



InfoCure

Creating Health Information Integrity

- ➔ **Strengthening reliable information:**
InfoCure is developing a digital trust infrastructure designed to improve access to trustworthy health information.
- ➔ **Independently assessing quality:**
Certification focuses on the quality management of digital health information services, assessed against clear requirements for their structures, processes, and production practices.
- ➔ **Making trust machine-readable:**
The certificate enables search engines, social networks, and AI systems to automatically identify and use verified quality attributes.
- ➔ **Testing regionally, developing globally:**
The pilot in Germany, Austria, and Switzerland is intended to lay the groundwork for quality standards that can be applied worldwide.

Summary

Never before have so many people been able to access health knowledge so quickly. Search engines, digital platforms, and AI systems offer an opportunity to make information available precisely where people make decisions about their health. Yet for this technological progress to improve population health, the quality of digital health information must be assured and made recognizable to machines and IT systems. Search engines, social networks, and AI systems exert considerable influence over the information we find and the claims that capture our attention.

Many people find it difficult to critically assess health information and judge the credibility of its sources. At the same time, misinformation and disinformation reach large audiences at great speed in a digital information environment where popularity, emotional intensity, and polarization often carry more weight and attract greater visibility than substantive quality. The consequences extend to patient safety, public health, and trust in public institutions and healthcare organizations.

Correcting and moderating misleading content remain essential to addressing health misinformation and disinformation, as do public awareness and education. But the response cannot be confined to rebutting individual false claims after they have spread. A resilient information environment must also strengthen trustworthy information and make it visible and accessible. This requires an infrastructure that identifies quality-oriented information services, independently assesses the structures and processes through which they assure quality, and publishes the findings in machine-readable form.

This is where InfoCure comes in. The nonprofit initiative is developing an internationally oriented certification system for digital health information providers, translating the results of independent certification into a machine-readable trust signal: a dataset of verified quality information that digital platforms can process automatically. Certification does not validate every individual statement, nor does it confer blanket approval on an organization. It applies instead to a clearly defined digital health information service, such as a website, app, database, social media service, or chatbot.

The assessment focuses on quality management: the structures, processes, and production practices that shape the quality of a service. These include responsibilities, research and approval procedures, the handling of evidence and conflicts of interest, updating processes, moderation, technical reliability, and accessibility. Individual content items are sampled to establish whether the procedures described are actually followed in day-to-day practice. The certificate therefore does not attest to the accuracy of every statement. Rather, it demonstrates that a service meets specified, transparent minimum requirements for quality management.

An adaptive indicator framework forms the methodological core of the certification system. Each service is first described in a structured profile: Which topics, target audiences, languages, components, such as a text-based website, chatbot, or video service, and production models, such as editorially produced versus machine-generated content, does it encompass? Cross-cutting quality requirements then apply to all services, supplemented by the production- and component-specific indicators relevant to each one. Editorially produced content is therefore assessed against different criteria from generative AI applications, databases, or user-generated services. This modular approach allows for differentiation while retaining a common quality baseline.

The certification model avoids a false precision that would be of little value to digital platforms. InfoCure deliberately does not produce a single aggregate score. Instead, it combines mandatory minimum standards, for which shortcomings cannot be offset by strengths elsewhere, with unweighted quality attributes. This preserves a clear picture of each service's specific strengths. Currency, accessibility,

or methodological transparency may matter differently depending on the context of use. Certification bodies therefore provide verified digital quality signals with precisely defined meanings. Individual platforms and systems decide for themselves whether and how to incorporate these signals into ranking, recommendations, labeling, or curation.

A crucial step is to make the certificate technically usable. Different representations are derived from a single certificate record through adapters aligned with established technical standards and vocabularies, including Schema.org and W3C Verifiable Credentials. Stable identifiers, clearly defined validity periods, a central source of status information, and reliable suspension and revocation mechanisms are intended to prevent expired, tampered-with, or invalid certificates from being processed as positive signals.

The division of responsibilities remains clear. Certification bodies assess compliance with specified quality requirements. Operators of digital platforms, including search engines and AI systems, alone decide whether and how to use the resulting signals for ranking, recommendations, labeling, or curation. Institutionally, InfoCure is designed as a federated system. A global core standard is intended to ensure internationally comparable minimum requirements and procedural rules, while national certification bodies take account of the languages, legal frameworks, health systems, and information cultures in their respective contexts. Germany, Austria, and Switzerland will initially provide a shared setting for development and learning.

In coordination with Gesundheit Österreich GmbH and the Careum Foundation, the Bertelsmann Stiftung is developing a prototype designed for international adoption. Indicators, assessment rules, audit guidance, the certification workflow, technical infrastructure, and certificate formats are being developed and tested incrementally. They are intended to provide the foundation for subsequent international scaling. Close collaboration with the World Health Organization (WHO) and the International Telecommunication Union (ITU) is being sought to support this process.


