, the aspects that align with the DPI components are: scaled health service for the collection of aggregate data, data repository for storing those data, and, in some cases, health facility registry.

DHIS2 is categorized by OpenHIE as a health management information service, and while it originally facilitated capture, management and analysis of aggregate routine monitoring and evaluation indicator data, entity (client)-level functionality is provided through the Tracker component. DHIS2 has functioned as a large-scale stand-alone health information system, enabling capture and analysis of routine and event-based data collected through DHIS2 Web and mobile. The DHIS2 community has recently indicated an increased appetite to mature standards-based interoperability features in addition to the ADX format, and is looking to implement FHIR to enable data exchange with other eHIS. With significant long-term funding from international donors, DHIS2 has developed a large footprint in LMICs and is used at scale in over 100 countries, often as a core piece of the national HIS. Although the DHIS2 platform is designed to be configurable for country requirements via a user interface requiring no coding skills, a robust community of DHIS2 experts and consultants support its implementation. The community includes several national health information service provider (HISP) organizations and a number of large, for-profit companies that provide hosting and configuration services. Many ministries of health have in-house DHIS2 capacity but also utilize HISP nodes and/or companies like BAO and BlueSquare for technical assistance and hosting services.

Central Data Repositories

Offline data warehouses and analytics for health

Multiple products offer data repository and analytics functions, including platforms primarily for surveillance such as SORMAS and EpiVerse. Other DPGs offering repository and analytics functions include DHIS2, Laso, InaSAFE, Harmony, and OpenHexa. Harmony and OpenHexa are reviewed in more detail below, as examples.

Harmony

Harmony is an open-source version of Zenysis' data integration platform. A more functional version is available through a paid-for licensing agreement with Zenysis, which is a for-profit company based in San Francisco, USA. The platform enables data integration, analytics, and application of ML models and data quality assurance functionality and has been used by partners and governments in Ethiopia, South Africa, Rwanda, Zambia, Brazil, Vietnam, Benin, Mozambique, and Pakistan. It has primarily been used for the integration of health data sets, but it can be used for data from other domains.

Harmony has a relatively complex technology stack, requiring an engineering team skilled in new database technologies such as Druid and with significant Python coding competence, and it relies on support from the Zenysis software team. However, its social enterprise business model, with significant philanthropic investment, enables Zenysis to co-invest with international donors to provide services to LMICs.

OpenHexa

OpenHexa is a cloud-based, open-source data integration platform developed and provided by BlueSquare SA, a for-profit social enterprise based in Brussels, EU. The platform is implemented in ten countries, primarily in West Africa, with 400 active users across implementations. OpenHexa enables the extraction, cleaning, and storage of health data from multiple sources, including DHIS2, for visualization and analysis. A BlueSquare, in-house data science team is required to operate the system, which is based on JupyterHub and Jupyter Python notebooks. Data are stored in PostgreSQL and cloud blob stores. OpenHexa uses generic business intelligence tools for visualization (e.g., Tableau, PowerBI, Shiny, Dash, and Voilà).

Point of Service Applications – Electronic Health Record Systems

OpenMRS – Scaled Point-of-Service EHRs

A 20-year investment in OpenMRS has resulted in enormous adoption of the product, an EHR designed for use in LMICs. As of October 2023, it is implemented in over 8000 sites across 40 countries, servicing more than 15 million active clients, and it is the core software component of national EMRs including RwandaEMR, UgandaEMR, and KenyaEMR. While a huge benefit has been realized, the full potential of the software has not been achieved as instances are predominantly stand-alone, have limited participation in national HIEs, and rarely contribute to a shared health record.

In some cases, OpenMRS contributes data to a national data repository. Records need to be manually transferred from individual OpenMRS instances to the repository, either digitally or via data export to physical media.

In addition, as use of the software has grown, new challenges have emerged. A 2020 security assessment conducted by MEASURE Evaluation highlighted some gaps in operational and system security, findings that can be generalized to other implementations of OpenMRS. The stand-alone implementation of OpenMRS leads to divergence in concept dictionaries across instances, creating a barrier to interoperability. Significant effort and investment have been contributed to enhancing OpenMRS interoperability, most recently in the development of the OpenMRS FHIR API and in the integration of OpenMRS with OpenConceptLab, but these are predominantly used in laboratory settings.

A large community of developers contributes to the codebase, but contributions are often not aligned with coding standards, and a large amount of technical debt – the implied, future cost of not fixing problems in the present – has been introduced into the code base. The community struggles to support skilled core developers who can ensure code quality of community contributions as well as contribute code themselves.

Summary of Product Review

Across DPI-H component categories, there are few products mature enough to function in a scaled, interoperable infrastructure. Further, few products in the current landscape are leveraged by other DPGs and thus cannot be considered as DPI-H. DHIS2 is exceptional in scale and maturity given the longevity of use, level of investment, provision of training, and available technical support. However, it is often implemented as a stand-alone platform with limited standards-based interoperability. Products in other categories have varying levels of "shelf-readiness." They often require configuration support and strengthening for deployment and use at scale and additional development of core software to add required functionality or fix bugs. The open-source communities supporting these products vary significantly. The University of Oslo employs a large team of software developers, analysts, product owners, DevOps engineers, and other staff needed to make up enterprise software development teams; however, there is limited contribution of code to DHIS2 from external contributors. OpenHIM and other registry products have much more limited communities and funding mechanisms. Of all the DPI-H components, products supporting health data security frameworks are relatively underdeveloped and not routinely used or integrated with other digital health systems, representing a critical gap in the product landscape for DPI-H. Altogether, while products exist in each category of DPI-H, they are often not fully shelf-ready for easy implementation by countries. Implementing these products is inherently challenging, especially considering the complexity of health terminology, messaging standards, and local workflows. However, strengthening the overall shelf-readiness of these products may reduce the additional technical burden required to implement them in a specific country context.

2.3 Case Studies of DPGs Deployed in LMIC Contexts

The challenges noted with available products are further complicated by the nature of investment in DPGs at the country level. The case studies below look at several countries using one or more DPGs at scale to explore on-the-ground limitations to functioning as infrastructure in an LMIC context. While all have some global goods in use at scale, in particular EMRs, the case studies show how differences in governance, adoption of an architecture and planned approach to system development, and history of disease-oriented vertical investments have hampered the development and use of products as DPI-H.

Case Study: Nigeria

Nigeria has a very successfully scaled point-of-service EMR application, yet it is largely operating in isolation, without the ability to connect to other health services or components of health information exchange. NigeriaMRS, derived from OpenMRS, is the single highest adopted electronic client record system in the country. To date, over 98% of all CDC-supported healthcare facilities in Nigeria use NigeriaMRS to support their HIV programs. Yet, of the total electronic health records in the country, 99% are used exclusively for HIV programs and have historically operated independently of other scaled systems.

