3T — Trans-Forming, Trans-Coding, Trans-Lation

Position paper on the intermediate canonical layer

The intermediate layer in many-to-many health data exchange

Dr. Kai Heitmann · Release 1.33.0 – 22 September 2026

Contents

Summary

Patient summaries already cross European borders through MyHealth@EU, voluntarily and between some Member States. The European Health Data Space Regulation makes that connection obligatory for all of them from March 2029, and adds a second group of data categories from March 2031. One of the tasks for every Member State will therefore be to transfer several inbound formats to several outbound ones. Either every pairing is built and maintained on its own, or everything is routed through one shared description of how the data is structured and what its values mean, an intermediate layer.

Four categories of questions decide which.

- count – how many formats and versions, over how many years?
- costs – what does one translation cost to build and to keep working?
- skills – who is competent to decide what it should say, how scarce are those people?
- risk – what reaches a patient when a translation is wrong?

The third is usually the binding constraint and is rarely costed. The fourth is usually missing from the business case altogether.

Built pairwise, every sending-to-receiving combination is its own translation. Below roughly three formats on each side that is cheaper and stays cheaper. Above it the count turns, sooner where versions of the format specifications revise often. In any case it has to be ensured that clinical review is funded and that an owner can be named.

A shared description demands concentration. An error nobody catches reaches every recipient at once, where a direct translation would have reached one. If it is also operated as a shared service and not only agreed on paper, availability and security have to be guaranteed as well.

What it buys is independence from each other's schedule. The obligation arrives on a fixed date, while the suppliers who have to meet it work to timetables of their own. With a description in the middle, one side can move to a new version while the other has not, and the exchange keeps working in between.

On these criteria, the intermediate layer is the right shape for cross-border exchange under the European Health Data Space. What decides it is the variety of representations brought together at the points of mediation, the continuing movement of the specifications, and the fact that binding deadlines do not synchronise the transition of the systems. It holds only if conditions are met that are often not met today.

The occasion

Patient summaries, electronic prescriptions and the dispensations that follow them already cross borders today. MyHealth@EU carries them between those Member States that have chosen to connect, and that connection is voluntary. Under the European Health Data Space Regulation it stops being voluntary: from March 2029 every Member State has to be connected for the first group — patient summaries, electronic prescriptions and electronic dispensations — and from March 2031 for the second, which adds medical images and image reports, laboratory results and hospital discharge reports.

The first group is three categories rather than one, and the second adds three more. That is not the same as six formats: each category brings its own national formats with versions of their own, so what enters the arithmetic below is formats and versions per category per Member State, not the bare number of categories. The set is also expected to grow.

The patient summary is the prominent case. It is carried today by a document format specified for the present infrastructure, while the exchange format the new Regulation itself requires is still to be adopted in an implementing act. A running exchange therefore has to move from the current specification to the new one, and every national system holds the content for it in its own way.

Two features of that schedule bear directly on the question here. The obligation arrives on a fixed date rather than when a project happens to be ready, and it arrives in stages, so the scope is known in advance to grow. An architecture chosen now will be asked to absorb a second wave of formats two years after the first, and that is a question about the shape of the architecture rather than about its schedule. It also lands on something already running rather than on a blank sheet.

> Mapping is not a one-off project. It is an architectural decision.

In that sense the following discussion is not a hypothetical design exercise.

What is being chosen

Health information has to travel between systems that were never designed for each other: hospital records in one country, a national uniform health record system in a second, a regional exchange platform in a third. Each writes the same clinical facts down in its own way, and typically none of them was built with the others in view.

The choice is between two shapes, also called architectures:

- every sending-to-receiving pair can be built as its own translation, with the pair's two formats wired into it, or
- every format can be translated into one shared description of how the information is structured and what its values mean, which can then be passed on.

The second shape is what this paper calls the *intermediate canonical layer*, or the *intermediate layer* for short. The description at its centre keeps its own name, the *shared description*. The two are not the same thing. The shared description is the model: it fixes how the information is **structured** and what its values **mean**. It is thereby both the place where a wrong decision can sit and the place where a contested one becomes visible.

The intermediate layer is the arrangement around that model, and it answers for all three kinds of transformation. The model can settle two of them, structure and semantics. For language it can carry provenance and the original text along, but the quality of a translation must be secured by the third part. What a clinician reads in their own language is therefore not a model question. The difference returns in the section on the three kinds of transformation.