1. The Crisis of Trust in the Digital Health Information Ecosystem

To use health information safely, people need to be able to find it, understand it, assess it critically, verify its sources, and apply what they learn to their own circumstances. For many, the transition from a plausible-sounding statement to a well-informed decision is particularly challenging. Health information is often complex, must be interpreted in light of individual circumstances, rests on statistical probabilities, and may offer no single, unequivocal truth or answer. Consumers can therefore struggle to distinguish a scientifically grounded claim from one that is simply well crafted but ultimately serves commercial or ideological interests.

This uncertainty is far from marginal. Empirical research on population health literacy shows that more than two-thirds of people experience substantial difficulties with digital health information. Assessing whether information is credible and of sound quality is particularly demanding. Limited health literacy is therefore not merely a matter of individual education. It creates a structural vulnerability in health systems and online public discourse. Where differences in quality cannot be reliably identified, emotional appeal, familiarity, or algorithmic reach can take the place of substantive expertise.

The consequences go well beyond a misguided search. Health-related disinformation can encourage risky behavior, undermine effective prevention, and delay necessary treatment. At its worst, it can put lives at risk. It also erodes trust in science, public institutions, and organizations responsible for healthcare and prevention. The harm is not confined to individual users: it can weaken the capacity of health systems to act, reduce acceptance of preventive measures, and make it harder for societies to reach a shared understanding of health risks.

FIGURE 1

Strategies for Curbing Misinformation and Disinformation

Digital transformation intensifies this challenge in two ways. First, platforms, search engines, and social networks have become primary gateways to health information for much of the population. They organize sources, prioritize content, and thereby shape what people see. Second, generative AI is changing the production of information itself. Chatbots select sources, synthesize content, and generate new health-related responses. This creates a new governance challenge: risks must be contained, while search, recommendation, and generative systems need reliable indicators of the quality and trustworthiness of their sources.

Reliable information enables informed choices about nutrition, physical activity, prevention, and treatment. It helps people put risks into perspective, understand therapies, and strengthen their ability to manage illness. Conversely, unequal access to understandable, dependable information can deepen health inequalities. Information quality is therefore not a secondary communications concern. It is an integral part of contemporary public health.

There is no single solution to misinformation and disinformation (see Figure 1). Fact-checking refutes specific claims. Prebunking prepares people to recognize common manipulation techniques and strengthens public judgment. Content moderation labels, restricts, or removes content that violates platform rules. InfoCure replaces none of these approaches. It complements them by enabling trust-based visibility: reliable providers are to be identified in advance, and their quality attributes made recognizable so that digital systems can take them into account in search, recommendations, curation, and content generation.

2. The InfoCure Approach: Making Reliable Sources More Visible

InfoCure shifts attention from the problem itself to the conditions under which reliable information is produced. It takes a preventive approach intended to strengthen the resilience of the information ecosystem. The initiative aims to improve public access to dependable health information, contributing to public health, protection against the health risks posed by misinformation and disinformation, and consumer protection. At its operational core is an

internationally scalable certification system for digital health information providers.

The system is designed to describe and assess both editorially produced, human-authored information services and those based on machine-generated content, databases, or user contributions. Initially, however, the focus is on editorially produced information, which is also shared on social networks and serves as a source endpoint for AI language models. The scope is intended to expand progressively over time.

During the preparatory phase, existing tools for assessing the quality of health information were analyzed, and quality principles were translated into criteria from which auditable indicators could be developed step by step. In this way, abstract principles and criteria become operational requirements that can be assessed transparently in an audit. This work builds on the Global Principles for Identifying Credible Sources of Health Information, developed by the US National Academy of Medicine in collaboration with the World Health Organization and published in 2021.


InfoCure has further developed these principles and translated them into a more detailed set of criteria. The work also draws on established tools from the German-speaking countries: the German Health Literacy Network's Quality Criteria for Reliable Health Portals; the MAPPinfo - Mapping the Quality of Health Information checklist from Stiftung Gesundheitswissen; and the German Network for Evidence-based Medicine's Good Practice in Health Information criteria. Expert input from the Nature Medicine Commission on Quality Health Information for All strengthens both the framework's international relevance and its empirical foundations.

The indicators underpin a certification approach focused on structures, processes, and production practices. Providers must explain, for example, how they research information, subject it to expert review, approve it, keep it up to date, moderate it, and make it technically available. A multistage audit examines whether these procedures are robust and are actually applied in practice. Where the mandatory minimum requirements are met, the information service receives a certificate. Designed as a standardized dataset, the certificate enables providers to communicate the quality of their certified service to search engines, platforms, and AI systems.