Nigeria's electronic laboratory information management system (eLIMS) is used nationwide for viral load management through a PEPFAR program, but it has only recently integrated with the EMR system (2021) to begin automatically recording client viral load results instead of having them manually transcribed. The Federal Ministry of Health developed a National Health ICT Framework in 2015, which called for an architecture to unify disconnected functional applications through an infrastructure layer, but this is only beginning to be implemented.

Nigeria reflects a context in which individual DPGs have scaled, but without the strong governance, architecture, and funding dedicated to a holistic architecture, it is not able to leverage the systems it has within HIV for broader digital health public infrastructure.

Case Study: Kenya

The KeHIMS project in Kenya is another example of scaled use of DPGs in-country, designed with the intent to ultimately support many connected systems and data exchange. Initially funded by PEPFAR, the KeHMIS project has supported Kenya in standardizing and implementing over 1,200 instances of KenyaEMR, which is a modified version of OpenMRS. In addition, through support from Palladium, KeHIMS includes a data exchange layer that provides interoperability for eight systems, including other DPGs: DHIS2, automated appointment scheduling, SMS adherence messaging, antiretroviral dispensing, SMS laboratory notifications, and a partner progress monitor. The MoH maintains a national data warehouse for data collected through KeHMIS-connected applications. KeHMIS is now supported through both domestic and donor funding, and the MoH has created standards that other digital products must align with to work with KeHMIS. There is strong governance by the Kenya HIM task team, which has prioritized and commissioned six use cases that were jointly funded. This is the result of over a decade of investment, training, and capacity-strengthening efforts through large PEPFAR programs.

While Kenya is better positioned than other countries in having relatively strong governance and an extensible architecture in place, it still faces significant implementation challenges to achieving data exchange between multiple scaled health services and registries. A 2020 evaluation noted many of the facilities, 183 of 213, lacked the interoperability layer (IL) module and, hence, had no capability to exchange health data with external systems. Further, of the approximately 14% of facilities that had data exchange capability, the majority were in one county. Thus, while theoretically possible to enable data exchange between KenyaEMR and other services, implementation has been slow. Challenges in underlying technology infrastructure still hamper the ability to connect many facilities and counties with inconsistent or no connectivity. The same evaluation found approximately 50% of facilities lacked a nationally standardized client identifier format, which also limits the utility of a client registry.

Case Study: Rwanda

Rwanda was an early leader in implementing an enterprise architecture approach and a national health information exchange. While investment from 2012 has not resulted in a fully functioning HIE today, Rwanda has many characteristics that support the ability to realize DPGs for DPI-H. Rwanda has strong digital health leadership and governance. Digital health is being implemented as part of the national health strategy, and governance structures are fully functional, government-led, consultative of other ministries, and responsible for monitoring the implementation of digital health. The MoH's Chief Digital Officer and team coordinate activities between the government-wide body working on digital transformation, the Rwanda Information Society Authority (RISA), the Digital Health Committee, and the technical working group, reviewing and managing all potential partnerships between the MoH and technology implementers.

With its strong governance, Rwanda has implemented multiple DPGs throughout its health system, including OpenMRS, OpenClinic, e-Heza, and RapidPro for client care. Rwanda has a mature DHIS2 implementation for routine indicators, as well as multiple DHIS2 tracker programs covering noncommunicable diseases and IDSR. However, Rwanda is experiencing challenges in maintaining open-source products that have been so successfully scaled. RapidSMS, for example, grew to a point where it lost functionality and became unstable, requiring recent calls for additional investment to regain utility.

Although many of the point-of-service applications are scaled, they operate in isolation as the interoperability layer is not fully functional.

Rwanda is in the pilot stage of implementing OpenHIM as an interoperability layer. It is currently undertaking efforts to connect its client registry to a client data repository that would connect to multiple point-of-service applications and serve as an SHR. As of July 2023, there are 65 connected instances to the SHR. Implementation challenges include reducing the time to sync data to ensure the SHR can function in real-time, technical challenges integrating existing systems with OpenHIM, and strengthening trust with clients to encourage their consent to share data at a national level.

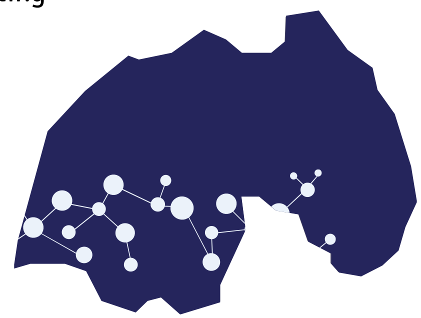

Rwanda's strong governance has been relatively successful in coordinating donors around planned systems, yet historically, funding from donors and partners has followed fairly narrow, program-oriented goals. Despite being an early leader in implementing an enterprise architecture approach, the investment in a full health information exchange was not sustained, and it is only recently getting to pilot stage. To move from the pilot stage to a fully operational DPI-H, Rwanda needs sustained, program-agnostic funding to operationalize the interoperability layer and canonical registries, maintain and update digital tools already used within its system, and strengthen regulatory and governance capacity to ensure new tools introduced into the system align with the planned architecture.

2.4 Limited Landscape of Key Supporting DPGs, Global Goods, and Building Blocks

While DPI-H components, the infrastructural DPGs, and building blocks are foundational and most commonly leveraged within the digital health ecosystem, several key Supporting DPGs and building blocks were identified as important to enable the user journeys described in Part 1. There are relatively few data, content, and AI model products listed as DPGs, global goods, or as building blocks on the Digital Impact Exchange. These types of building blocks are far fewer overall, regardless of SDG or sector, and especially in the health sector. Specifically, there are seven products identified as open data DPGs and one open AI model. Notably, there are no open data sets listed in the Digital Impact Exchange for One Health use cases.

While not comprehensive, the limited representation of open data and open AI models in the three product lists may contribute to difficulty finding such assets. Separate from lists of DPGs, GGs, and building blocks, there are many other platforms that host open data. For example, datacommons.org offers a platform to find open data, and Github hosts AI model packages and documentation for contributors who wish to use them. However, the landscape of data, content, and models is more diffuse and less standardized than that of products that could function as DPI-H. The DPG and Global Goods certification process has been primarily used for evaluation of software systems. The dearth of data sets and models in health may signal a need for more socialization around the value of these assets for the kinds of predictive analytic needs identified in the user journeys, as well as refined guidance for how to evaluate them for use in conjunction with a country's DPI-H.