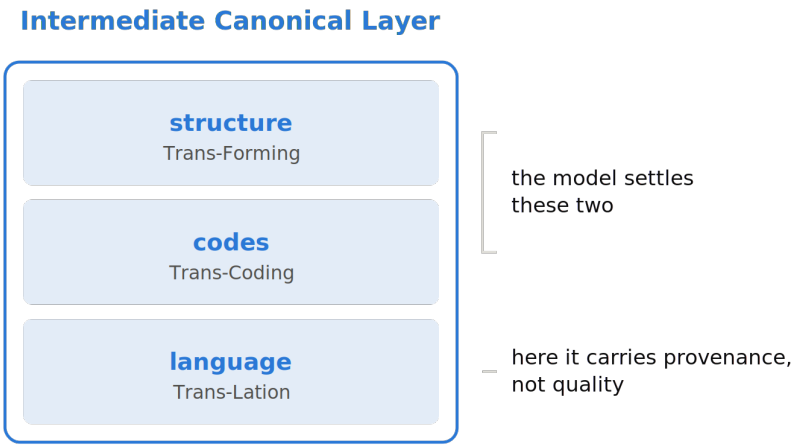

Figure 1 — The intermediate layer and its three compartments, which are the three kinds of transformation. The shared description at its centre fixes the structure and the meaning of the values. For language the layer carries the source language, the original text and the provenance of a translation; the quality of the translation itself is assured elsewhere.

The intermediate layer architecture costs more before it delivers anything and less at every step afterwards. That is why the decision turns on how long the arrangement has to last and how much it has to absorb.

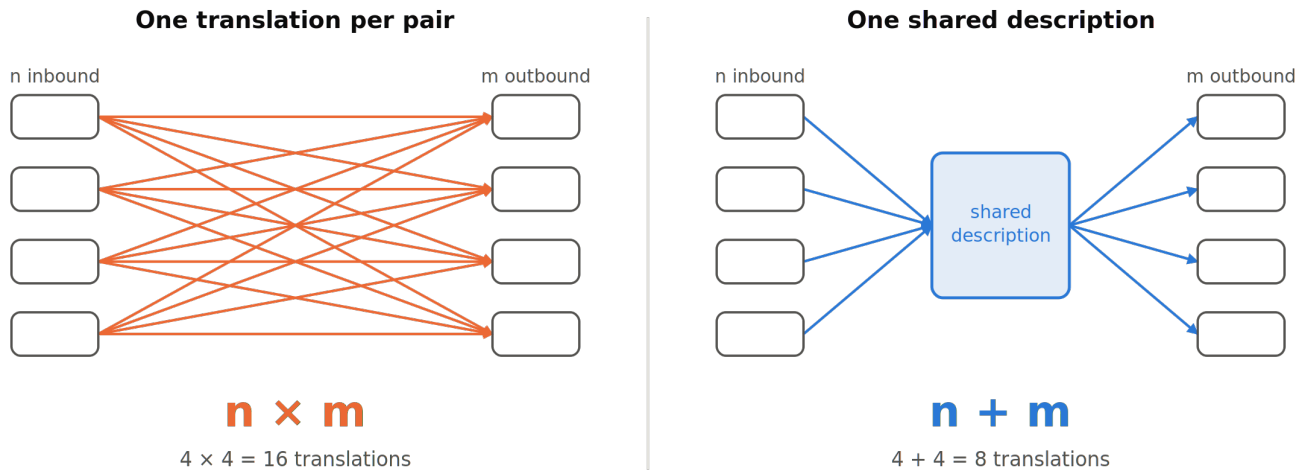

Figure 2 — The two architectures. Four inbound and four outbound formats need sixteen translations when every pair is built on its own, and eight when everything passes through one shared description. The figure counts pieces of work; it does not account for the effort inside each one.

The shared description is a specified information model with a conformance definition: something an implementation can be tested against, rather than a codebase or a database schema. It has to be written in a suitable modelling formalism, and the candidates in this field are the familiar ones — the HL7 FHIR ecosystem with its profiles, terminology and mapping mechanisms, and the European electronic health record exchange format with its established models.

What follows is about criteria for deciding the architecture. Throughout, it has to be borne in mind that this choice is itself a set of specifications with versions of its own. The specifications together with their respective version axes therefore have to be managed dependably.

Four questions decide it

The count. How many formats have to be supported on each side, how often do their versions change, and over how many years?

The cost. What does it cost to develop a translation and to keep it working?

The skills. Who is competent to judge what a translation should say, and how scarce are the people qualified to do so?

The risk. What consequences can a faulty translation have for patients?

The first two can be estimated on a spreadsheet. The third is often where the decisive bottleneck sits: available expertise. Its scarcity is rarely costed. The fourth question helps determine what effort is justified, and it is precisely the one that tends to be left out of the business case.

The answers also depend on what work a translation requires and how demanding that work is. The following section therefore looks at the associated tasks alongside the four criteria.

One — the count

Four incoming and four outgoing formats make sixteen possible pairings. Under direct translation each of them needs its own rules, tests and maintenance. A shared description in the middle reduces the number to eight: four translations inward and four outward. When a fifth format appears on one side, it takes one additional translation rather than four.

The number of formats is not the only thing that counts. Their versions change over time as well. With direct translations, every connection that uses the changed format then has to be reviewed and where necessary adjusted. With a shared description that effort can be confined to a single connection. Because maintenance runs for the whole operating life, frequent version changes can weigh more heavily than the initial number of pairings. The advantage of the intermediate layer can then appear at fewer formats.