Whether the signal is ultimately used, how it is processed, and what weight it is given remain matters for the platforms themselves. They can incorporate the trust signal into selection, prioritization, and recommendation systems. This could help displace high-risk misinformation from especially prominent positions. InfoCure does not, however, guarantee any particular ranking or intervene in platform algorithms. It creates the conditions for digital intermediaries, such as AI systems, search engines, and social networks, to integrate independently assessed sources into their own systems.

The strategic distinction lies in connecting two domains that have often been treated separately: scientific evidence, institutional accountability, and independent assessment on the one hand, and algorithmically mediated information environments in which machines identify and organize sources on the other. InfoCure translates certification findings into a machine-readable signal that can be processed within digital information environments. It thus bridges quality assurance and algorithmic information processing, extending the focus beyond countering harmful content to a complementary governance question: How can reliable information services be strengthened systematically?

FIGURE 2

Scope of the Certificate

3. What Certification Covers: Scope and Limitations


Certification is available to organizations that operate, and are legally responsible for, a digital health information service (see Figure 2). These may be companies, associations, foundations, or public institutions. What matters is that responsibility can be clearly assigned to an identifiable legal entity. Certification does not, however, apply to that organization as a whole. It covers a specifically designated information service, such as a particular website, app, or platform. Where an organization operates several services, these must generally be assessed separately.

A service may comprise several components. A website, for example, may be linked to a newsletter, database, social media channel, podcast, or chatbot. The certificate must define precisely which components, production models, languages, target audiences, and access channels fall within its scope. This boundary primarily protects users: it prevents the quality assurance attached to an assessed service from being implicitly extended to other, unassessed activities of the same organization.

InfoCure defines health information through a broad understanding of health. Services covering health, well-being, diseases, and their causes and consequences are eligible for certification. This includes prevention, diagnostics, treatment, therapies, and medicines. The scope also encompasses health and social care systems, nursing care and rehabilitation, health-related products and services, and legal and financial matters. Research and development, as well as living with chronic conditions or disabilities, may also be covered. Together, these topics form each service's individual subject profile.

Digital health information serves an equally wide range of purposes. It can inform, educate, raise awareness, motivate, or encourage people to take action that supports their health. Some services bring information together to support research, orientation, or navigation. Others offer individualized advice, support self-management, or enable exchange through forums, messaging groups, and online communities. Others again pursue legitimate organizational interests, maintain relationships with target audiences and stakeholders, or facilitate access to products, services, and other transactions. Certification must therefore consider not only what a service discusses, but also what it is intended to do.

FIGURE 3

Overview of Content Production Models

InfoCure distinguishes four basic content production models (see Figure 3). Editorially produced information is researched and reviewed by authors, journalists, or editorial teams. Machine-generated information is produced through rules-based or AI-enabled systems. Database-driven information draws on structured data that are collected, organized, and retrieved for specific purposes. User-generated information comes from individual users or communities and may be moderated. Hybrid services combine several models, for example articles, a searchable database, a community, and a chatbot.

The delivery model describes the components and formats through which information reaches its intended audiences. These include text, graphics, images, audio, and video, as well as databases, e-learning services, and interactive tools. Some components provide information in one direction; others personalize it or allow users to contribute actively. A service might distribute editorial articles through a website, newsletter, and podcast while also offering a conversational chatbot. It is the combination of these components that constitutes the digital service as users actually encounter it.

Certification centers on the system through which health information is produced: the structures, processes, and production practices behind the service.

These encompass content creation, expert review, approval, publication, updating, moderation of user-generated content, and technical delivery. Where a service uses several production models, the requirements for all of them are taken into account. Individual content items are examined only on a sample basis to verify that the procedures are being implemented. The assessment does not address how a particular person subjectively perceives the information or the effect it has on that individual. InfoCure therefore certifies the institutional conditions for reliable information, not every statement or every possible situation in which it might be used.

4. The Adaptive Indicator Framework

A rigid set of criteria would not do justice to the diversity of digital health information services. An editorial article is produced under different conditions from a generative chatbot response; a curated database carries different risks from an open community. InfoCure therefore combines a common quality core with adaptive modules. The indicator framework consists of three coordinated elements: a descriptive service profile, cross-cutting quality requirements, and production- or component-specific indicators.

The starting point is a robust profile of the service. It records production and delivery models, topics, target audiences, languages, conditions of access, the pricing model, opportunities for interaction and personalization, funding, and data processing. Existing accessibility features are also described. This information does not yet constitute a quality assessment. Its purpose is to characterize the service precisely and identify the indicators that apply in each case.

This detailed service profile is followed by cross-cutting requirements that apply, in principle, to all digital health information services. They primarily concern structures, processes, and production practices, including a demonstrable commitment to quality, clearly assigned responsibilities, appropriately qualified staff, and other relevant aspects. These requirements establish a common quality baseline, regardless of how content is produced and delivered.