There is growing interest in and support for Digital Health Content Global Goods. The WHO SMART Guidelines can be classified as content global goods, and significant effort is being invested into the development of digital adaptation kits and FHIR Implementation Guides for several guidelines. These initiatives require mature ecosystem enablers to succeed.

Part 3:

Challenges Using DPGs as DPI-H: Role of Enabling Environment

Environmental Health Officer for Gwembe district, Zambia, reviews malaria surveillance data on a laptop computer.
PATH/Gabe Bienczycki (woman shown is Marie-Reine Rutagwera, Malaria Surveillance Specialist for MACEPA)

Part 3: Challenges Using DPGs as DPI-H: Role of Enabling Environment

Given the limitations in available DPGs and the context of their use in LMICs, numerous challenges to the use of DPGs as DPI-H become apparent. Many of these are familiar challenges to the global digital health community, yet they merit reiteration given the challenges they present for the realization of DPI-H.

3.1 Countries still experience insufficient leadership and governance capacity to implement DPI-H

Many LMICs have made significant progress over the last decade in developing digital health governance structures, planning for digital health systems, and drafting and adopting digital health strategies and enterprise architectures that provide a backbone for digital health governance.

Country efforts have also been strengthened by investments in leadership training and capacity. The United States Agency for International Development (USAID), World Health Organization (WHO), International Telecommunication Union (ITU), and Digital Square collaborated on developing a training course for planning national digital health systems that creates awareness and buy-in on the idea of an enterprise architecture approach, which would facilitate the implementation of multiple components of DPI-H within one system, connected by an interoperability layer.

Despite these advances, there remain many countries with no plan or strategy, and many struggling to maintain leadership support through drafting of strategies to implementation. A recent review of Asia eHealth Information Network (AeHIN) countries that have health enterprise architecture frameworks noted implementation challenges related to limited enterprise architecture knowledge and lack of senior leadership involvement. Turnover within ministry of health officials results in need for continued investment in training and capacity development at leadership level.

3.2 Countries that do have strong leadership are facing technical and regulatory implementation challenges

Even where countries have strong leadership support, technical and regulatory capacity challenges can limit implementation. According to the Digital Health Monitor, most countries in Sub-Saharan Africa have an approved digital health strategy or framework (orange countries, Figure 7) and many have a costed plan (yellow countries, Figure 7). However, few have moved into the resource mobilization and implementation phases of their strategies (green countries, Figure 7).

Several factors contribute to this. First, most LMICs have limited in-country implementation capacity needed to configure existing products in order for them to function as infrastructure. Leadership capacity development and training activities to date have largely focused on advocacy around the idea of enterprise architecture and the need for common, shared infrastructure. There has only recently been a significant effort on developing the technical aspects needed to bring together health data terminology and messaging standards to enable functional interoperability. FHIR has emerged as a messaging standard that could make this possible. It has significant momentum from donors and, to some extent, at country level; however, it needs further commitment and regulation from country bodies to succeed.

In addition, countries have limited regulatory capacity to effectively shape the national product landscape to align with the standards and architecture needed for DPI-H to function. At a global level, the digital health community has produced numerous resources, guidebooks, and certification processes to begin to move the market of digital health interventions in the direction of DPGs, including the subset of those that could function as DPI-H. Resources like Digital Square's Global Goods Guidebook, DIAL Digital Impact Exchange, WHO Digital Clearinghouse for Solutions, and the Digital Implementation Investment Guide, provide high-level guidance and identify products that meet the specific criteria that align with DPI-H requirements. However, none of these bodies or tools is a regulatory mechanism. Few countries have a body with the authority to effectively enforce alignment with established standards or guidelines. Several are making strides with this authority, such as Rwanda's Digital Health Office, yet even when a body exists there may not be sufficient capacity to effectively enforce alignment with the adopted approach.

Figure 7: **Status of Digital Public Health Frameworks**

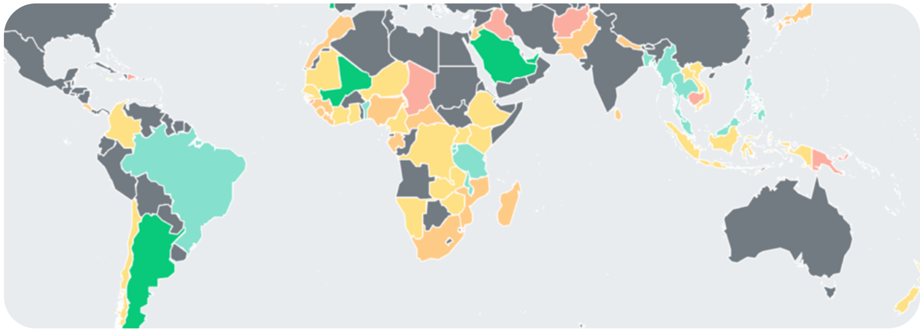

Image Credit: Digital Health Monitor Strategy & Investment Indicator, July 2023.

3.3 Implementation challenges are exacerbated by technical deficits in key DPGs and technical capacity in-country

While countries have been encouraged to adopt open-source products, funders, and implementing partners have under-emphasized the operational expenditure needed to develop and maintain software products as part of their health information exchanges. Without a full understanding of the total cost of ownership for open-source products, countries may begin integrating open tools into their health system with insufficient resources to maintain and improve them.

DPGs are sometimes provided as codebases rather than as 'shelf-ready' products, and need further work to configure and maintain them as part of country DPI-H. Configuration and maintenance of these products requires technical skills that are often scarce in LMIC contexts. For example, these include mature software development and information and communication technology (ICT) skills, as well as digital health technical teams to contribute to core development of DPGs, including feature development, security patching, and integration. While maturing FHIR functionality makes integration less complex, FHIR itself is technically complicated. Successful adoption requires skilled and fully capacitated software teams that have the necessary background in health, integration, and modern software languages.

FHIR offers distinct potential to reduce proliferation of bespoke tools by simplifying data exchange and enabling sharing of common tools and machine-readable business logic between DPI-H components and the functional applications that provide services. For example, FHIR-based SMART Guideline Implementation Guides enable configuration of digital health applications without the need to change underlying application code. Fully realizing this potential requires an ecosystem of services and applications that are FHIR-native, and utilize a common set of FHIR libraries and tools. However, SMART guidelines are newly emerging and require technical expertise to implement. Implementation of FHIR messaging standards alone is not sufficient to enable interoperability as there are few FHIR-native apps or services available.

Further, FHIR-based interoperability solutions must still work with products using multiple data standards and code lists, increasing the need for a workforce comfortable with FHIR in addition to LOINC, SNOMED, ICD, and other semantic standards.