The specific modules address the points at which risks diverge. For information researched by people and curated editorially, transparent research methods, expert review, clear source attribution, and defined approval procedures are essential. For generative AI, the emphasis shifts to the quality of training and underlying data, human oversight, the labeling of automatically generated content, and procedures for handling erroneous outputs. Database-driven services must ensure data quality, sound organization, and regular updating. For user-generated content, the priorities are moderation, accountability, and the handling of problematic contributions.

The service as it actually operates determines which modules apply. Where several production models or components are used, multiple sets of specific indicators are applied in parallel. A hybrid service is therefore not forced into a uniform template but assessed according to its actual functions and the way its content is produced and delivered. At the same time, certification status remains comparable because the cross-cutting minimum requirements apply to all services.

A comprehensive review of every text, video, database entry, or conversational output would be neither scalable nor financially sustainable, and it could not be kept continuously up to date. It would also shift InfoCure's central purpose toward that of a universal fact-checking system. Sampling nevertheless remains indispensable. Where a provider describes formal research and updating procedures, selected content must demonstrate that those procedures are visibly and consistently applied.

To turn quality principles into a robust audit, each indicator must be operationalized. Every assessment item is defined by a guiding question, appropriate evidence, and clear assessment benchmarks. Evidence may include internal or publicly available policies, methodological descriptions, minutes, reports, documented responsibilities, and examples of content. Consistent guidance for auditors prevents the same circumstances from being assessed differently. The indicator framework is therefore more than a catalog of good intentions: it provides the methodological basis for replicable, reliable, and internationally comparable conformity assessments.

The framework's modular architecture also supports its continued development. New production models, formats, quality attributes, or topic-specific risks can be added without redesigning the entire standard. In an environment where AI applications, platform features, and media formats evolve rapidly, this extensibility and agility are essential to maintaining the certification system's methodological precision and technological compatibility.

5. From Certificate to Digital Trust Infrastructure

A certificate that exists only as a PDF in a download folder does not, by itself, change a digital information environment. For InfoCure to have an impact, assessment findings must be understandable, unambiguous, and machine-readable. The certificate is therefore more than a document: it is a structured dataset that reliably records a service's certification standing, verified quality attributes, scope, and current status. Providers can use it to demonstrate the independently assessed quality of their information service to search engines, platforms, and AI

systems. These systems can then read the additional trust signal and take it into account when selecting and assessing sources.

A single aggregate score would be ill suited to this purpose because the importance of a quality attribute depends on the context in which it is used. A search engine might place particular weight on currency; an accessible service on clarity and accessibility; and an AI system on traceable evidence and institutional accountability. Reducing all these dimensions to one number would require the certification body to make normative weighting decisions, sacrifice semantic detail, and impose a hierarchy that would not suit every use case. Exceptional strengths could also mask serious weaknesses, for example in scientific rigor, neutrality, or independence.

InfoCure therefore separates the question of whether a service qualifies for certification from the differentiated description of its additional quality attributes. This leads to an architecture with three substantive layers, combining a clear trust threshold with granular information without assigning weights:

- **The first layer consists of Mandatory Standards.** Every mandatory requirement must be met. A failure to meet a non-compensable minimum standard cannot be offset by particular strengths elsewhere. The outcome is unambiguous: met or not met. Platforms can use this layer as a threshold, a filter, or a trigger for labeling.
- **The second layer comprises Quality Attributes.** It describes features such as currency, multilingual provision, clarity, accessibility, user orientation, and other dimensions that may be present to varying degrees. These attributes are reported individually in machine-readable form rather than condensed into a score. This preserves the specific strengths of each service and allows external systems to apply their own context-sensitive weighting.
- **The third layer consists of Rule-Based Badges.** These can recognize distinctive achievements, such as above-average quality, a strong focus on a particular audience, extensive

personalization options, or specific evidence of quality in individual areas. Badges are based on transparent, verifiable rules. They complement the minimum standards but do not replace them. The attributes on which they are based must remain traceable at all times.


This substantive framework is represented in a clearly structured certificate record. Identity and attribution data include, in particular, the certificate ID, provider ID, name of the responsible legal entity, issuer ID, issuing body, region or jurisdiction, and applicable policy version. Status and validity data indicate whether a certificate is valid, expired, suspended, or revoked. They include the start and end of the validity period, the dates of the most recent and next review, and suspension or revocation status where applicable. Quality data record the Mandatory Standards, Quality Attributes, and Rule-Based Badges. Verification data point to the authoritative source of status information and may include signature or hash references.

The technical architecture follows a simple principle: creating the substantive certificate record, delivering it in different formats, and verifying its current status are separate functions. All representations nevertheless derive from a single authoritative data model. This canonical certificate prevents any website seals, registry entries, machine-readable metadata, and verification services from making inconsistent claims about the same service. Security, data protection, traceable changes, and the need to limit critical dependencies are built in from the outset.

The Certification component forms the protected core of the certification process. It receives applications, service information, and supporting evidence; documents audits; tracks improvement measures; and records certification decisions. Its final output is the complete, authoritative certificate record. Because this component may process sensitive internal documents, personal data, and business-critical information, clearly defined roles, secure access permissions, and a complete procedural record are indispensable. It is therefore supported by proven audit-management software that is itself certified.

FIGURE 4

The Machine-Readable Trust Signal



The Delivery component takes the approved certificate and converts it into different technical formats without making any new judgments about quality (see Figure 4). It enables different systems to use the same information. Depending on the intended use, adapters can align with standards such as Schema.org and JSON-LD, W3C Verifiable Credentials, Open Badges, or C2PA. A change in a receiving system's technical requirements does not require the entire architecture to be rebuilt. Only the relevant adapter needs to be adjusted; the canonical certificate remains the authoritative source of the substantive information.

Certificate information is then made available where the receiving systems can actually detect the trust signal: within the certified information service itself. Providers must therefore integrate technical elements of the certificate into their own services. This may involve embedding structured data using Schema.org or JSON-LD in a website, linking a Verifiable Credential to a platform or provider account, or making certificate information available in other supported formats. InfoCure provides the technical foundations required for this integration.

The Verification component provides a public, authoritative status response. Platforms and other third-party systems must be able to determine

automatically whether a certificate is currently valid, expired, suspended, or revoked, which version applies, and who issued it. Verification does not constitute a new quality assessment. It confirms only the status of a certification decision already made. Since many systems may request this information simultaneously, the service must be lightweight, highly available, and cache-friendly. Authority over validity, suspension, and revocation remains with the responsible InfoCure technical trust authority.

For major digital platform operators, including those running search engines and AI models, the barriers to integrating the machine-readable trust signal are low. They do not need to establish their own audit procedures, acquire specialist expertise, or implement bespoke proprietary solutions to read the certificate. The certification system provides stable provider and certificate IDs, a reliable registry, defined validity periods, transparent revocation logic, and standardized trust endpoints: trusted interfaces for identity verification.

The InfoCure certificate supplies structured, machine-readable trust signals. Platforms nevertheless retain full discretion at all times over whether and how to integrate them into search, recommendations, AI source selection, labeling, or

knowledge systems. Conversely, platform operators must not control the quality standard, the recognition of certification bodies, or individual certification decisions. Independent assessment and platform autonomy are the cornerstones of InfoCure's trust architecture.

6. The Certification Process

Every trust signal rests on a controlled, transparent process (see Figure 5). Certification begins with an inquiry from a provider. The first step is to establish which specific health information service is to be assessed, which components it includes, and whether it is eligible for certification in principle. The scope of work, timeline, and anticipated effort are then agreed. Once a contract is in place, a suitably qualified auditor is selected, with any potential conflicts of interest ruled out.

The provider then receives access to a secure software portal, where it describes its service and submits the required evidence. The information collected covers production and delivery models, topics, languages, target audiences, conditions of access, the pricing model, interaction and personalization options, data protection, funding, and accessibility. These descriptive details determine which modules of the adaptive indicator framework apply. The cross-cutting quality requirements are assessed as well.


Evidence includes methodological descriptions, internal policies, documented responsibilities, procedural instructions, and records of content creation, expert review, and updating. The certification body first checks whether the information is complete and plausible. It requests any missing information or clarification of inconsistencies. Because the process involves sensitive internal documents and potentially business-critical information, it must comply with data protection requirements, use role-based access, and maintain a secure audit trail.

The audit itself has two phases. In the first, the self-reporting phase, submitted documents are checked for completeness. The materials are then reviewed, content samples examined, and information gathered for the second phase. For each quality criterion, the auditor records the evidence available and whether the requirement is met. An interim report identifies unresolved questions, nonconformities, and necessary improvements.

In the second phase, auditors interview the staff responsible. They examine how the stated quality processes are implemented in everyday practice and, where needed, review additional samples and internal evidence. Together, the self-reported information and the audit interview constitute the substantive conformity assessment.

ABBILDUNG 5

The Certification Process



Any nonconformities with the certification criteria identified during the audit are documented in the system and communicated to the provider. The provider is given an opportunity to submit corrections or corrective actions. The auditor then reviews the measures and supporting evidence submitted. The audit is completed only once the identified nonconformities have been adequately addressed or it has been established that the requirements are not met. The handling of nonconformities, supplementary submissions, and corrections remains fully traceable in the case file. Once the substantive assessment is finished, the auditor formally closes the audit. The audit report records the findings and provides the basis for the subsequent review and certification decision.

The final certification decision is not made by the person who conducted the audit. An independent panel reviews the documentation under a dual-review principle. It may grant the certificate, defer the decision until further conditions have been met, or refuse certification. Separating assessment from decision-making safeguards independence and impartiality. Both the decision and its rationale are fully documented.