DPG community contributions are intended to provide some of the time and skill needed to configure and maintain open-source products, but, as noted in the product case studies, many existing DPGs have technical debt from underinvestment in core product development. Moreover, their community of contributors are generally volunteers from other LMIC contexts, so they do not have the benefit of paid software developers or a global community of users that contribute to other open-source products used in high-income countries (e.g., MongoDB, Django, Apache, Mozilla). Together, the complexity of implementation, limited country capacity, and existing technical debt present high hurdles to the use of DPGs as DPI-H.

3.4 Engagement between public and private sectors is not clearly regulated

Private and public sector health services are not often aligned, with limited, if any, participation of private sector HIS in public sector health systems or plans for the exchange of health data between sectors. Data sharing between the public and private sectors needs to be well-defined and regulated by the ministry of health. This area of data governance is not mature, and the inability to link records across the sectors limits the continuity of care for clients who may use both public and private services for different health needs.

Ministries of health are sometimes not comfortable with the utilization of proprietary software solutions because of concerns about loss of ownership of data held within vendor systems. This concern may also apply to the intellectual property that a country has contributed to a product during the development and use of the product in the country context.

Both concerns may be mitigated through the design and enforcement of appropriate procurement contracts with vendors, but ministries must have sufficient capacity and robust regulatory frameworks to negotiate contracts to their advantage.

Despite these challenges, countries may opt for a mix of open-source and proprietary products in their ecosystem. Ensuring countries have a full understanding of the total cost of ownership of a software product, including long-term capital and operational expenditure, can facilitate fully informed choices. At the same time, ensuring DPI-H is developed and regulated so that private sector actors can also leverage it will support a more robust market and will be particularly important for use cases around health payments where private sector actors are a significant contributor to country health service delivery.

3.5 Donor investments in country-adoption of DPGs have been in disease silos rather than with an ecosystem approach

As seen in several country case studies, investments in DPGs and global goods have typically been through disease-oriented programming, focusing on one product or one service rather than an integrated architecture. This creates additional complexity in the system for the HCWs, who must use multiple tools or a mix of paper and digital tools across their client load (as in User Journey 2), and makes it more difficult for countries to advance an agenda of harmonized, interoperable systems.

Research on fragmentation in global health has identified additional challenges related to the alignment of donor funding and national priorities. For example, according to the visualization tool for financing global health from the Institute for Health Metrics and Evaluation (IHME), donor funding for health systems strengthening has accounted for a small percentage of the total health spend, and, more importantly, that percentage has not grown significantly as technology has become more pervasive. While there are investments in digital systems through disease-oriented programs, it remains difficult to measure the magnitude of investment in digital systems and whether it is keeping pace with the role of technology in health systems.

The absence of scaled, interoperable tools functioning as DPGs that serve the entire health sector may reflect both the tendency to invest in disease verticals rather than ecosystems, as well as an overall underinvestment in ecosystem-level activities. Both factors make it more difficult to advance agendas around harmonization, even where there are strong country-level governance mechanisms and existing digital health strategies in place.

3.6 Current digital systems are designed for data collection rather than health data exchange and use

Where specific point-of-service applications are operating at scale, they have primarily been designed to provide routine monitoring and evaluation data and not as part of an interconnected digital infrastructure that would directly contribute to client care. A focus on data collection, often to satisfy the needs of donors, has resulted in fractured digitization of individual health programs. This places an increased data collection burden on health workers, who often need to use multiple digital tools as well as paper records and tally sheets. Submission of disparate data to central data repositories is useful and necessary for program management, but it fails to address the needs of clients and health workers. Requisite data could be sourced from point-of-service applications participating in an HIE and contributing to an SHR, without additional burden.

Additionally, the absence of a consent mechanism in most point-of-care applications hinders the ability to ethically contribute to a shared health record, even if the application is technically able to do so. Data governance protocols for personal health data are only starting to emerge, and, in many cases, data sharing is limited by default rather than finding compliant ways to share data across systems.

Further, as there is increasing capability to use advanced data analytics for program strategy and design, there is an increasing need to work with data from multiple reporting databases. For example, determining the effectiveness of CHW visits on client retention and health outcomes requires CHIS, EMR, and LIS data. Enabling health system managers to adjust testing guidelines dynamically based on the availability of test kits requires LMIS and EMR data. To provide high-value information, data from several operational data sources need to be linked and queried – operations that can be challenging without a dedicated data repository for analytics.

While existing operational data is a significant source of data for analytics, there is a widespread perception that much of the data produced is underutilized. Concerns about data quality, the regulatory implications of sharing client-level data across systems, and unfamiliarity with available open data all contribute to the limited implementation of the data-analytics functionalities identified in Part 1.

When looking specifically for data that could be leveraged for machine learning and AI models, having data representative of the context in which it will be used is critical. The limited availability of labeled training data sets for machine learning models is another limitation to achieving predictive analytic functionalities at scale. While some programs are able to analyze data originating from EMRs implemented as part of that program, most countries do not yet have widespread digitization of health system data across disease verticals, and governance for the management and use of that data is still emerging.

Prior to the launch of Lacuna Fund, a dedicated fund for the creation of labeled training data sets for LMICs, a landscape review[3] noted that privacy protections for labeled data sets, licensing considerations for labeled training data sets, and platforms for discoverability are all emerging issues that contribute to the challenge of creating and using open, labeled training data for machine learning models.

In some cases, organizations may have data they are willing to make open, but it is "messy" or otherwise resource-intensive to prepare for release, and organizations have little incentive to do so. Where for-profit organizations invest in primary data collection to develop a training data set, licensing for limited access may be necessary to maintain a sustainable business model.

Summary of Challenges

Altogether, these challenges reflect both familiar, persistent challenges of strengthening enabling environments for digital health through leadership, capacity, and governance, as well as more specific challenges for DPI-H, namely, the technical debt of existing global goods, the nature of investment in digital systems, and the emergence data governance challenges for data exchange and data use in the health sector. To realize the potential of DPI-H over the next five to ten years, both sets of challenges will need to be addressed.

[3] *Meridian Institute. 2 March 2020. Universal Labeling Project Landscape Assessment. Shared via personal communication.*

Digital Innovation in Pandemic Control (DIPC) project - PATH

Part 4:

Recommended Approaches to Support DPI-H Implementation

Part 4: Recommended Approaches to Support DPI-H Implementation

The strategies outlined below are envisioned as a collective approach for the digital health community rather than one for a single organization. Specific approaches will likely align to varying degrees with a wide range of digital health funders, and this report anticipates that different organizations will lead in different approaches. Potential investment strategies are presented according to three broad approaches for strengthening the implementation of DPI-H. The first approach includes a set of global-level strategies that will facilitate the implementation of DPI-H by strengthening the market of existing products for use as DPI-H, addressing market gaps in the operationalization of health data security and consent, strengthening the market of supporting DPGs and building blocks, and enhancing governance approaches for a future in which predictive analytics will be layered on top of DPI-H.