A successful service receives a certificate with a unique identifier, a defined scope, and a specified validity period. It is also entered in a certification registry. The provider receives a final audit report and binding rules governing the use of the certificate and any associated seals. All documents, decisions, and evidence are retained in the case file under version control.

Responsibility does not end when the certificate is issued. Reviews generally take place annually throughout its validity period. Complaints or indications of quality problems may trigger additional checks. Certification status may be maintained, temporarily suspended, or revoked. Before the certificate expires, recertification begins with an updated scope and a new audit plan. The certificate thus remains a current trust signal rather than becoming a one-time, permanent authorization.

7. A Global Core Standard, Implemented Regionally

InfoCure can be effective internationally only if its certificate provides the same fundamental assurance regardless of the country in which it is issued. At the same time, a credible assessment cannot disregard differences in language, legal frameworks, health systems, and information cultures. The long-term objective is therefore not a centralized system in which a single international body directs every audit. InfoCure is intended to grow as a federated network: a common global core implemented by competent national or regional certification bodies.

The core standard establishes the quality principles and procedural rules that must apply everywhere. These include transparency, professional accountability, the handling of evidence, disclosure of interests and funding, reliable updating processes, and traceable editorial and technical procedures. Equally binding are the principles of independent auditing, separation of assessment and decision-making, clearly defined validity periods, and effective mechanisms for surveillance, suspension, and revocation.

Certification in practice should take place where the language, legal environment, healthcare system, and information landscape are understood. National bodies guide providers through the process, review evidence, conduct audits, and can take account of context-specific considerations. This proximity is not a departure from the global standard; it is a condition for applying it credibly. Regional autonomy concerns implementation, not discretionary changes to non-negotiable minimum requirements.

To ensure that certificates remain internationally comparable, certification bodies must themselves be recognized and supervised under common rules. This requires qualified auditors, documented procedures, robust separation of roles and decision-making responsibilities, and consistent assessment methods. Experience from different countries should inform further development without allowing the common core to fragment into competing national variants.

A pluralistic global governance structure is intended to safeguard this balance. It should include national and regional certification bodies, independent experts, public health expertise, and representatives of patient, consumer, and wider public interests. Multilateral and technical organizations should contribute in an advisory capacity. Long-term responsibilities include further developing the core standard, maintaining common procedural rules, recognizing certification bodies, ensuring consistency, and overseeing the global registry and digital trust signal.

The federated model avoids two opposing risks. A centralized structure would struggle to reflect the diversity of national contexts and could quickly be perceived as detached from local realities. Isolated national systems, by contrast, would produce divergent quality assurances and technically incompatible certificates. InfoCure therefore combines international consistency with regional expertise, not as a compromise but as a deliberate organizing principle.

8. Multilateral Governance

An international trust infrastructure for reliable health information requires clear responsibilities and an institutional foundation that supports worldwide legitimacy and adoption. Developing quality indicators, setting technical standards, conducting certification, and using digital trust signals are distinct functions, each with its own requirements and operating conditions. InfoCure therefore envisages a governance architecture in which these functions are clearly separated yet systematically connected.

The proposed international institutional framework for InfoCure builds on WHO's existing mandates to set standards in health and make science-based information accessible in the digital environment. Strengthening a reliable digital information environment is not merely a project objective; it is part of a broader international public health agenda. On the basis of these mandates, WHO could assume the central substantive and institutional role in taking the approach to global scale.

A WHO Technical Reference Group could, for example, be established to assess and refine the regionally developed quality indicators, audit foundations, and certification principles, and translate them into a globally applicable reference framework. This work would draw on perspectives from different world regions and expertise in public health, evidence-based health information, quality assurance, certification, health literacy, digital health, and civil society.

Over time, WHO could also take on a central role within the international certification architecture. One option would be to establish common requirements for national certification bodies and entrust their recognition to WHO. This could combine a globally comparable quality framework with decentralized, context-sensitive implementation. Individual audits and certification decisions would remain at the national or regional level.

Independent certification bodies assess information services against the global quality framework while taking account of the relevant legal, linguistic, and institutional context. Under this model, WHO would not certify individual information services itself. Rather, it would safeguard the common reference framework and, in time, the quality and comparability of the certification architecture. The separation between global standard-setting, recognition or accreditation of certification bodies, and individual certification decisions would be preserved.

In parallel, the technical dimension of the trust infrastructure could be embedded in the international standardization structures of the International Telecommunication Union. The Bertelsmann Stiftung could participate in the relevant ITU working groups and bodies as an Associate or Sector Member. The aim would be to standardize the technical requirements for the machine-readable trust signal internationally, enabling interoperable use across different digital ecosystems.

Key requirements include stable identifiers for providers and information services; a canonical, authoritative certification record; automated status checks; versioning; time-bound validity; revocability; security mechanisms; and standardized, cache-friendly delivery formats. This would create a

technical infrastructure through which search engines, AI applications, social networks, browsers, and other digital services could automatically identify, verify, and process certification information.

Under this proposed model, WHO and ITU would have complementary roles. WHO would establish the globally substantive framework for quality and certification, while ITU would set the technical standards for representing it in machine-readable, interoperable form. National or regional certification bodies could build on these foundations to conduct the assessments. Substantive legitimacy, technical standardization, and operational certification would thus remain institutionally separate while forming part of a shared trust infrastructure.

Responsibility for using the resulting trust signals remains with digital platforms. National certification bodies provide a verified, structured, machine-readable trust signal. Information providers can use it to communicate the quality of their certified service to search engines, AI systems, browsers, and social networks. These platforms may then incorporate the InfoCure signal into their algorithmic processing at their own discretion. This preserves a distinction central to InfoCure: providing standardized quality data is separate from determining how those data are weighted by algorithms in a particular context.

9. From Regional Pilot to Global Certification System

The international ambition deliberately begins at the regional level. Before a global certification system can emerge, the approach must demonstrate its value under real-world conditions. Germany, Austria, and Switzerland therefore constitute a shared setting for development and learning, not the geographic limits of the future system. In coordination with Gesundheit Österreich GmbH and the Careum Foundation in Switzerland, the Bertelsmann Stiftung is developing a prototype that responds to the rapidly growing near-term need for systematic quality assurance while being designed for international adoption from the outset. BSt Gesundheit gGmbH, a wholly owned nonprofit subsidiary of the Bertelsmann Stiftung, is intended to take responsibility for the practical design and operational

implementation of the certification approach and to serve as the national certification body for Germany, Austria, and Switzerland.

The findings from this proof of concept will provide the basis for subsequent international development. The aim is not to export a regionally developed model unchanged to other countries. Rather, the indicators, procedures, and experience gained are to be assessed internationally and developed into a quality framework suitable for global use. That framework must define common minimum requirements while retaining sufficient flexibility to accommodate different linguistic, cultural, legal, and institutional contexts.

An initial regional pilot uses a limited set of core quality indicators. This focus allows the central mechanisms to be tested early under practical conditions, experience to be gathered across different services, and the process to be refined accordingly. During the pilot phase, indicator development is supported, among others, by an international commission: the Nature Medicine Commission on Quality Health Information for All. It brings together international expertise in science, public health, and health communication. This work on internationally applicable quality indicators creates an important link between scientific evidence, practical quality assurance, and subsequent certification.

Development is also supported by the InfoCure Sounding Board. Its members contribute perspectives from academia, patient and consumer advocacy, public administration, civil society, health communication, and the practice of digital information provision in Germany, Austria, and Switzerland. The board advises in particular on quality requirements, institutional arrangements, and the practical design of assessment and certification procedures.

InfoCure also draws on the perspectives of the Health Transformation Hub's 30under40 network. This interdisciplinary network brings together young professionals from politics, academia, business, associations, health insurance, nursing, and patient advocacy, explicitly contributing the perspectives of younger people. The patronage of Stefan Schwartze, the German Federal Government's Patient Commissioner, further underscores the

initiative's commitment to patients' interests and information needs.

The pilot tests more than individual criteria. It examines the entire certification chain: describing the service and defining its scope; submitting information and evidence; selecting the relevant indicators; reviewing documents and conducting audit interviews; reaching an independent certification decision; issuing the certificate; and making it technically available. This will show whether the substantive, organizational, and technical elements of the architecture work together in practice.

The central questions are practical. Are the indicators clearly and unambiguously formulated? Can different providers reliably supply the information required? Can the same requirements be assessed consistently? Does the process adequately reflect the diversity of digital health information services? Can the findings be structured so that search engines, platforms, and AI systems can process them automatically? The pilot in the German-speaking countries is therefore the proof of concept for the entire theory of change.

The lessons learned will feed directly into the next stage of development. The indicator framework is expected to expand incrementally in 2027. Additional providers in the German-speaking countries are to be included and certified. In parallel, assessment rules, audit guidance, the certification workflow, technical infrastructure, and digital certificate formats will be further developed. This will generate practical evidence on applicability, consistency, resource requirements, and technical compatibility.

The next step will not be immediate global scaling, but a carefully managed phase of international validation and capacity-building, undertaken in close coordination with international and multilateral partners. Work in selected countries and regions will establish which elements can be transferred directly and where linguistic, legal, cultural, or institutional adaptations are needed. International scaling involves more than translating a set of criteria. It requires joint scrutiny of quality requirements, the establishment of suitable certification structures, the training and qualification of

auditors, and the testing of technical infrastructure across different information environments.