The second approach includes a set of strategies that work to strengthen the country enabling environments to implement DPI-H. These country-level strategies respond to the familiar, yet persistent, challenges of leadership, technical, and regulatory capacity that will hinder DPI-H implementation efforts if not addressed.

These two approaches support a third, direct approach to realizing DPI-H: coordinated, country-level investments to implement multiple DPI-H components. This approach is a key connection between the global market strengthening and country capacity strengthening approaches. Coordinated country-level investment in DPI-H implementation can provide an opportunity to leverage existing global assets and also strengthen them with contributions from real-world implementation experience. Figure 8 shows the complementary relationship between the three approaches. Each of the investment approaches is further detailed below with specific illustrative strategies, followed by high-level cost estimates.

Figure 8: Relationship Between DPI-H Investment Approaches

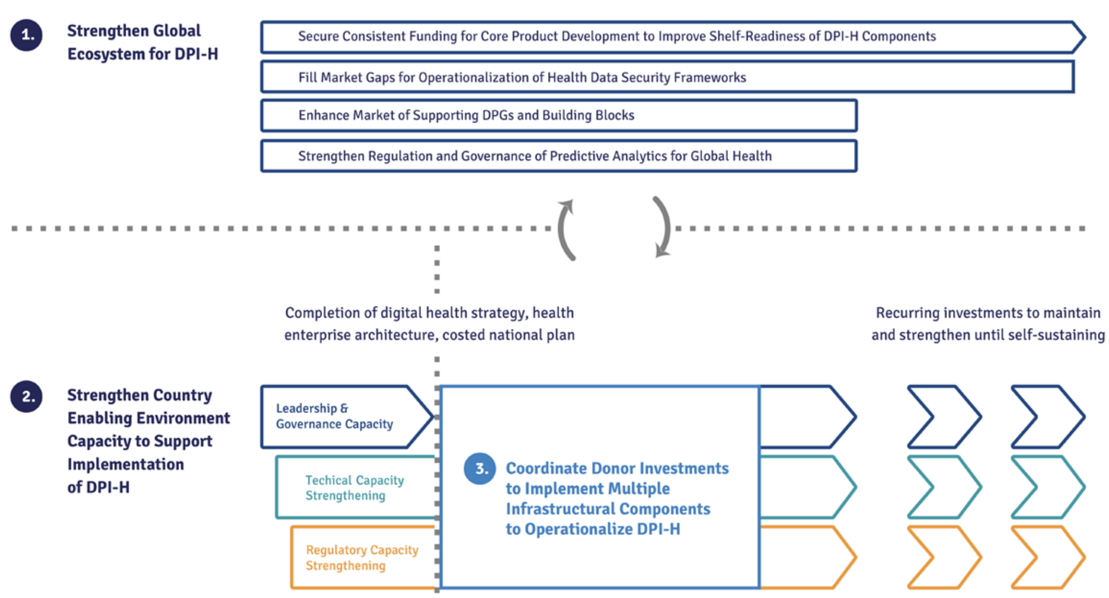

4.1 Approach 1: Strengthen Global Ecosystem for DPI-H

Support a funding model for consistent core product development of identified DPI-H components to strengthen shelf-readiness.

As noted in the product landscape above, many DPGs with the potential to serve as DPI-H are not technically ready to do so. Most DPGs, while meeting certification to become a global good or DPGs, were designed as stand-alone products rather than one set of products that, together, can perform the functions of a health information exchange. Even where open-source tools are successfully scaled and used in multiple countries, many do not have the necessary deployment and maintenance capabilities or are not compliant with the necessary messaging standards or security measures to enable them to serve as DPI-H. As a result, countries investing in tools for their digital health infrastructure currently must choose between open-source tools that come with less mature security, maintenance, and deployment support or commercial products that have proven maturity but can introduce dependence on software and service vendors. In the present landscape, the cost of commercial products may be preferable in some contexts, depending on licensing costs and terms of commercial contracts. Strengthening core DPG applications so that they offer comparable security, interoperability, deployment, and maintenance to commercial products will enhance their viability as DPI-Hs.

Illustrative investments include:

- Supporting a governance body (e.g., an organization or secretariat) with full-time technical staff to support core product development for GGs and DPGs that could serve an infrastructural function. The body would transparently distribute funds to product custodians to invest in strengthening the product for use as DPI-H. Core development would include improving identity and access management, security, consent capability, and use of messaging standards such as FHIR. DPI-H readiness largely aligns with Digital Square’s definition of “shelf-readiness,” which may be an initial guide to the kinds of investments this governance body would support. The governance body could provide technical support as needed and should be responsible for assessing outcomes of funding.
- Support the development of the FHIR-ecosystem, including work on FHIR-servers, FHIR libraries and FHIR-based analytics as well as new FHIR-resources for needed in LMICs contexts.
- Investing in an ecosystem of organizations to provide sustainable technical implementation support to countries to utilize GGs and DPGs as DPI-H. These organizations should have the technical capacity to configure and implement these products as part of a health information exchange infrastructure, not as stand-alone implementations.

Strengthen the global market of Supporting DPGs: representative training data, open content

The product landscape also revealed a dearth of DPGs that fill the role of Supporting DPGs on open data, open content, and open AI models. These Supporting DPGs make use of underlying DPI-H and enable the creation of a diverse set of functional applications that can address a range of health challenges. A robust digital health ecosystem should include a wide variety of such DPGs – open data and representative training data for AI models that respond to a wide range of health priorities, including but not limited to the personalized care, risk-based prioritization, and outbreak forecasting use cases identified in the user journeys. Similarly, a wide range of health interventions would benefit from open content that can support digital applications, such as machine-readable clinical guidelines. While there are emerging initiatives to address these challenges, they are relatively small. Lacuna Fund is beginning to fill gaps in representative health training data for AI models,[4] supporting a recent call for representative training data sets for health equity and climate and health. Out of that call, Lacuna Fund was able to award just under $3 million to support the creation of 15 datasets. However, representatives estimate an unmet need of over $5 million among the pool that applied.[5] As predictive analytics use cases grow along with urgency around priorities such as outbreak forecasting and One Health, there will be an even larger need for discoverable, representative, accessible training data sets.

With respect to open content global goods, SMART guidelines are another area where more investment may be needed. WHO SMART Guideline Layer 2 Digital Adaptation Kits and Layer 3 Implementation Guides provide open content that enables digitization and, eventually, automation of care plans for many interventions. There remains significant room for expansion to additional clinical areas as well as disease surveillance and outbreak response.