Over time, the regionally tested approach is intended to evolve into an international standard available to all countries as a global public good. As a multilateral organization, WHO could coordinate this development within its health policy mandate and confer broad international legitimacy on the standard. The objective would be a jointly developed reference framework that can be used worldwide, combining consistent quality requirements with the flexibility needed to respond to national and regional circumstances.

To ensure international transferability, diverse perspectives must be involved early: countries with very different health systems, civil society organizations, patient and consumer representatives, academic institutions, and national public health stakeholders. This dialogue must establish which elements should remain unchanged as part of the global core and which can be tailored to context. The regionally tested model should evolve into a lasting global system only once the standard, audit practice, recognition arrangements, technical infrastructure, and governance function reliably together.

During the development phase, the Bertelsmann Stiftung and its nonprofit subsidiary BSt Gesundheit gGmbH serve as initiators, incubators, and development partners. They are responsible for the regional proof of concept, work with partners to establish the operational foundations of the certification system, and bring the resulting insights into the international processes of WHO and ITU. Their distinctive role is to connect the scientific development of indicators, certification in practice, and technical standardization. As internationalization progresses, responsibility is intended to pass gradually to the designated multilateral and national structures.

10. Strategic Outlook

Not all international roles have yet been formally agreed. The composition of the future governance structure, the institutional home of the secretariat, and the specific forms of cooperation remain under development and discussion. The intended division of functions is nevertheless clear: a common core standard to ensure comparability; national bodies to deliver context-sensitive implementation; and multi-lateral governance to safeguard independence and technical compatibility. Digital platforms, including search engines, browsers, and AI systems, ultimately make the trust signal usable within the information environment.

The proof of concept in the German-speaking countries is not the finished product. It is the first robust building block of a trust infrastructure designed for global adoption. Its strategic significance extends beyond the certification of individual information services: InfoCure makes a tangible contribution to the integrity of health information. Integrity means more than the factual accuracy of individual pieces of content. It encompasses a trustworthy information ecosystem across the entire health information lifecycle: scientific and editorial development, responsible publication, distribution through digital platforms and AI systems, ongoing updating, and, where necessary, correction or withdrawal. Scientific rigor, accuracy, accountability, and transparency are central guiding principles. They are complemented by governance mechanisms that promote equitable access, enable transparent oversight, and protect against manipulation, discrimination, and misuse.

Trust is at the heart of this integrity, but it is neither blanket nor unconditional. Robust trust emerges where quality can be substantiated transparently, responsibility can be clearly assigned, and compliance with common standards can be independently verified. InfoCure creates the institutional and technical conditions for this: providers disclose their structures and processes, an independent body assesses their conformity, and the result is recorded in a verifiable digital certificate. Trust is therefore not merely implied. It is underpinned by transparent criteria, verifiable procedures, and clear accountability.

The particular strategic value of this approach lies in connecting two domains that have often operated separately: institutional quality management and algorithmically mediated visibility. The digital certificate translates the findings of an independent assessment into an interoperable, machine-readable trust signal. Search engines, social networks, browsers, and AI systems can identify the signal automatically and integrate it into their own selection, recommendation, or knowledge systems. InfoCure neither prescribes a uniform weighting for platforms nor makes ranking, recommendation, or moderation decisions itself. Instead, it provides structured, reliable quality information that digital intermediaries, such as search engines and AI chatbots, can use according to context.

InfoCure thereby strengthens information integrity at a critical point. A trustworthy information environment is not created solely by removing false or harmful content. It also depends on making science-based, transparently produced, trustworthy information systematically identifiable, verifiable, and easy to find. InfoCure thus shifts the strategic focus from an exclusively reactive response to problematic content toward actively building reliable information environments.



Publishing Information

© Bertelsmann Stiftung, Gütersloh 2026

Publisher

Bertelsmann Stiftung
Carl-Bertelsmann-Straße 256,
33311 Gütersloh
www.bertelsmann-stiftung.de
trusted-health@bertelsmann-stiftung.de

Program

Healthcare Program

Responsible for content

Dr. Sebastian Schmidt-Kaehler
sebastian.schmidt-kaehler@bertelsmann-stiftung.de

Authors

Dr. Sebastian Schmidt-Kaehler
Senior Advisor

Malte Etienne
Project Manager

Esther Grothoff
Project Assistant

Sina Husemann
Project Manager

Dr. Inga Münch
Senior Project Manager

Inna Rifert
Project Manager

Dr. Tobias Silberzahn
Senior Fellow

AI-assisted text

Citation note

Schmidt-Kaehler, S., Etienne, M., Grothoff, E.,
Husemann, S., Münch, I., Rifert, I. & Silberzahn, T.
(2026). InfoCure - Creating Health Information Integrity.
Bertelsmann Stiftung (Ed.). Gütersloh

Coverimage

Image montage: © Miljan Živković – stock.adobe.com

Discover
InfoCure