Illustrative investments include:

- Continuing support for the development of representative training data sets to support machine learning models for global health problems.
- Increasing support to further develop WHO SMART Guidelines for additional health priorities to enable standardized care plans and produce digital clinical records that can be used in precision health and segmentation models.

[4] Lacuna Fund's health datasets are supported by Wellcome Trust, The Rockefeller Foundation, Google.org, Gordon and Betty Moore Foundation, Patrick J. McGovern Foundation, and Robert Wood Johnson Foundation.

[5] Unmet need estimate reflects the value of applications that made it to final review and were not ultimately funded. Estimate provided in personal communication and shared with permission.

Fill market gaps to operationalize health data security frameworks: consent tools and governance for sharing data across health systems through DPI-H

The product landscape also revealed that there are few existing, recognized DPGs or global goods that enable a consent mechanism or framework that would allow clients to control how their health data is shared across the health system. The lack of context-appropriate consent tools and frameworks becomes more concerning as clinical records are increasingly centralized in shared health records and central repositories, and the appetite for use of these data in ML and AI contexts grows. While some technologies are available, many of them are designed for use in research and academia, and they have yet to be adopted in digital health. Incentivizing the development and use of tools to facilitate client consent for their data use in an HIE will advance the availability of a suite of DPI-H tools. This may require socialization efforts with client communities as operationalization of data governance principles such as consent is not routine. Change management and client socialization activities are important complements to developing the technical tools for data security and consent for the use of client-level data across systems using DPI-H.

Illustrative investments include:

- Supporting the development of new (or modification of existing) tools that enable client-level consent for data sharing and use across scaled health services.
- Supporting the development of central consent frameworks to enable fiduciary management of client records.
- Supporting community-level socialization around individual data rights, norms, and expected safeguards as well as the benefits of data sharing across health systems.

Strengthen regulation and data governance for predictive analytics in global health

As health systems increasingly seek data use through predictive analytics, there is a need for enhanced guidance on appropriate governance approaches for open AI models and the training data sets that used to develop them. Preparing health system administrators to understand the implications of outbreak forecasting, or appropriately integrate client risk-scoring models into point-of-care systems will be enhanced by creating and operationalizing clear, endorsed guidance for how to evaluate model performance, safety, and interpretability. The WHO recently released a comprehensive report on Ethics and Governance of Artificial Intelligence for Health, which outlines a framework for the governance of AI in health, including regulatory approaches. Translating these recommendations into practical tools and specific resources will enable countries to strengthen their regulatory capacity and operationalize a governance protocol.

At the same time, developing specific criteria for certification of AI models as DPGs can translate recommendations into existing DPG and GG review processes. For example, standardizing the inclusion of documentation on model performance and risks through the use of model cards for AI models or datasheets for data sets, both of which are emerging from private sector AI leaders, can improve transparency and the ability of the global community to responsibly use open AI models.

Specifying governance approaches for predictive models can also improve usability. Encouraging, and even requiring, the use of technical standards for models such as the PMML format will ensure models are not siloed by the programming language in which they were built, and can run on top of DPI-H components that may already be in use. Similarly, using a standard data format for metadata and open data sets, such as those provided by Schema.org, can help improve the usability of data for analytic purposes. Codifying these approaches in endorsed data governance implementation guides will further strengthen the enabling environment for the use of predictive analytics.

Illustrative investments include:

- Developing contextualized, practical materials to support the operationalization of emerging AI governance frameworks and documentation approaches.
- Codifying the standards and documentation that should be provided for open AI models used in digital health systems.
- Incentivizing the development and licensing of AI models in ways that make them standardized and accessible, for example, through procurement language.
- Generating evaluations and evidence for the public health benefit of predictive models to facilitate safe and effective use of models in practice.

4.2 Approach 2: Strengthen Country Capacity to Prepare for DPI-H Implementation

Implementing DPI-H requires countries to have sufficient capacity at leadership, regulatory, and technical levels. This necessitates continued investment in familiar strategies to strengthen the digital health enabling environment at country and regional levels. While these approaches are not new, they nevertheless remain critical to realizing the benefits of DPI-H.

Continue leadership-level advocacy and management-level training for strong governance and development of digital health strategies, enterprise architectures, and costed national plans

Given the continued, siloed use of digital tools, supporting countries to develop a top-down plan of the digital systems and digital products that support their needs will advance the vision of DPI-H in countries that have yet to reach critical milestones of digital health system maturity. Investing in digital health strategies, architectures, and costed plans with public-sector leaders will facilitate the development of digital health infrastructure that is transparently governed, accountable, and aims to support an ecosystem of innovators and developers in the digital health space.

Illustrative investments include:

- Training all levels of the health system, including leadership and technical levels to socialize concepts of DPI-H, making use of common digital infrastructure, and advocating for an enterprise architecture approach to achieving health goals.
- Supporting country-level workshops and consultative engagements with ministries of finance and other key stakeholders involved in digital transformation to draft health-enterprise architectures, align with cross-sectoral digital transformation efforts, and facilitate creation of costed plans.
- Repeating training opportunities periodically to improve continuity of leadership in digital health and ensure there is a pipeline of digital health champions aligned on a common roadmap.

Support local workforce with technical implementation capacity needed to leverage DPGs as DPI-H

Implementing and integrating the components of DPI-H is technically complex and requires a workforce with the skills necessary to translate country-specific health requirements into software code that aligns with global emerging standards. Local workforces need to be able to support and contribute to an evolving ecosystem of DPGs and commercial products, to support implementation across the health sector and with other sectors.

Illustrative investments include:

- Adopting best-in-class training material for ICT and software development and deployment, and developing additional training on data exchange protocols such as FHIR, where needed.
- Supporting regular technical training on functional applications and services implemented in-country, as well as on standards and health information exchange strategies with practical applications.
- Working with local universities and job training institutions to develop and adopt existing curricula to create a pipeline of software engineers and data scientists skilled in working within an ecosystem of DPGs and shared digital infrastructure.
- Facilitating public-private partnerships to leverage professional software skills for developing and tailoring DPG products for use in a country's DPI-H.

Support countries to develop regulatory bodies with the authority to evaluate products against approved standards

While planning for national digital health systems that employ DPI-H will be top-down, the goal of DPI-H is to enable many different actors, public and private, to create solutions that work within it. To ensure products developed within the ecosystem align with an approved health enterprise architecture and associated standards, countries need the local capacity and authority to assess products as well as approve and facilitate the integration of approved products for use within the DPI-H suite of tools. An authoritative regulatory body needs the capacity to certify digital products and oversee the implementation, maintenance, and evolution of an approved enterprise architecture blueprint that maps out core DPI-H. This is a critical role needed to limit fragmentation and proliferation of solutions that will not support the vision of DPI-H. This body could also take on additional, important layers of governance in evaluating claims of effectiveness, benefit, or improved performance that solutions may make, in order to enhance the accountability of products leveraging DPI-H in a country's digital health ecosystem.

Regulatory capacity will also be critical for operationalizing emerging data governance protocols. From implementation of health data security frameworks to navigating emerging data sovereignty regulations and operationalizing guidance for use of data for predictive analytics, countries will need regulatory capacity to develop and support fit-for-context data governance practices and the protocols that support them.

Illustrative investments include:

- Supporting the establishment of a regulatory body within a ministry of health if none exists.
- Developing clear regulatory guidelines for requirements that align with the envisioned system as well as supporting tools, such as guidance around tradeoffs between open source and proprietary products, and checklists for approval of products for inclusion in public and private sector health systems.
- Supporting assessments of data governance approaches in alignment with adopted principles. This may include testing different approaches to optimize data protection and preservation of individual data rights and balancing these with availability of data to ensure quality service delivery.
- Consistently supporting established regulatory bodies with sufficient human resources, including staff with ICT, business analysis, and project management skills.

4.3 Approach 3: Coordinate Donor Investments for Implementation of Multiple DPI-H Components

Given the interdependence of DPI-H infrastructure and DPG products, investments in DPI-H should focus efforts on implementing multiple DPI-H components together. A holistic approach focusing on multiple, connected components of DPI-H can demonstrate the economy-of-scale benefit of country-level health infrastructure, providing cross-disease support that most countries have not been able to realize. Investment focused on use cases that necessarily engage multiple components of the DPI-H framework, such as the creation of an SHR to support continuity of care, remote consultation, and streamlined data collection, would necessarily involve deploying multiple parts of DPI-H together. However, these investments are most effective when they follow an enterprise architecture approach with a focus on interoperable and standardized products. Where countries may already have underlying DPI such as payments or identity infrastructure, additional investments should include integration with the existing DPIs.

Illustrative investments include:

- Coordinating donor investments and collaborative design of large, multi-year investments following an implementation roadmap for multiple, interoperable DPI-H components. This kind of investment is only appropriate where countries have already developed a health enterprise architecture or similar plan outlining how DPI-H will map to country-specific needs.
- Ensuring applications are tested, modified, and configured in appropriate test environments and sandbox environments prior to going live will enable important learnings and lower cost of necessary iterations.
- Setting up technical assistance units with strong business models that can assist countries with the implementation and configuration of DPGs to function as DPI-H and facilitate robust country-level feedback to global product development communities.
- Investing in the monitoring of the impact of services enabled by DPI-H across key user-segments to ensure equitable, inclusive impact for end-users of health systems.
- Supporting research and evaluation on specific applications and platform performance in different country environments to support context-appropriate investments.

4.4 Order of Magnitude Cost Estimates

The approaches outlined above vary in the magnitude of needed investment and the duration of investment. Several may be accomplished with investments made over the next five years to develop and strengthen key assets, which, once existing, may not require significant recurrent investment. Others are best suited to a consistent funding approach of smaller amounts on a recurring basis. Table 3 provides an overview of the order of magnitude costing associated with several of the illustrative investments identified for each approach above.

Table 3: Order of Magnitude Cost Estimates

Investment Strategy	Illustrative Activities	Magnitude of Investment*	Illustrative Investment
Approach 1: Strengthen Global Ecosystem for DPI-H			
Consistent core product funding to make DPGs and GGs "infrastructure ready"	Support an organization or governance body to administer consistent funding for select DPGs to improve shelf-readiness of product and support its implementation as DPI-H	++	$2 million per product per implementation per year
Fill landscape gaps to operationalize health data security frameworks	Invest in tools to operationalize consent in a health information exchange and socialize implications of data sharing with client communities	+	$50,000-$100,000 per year for next 5 years
Strengthen the global ecosystem of Supporting DPGs	Increase availability of open content and guidelines, open data, open AI models in global market of DPGs to enable DPI-H to support a broader range of digital health priorities	++	up to $5 million per year for next 5 years
Strengthen regulation and governance for predictive analytics in global health	Support operationalization of global guidance on governance of AI models in global health, including resources to support evaluation, documentation, evidence generation, and standards for use with DPI-H	+	up to $1 million per year for next 5 years
Approach 2: Strengthen Country Capacity for DPI-H Implementation			
Continued leadership training and advocacy capacity to motivate adoption of DPI-H	Support leadership-level training on planning national digital systems using DPI-H	+	up to $1 million per country per year
	Draft core documents such as a digital health strategy, health enterprise architecture, costed national plan	++	$2 million over 2 years per country
Strengthen technical capacity to implement DPI-H	Develop additional training materials and reusable technical assistance resources for implementation of DPGs as DPI-H	+	one time investment up to $1 million
	Offer periodic health data standards training courses including FHIR, SMART Guidelines, HL7 and others to build workforce fluent in health data standards and able to implement DPGs as DPI-H	++	up to $1 million per country per year
	Partner with university and job training programs to integrate DPI-H implementation-focused curricula	++	$1-5 million multi-year investment
Strengthen regulatory capacity to shape DPI-H ecosystem	Support training and capacity development to draft standards, regulations, effectively evaluate products against approved strategies, architecture, and requirements, working with private sector	++	$2 million per country per year
Approach 3: Coordinate Donor Investments for Implementation of Multiple DPI-H Components at Country-level			
Implement DPI-H at Country Level	Coordinate donor funding to implement a package of shelf-ready DPGs as DPI-H; includes implementation of a set of DPI-H components, founding of or augmenting support for necessary regulatory and technical support bodies, and integration of DPI-H with existing DPI	+++	$50-$250 million over 5 years per country, depending on size and existing fragmentation

*Legend: + = < 1 million, ++ = 1-10 million, +++ =10-100+ million

4.5 Associated Risks

The investment approaches above reinforce the direction of investments to harmonize and coordinate approaches to supporting digital health systems. However, the approaches include several uncertainties and assumptions that have important implications for the long-term impact of DPI-H implementation.

First, the recommendations above assume that not all DPG products will become DPI-H. In supporting investments in core product development to support use as DPI-H, the global health community will need to endorse clear expectations for the criteria products need to meet. Digital Square's definition of "shelf-readiness" is a starting place; over time, the global digital health community may refine this definition following country experience with DPI-H product implementations and their ability to support other solutions to leverage the DPI-H functionality.

Investments to strengthen the global DPI-H ecosystem may consolidate the market of available DPG products, resulting in some existing products being phased out if they are too resource intensive to improve. This market evolution will need a consistent, transparent, qualified, and objective process for allocating funding and assessing appropriate use of funds to maintain and encourage adequate diversity of products ready for use as DPI-H. Assessing the impact of this support will be an important accountability measure, as it should correlate with advances in country-level digital maturity.

Second, directly attributing changes in health outcomes to continued country-level investments in leadership and governance is likely to remain challenging. As opposed to service delivery or app creation, investments in leadership and technical training have a more complex theory of change, which may be more difficult to link to typical health outcomes. Continuing these investments assumes funders will accept monitoring and evaluation approaches that seek contribution to broader systems change indicators rather than direct attribution to specific health outcomes.

Third, DPI-H investments at the country level are necessary, yet there may be further efficiencies gained by standardizing DPI-H across regions. For example, regionally governed DPI-H could enable cross-border care for populations that are increasingly mobile or transnational sharing of disease surveillance data before and during pandemics. Development of country-level DPI-H should be guided by globally recognized standards, where possible, and aim to limit potential fragmentation between countries to ease cross-border and regional efforts to share DPI-H for selected use cases.

Finally, like all digital health interventions, DPI-H will rely on underlying technology infrastructure, including consistent power and connectivity at central-level and decentralized tools that work in both online and offline connectivity contexts. As digital divides persist across levels of urbanization, wealth, gender, age, and other dynamics, continued attention to broader digital inclusion and equity in access to DPI-H-enabled services will remain critically important throughout implementation.

Even with these uncertainties and underlying assumptions, failing to make investments to facilitate DPI-H implementation may be the larger risk. Lack of support for development and adoption of DPI-H risks worsening fragmentation and increasing the cost of operations without improving outcomes. Currently, countries are making investments in digital technologies as part of broader digital transformation efforts and without interventions to encourage a DPI-H approach, the global digital health landscape will become more fragmented. The longer this persists, the more difficult it will be to change course. Following the approaches outlined in this report will create the conditions that enable the development, adoption, and maintenance of DPI-H. Successful DPI-H implementations can thus become a viable alternative to an increasingly fragmented landscape.

Conclusion

The investment approaches identified above align with and reinforce the priorities of many stakeholders in the global health community. Aligning with country strategies, supporting an enterprise architecture approach, focusing on data exchange, and using open-source systems as key enablers of many health interventions are approaches that are championed in USAID's Vision for Action in Digital Health, CDC's Global Digital Health Strategy, and WHO's Global Strategy on Digital Health, and codified in the Digital Investment Principles.

Viewing the future needs of digital health through the lens of DPI-H offers a clear value proposition to align many varied interests: by supporting basic digital components that underlie many different functional requirements, DPI-H can support an ecosystem that reduces complexity, lowers the cost of implementing new systems, and enables local actors to efficiently innovate and manage products that evolve with system needs. It will support systems and applications across disease verticals and, when developed in coordination with foundational DPI, across sectors.

The last decade of investment in global health successfully promoted a vision for coordinated, planned digital systems in which DPI-H provides the foundation. Now is the time to support the operationalization of that vision through coordinated investments at the global and country levels. Working together, the global digital health community can support digital health ecosystems to effectively leverage DPI-H and accelerate progress toward the health outcomes that we all seek to achieve.

Appendix: DPI-H Considerations for Predictive Analytics

What is predictive analytics? Predictive analytics is a generic term that may encompass traditional statistical modeling approaches like linear regression as well as the use of machine learning models that are built from training data and machine-learning algorithms. It may also include capabilities such as natural language processing and computer vision that are often termed "AI." For the predictive analytics use cases identified in the user journeys in Part 1, the functionalities are assumed to use machine learning models to predict an outcome or natural language processing to generate a text response.

What's different about machine learning and AI? Traditional statistical models are built with predefined mathematical rules and applied to existing datasets. These models can produce probabilistic estimates or predictions and do not require large training datasets and significant computational power to create the model.

Machine learning models differ from traditional statistical models in that they are created by using learning algorithms and large sets of training data to enable a computer to develop the rules that define a machine learning model. Once the model is built, it can be applied to "live" data (data not used to train the model) to generate predictions based on the patterns it learned from the training data. Predictive learning models all require training data to identify the patterns that will dictate future predictions. Data needs vary depending on the algorithms used. Many machine-learning algorithms use quantitative data; natural language processing capabilities require textual data or "text corpora" to train models and predict words in sequence, and computer vision approaches use deep learning algorithms to analyze pixels in image data. In addition to training data, the process of developing machine learning algorithms requires infrastructure to store large volumes of data and the computing power to train models and to run trained models on "live" data. However, this report focuses on health-specific DPGs and building blocks related to the training and reuse of models rather than on the underlying server and network infrastructure.

Digital Public Infrastructure for Health: Considerations for Training AI Models. Training data must be representative of the context in which the model will be used if it is to be effective, so this report considers data requirements of finding locally representative data sets for LMICs (generally, data that is generated within a health system very similar to that in which the trained model will be used). It also focuses on the data repositories that facilitate use of data collected through operational data systems and recognizes the role of Supporting DPGs such as health analytics platforms to develop ML models.

Digital Public Infrastructure for Health: Consideration for Using AI Models. For the reuse of trained models, this report considers the availability of models and supporting documentation (either through open platforms or a licensing agreement) as well as the platforms needed to tailor pre-existing models with local data and run those models with live data. It also considers the governance implications associated with the responsible use of machine learning, which entail defining standards for model performance, evidence of its effectiveness in achieving a desired outcome, and documentation of limitations and potential biases stemming for the training data or model development process.

For this paper and the functionalities explored within it, generative AI does not feature prominently. Generative AI is a specific type of AI that uses a different learning process to enable it to generate entirely new content based on prompts. Examples of generative AI are large language models (LLMs) like GPT-3 and image-generators such as DALL-E. The underlying infrastructure considerations would be similar to those discussed above, except for significantly greater computing capacity requirements to train such models. In many domain-specific functions, an existing LLM would likely be re-used and potentially primed or retrained on a local dataset, rather than start from developing a new base LLM. For outputs of LLMs to be trustworthy, it would need text corpora representative of local context to be trained on, significant computational power to train and run, and governance mechanisms to review performance, bias, and trustworthiness of results generated. The use of LLMs in healthcare and service delivery is newly emerging and will need significant iteration to overcome doubts about the appropriateness and trustworthiness of LLMs in clinical settings, which may extend the timeline and limit the scope of their use in digital health service delivery.

Acknowledgements

This report was produced by Vital Wave with thought-partnership from PATH and Digital Square and support from Co-Develop. The authors would like to thank the digital health and country experts who shared their time and insights during the development of this report.