Nor does the number of parties alone reflect the effort. One supplier can serve several exchanges with the same product, each with its own requirements. The question of which translations can be reused therefore arises inside a supplier as well.

Two qualifications apply. First, a new version is rarely an entirely new format. Much of it remains compatible. Counting every version as an additional format therefore overstates the effort.

Second, the relief depends on what changes. Where a change concerns only the representation of content that is already mapped, it can be confined to the translation into the shared description. Where a version introduces new clinical content not yet mapped there, the shared description itself has to be extended. The outward translations that are to convey that content then have to be adjusted as well. The intermediate layer does not save this substantive work. How far the change reaches depends on which recipients need the new content and on how the shared description handles extensions and differing versions.

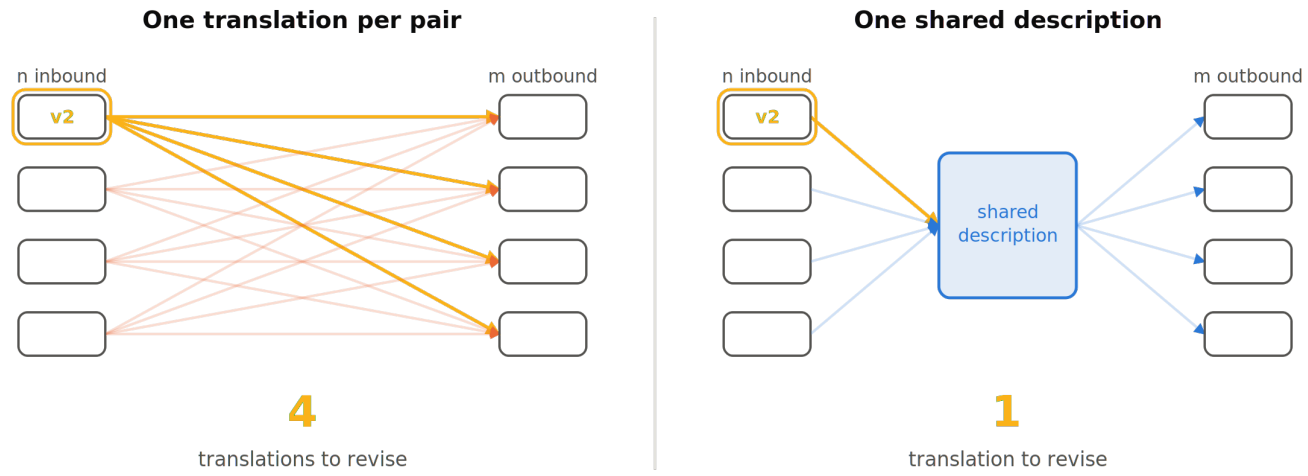

Figure 3 — The version axis. When an incoming format is revised, every connection that uses it has to be reviewed and where necessary adjusted. Built pairwise that can be four; through the shared description the effort can be confined to one. The figure illustrates the mechanism, not a measured saving. The qualification stated above still holds: it concerns changes in how something is written down, not a revision that adds new clinical content.

Two — the cost

A translation is more than program code. It needs a specification, an implementation, test cases, a clinical review and a release. After that it has to be maintained for as long as it stays in service. Development is therefore only part of the total cost. Fewer translations can lower the effort at several points, not only in programming.

The number of pairings alone is not enough to determine the saving. A translation into a shared description can be more demanding than a direct one: it has to map content independently of any single recipient's requirements. To that are added the development, coordination and maintenance of the shared description itself. Eight such translations therefore do not necessarily cost half of what sixteen direct ones cost. Whether and when the additional effort pays off depends on how much work is removed by reuse and by simpler maintenance.

Costs and benefits are also not distributed evenly across the parties. Recipients gain from data out of different sources being delivered according to one shared description. The work required for that can sit elsewhere, with Member States, regional platforms or private providers. A viable funding model therefore has to settle who pays for development, for substantive review and for continuing maintenance. An arrangement that is economic in total does not sustain itself when individual parties carry the cost while the benefit arises mainly somewhere else.

Three — the skills

Whether a transformation is right calls for two kinds of review. Technically it has to be checked that the data arrives in the right place and in the right form. Clinically it has to be judged whether it keeps its meaning on the way. These call for different competences, which are frequently distributed across several people. A technically flawless translation can be clinically wrong without the system reporting a fault.

The shared description can provide relief precisely where experts are scarce who can judge both the clinical content and the translation rules. Such competence cannot be created at short notice by additional budget. Settling recurring questions jointly, rather than examining them again for every pairing, can relieve those experts. At the same time every shared decision demands particular care, because one error can affect several recipients.

The group responsible can gain a broader view in the process: it brings together the requirements of the incoming and the outgoing formats. That makes it possible to settle jointly what counts as an active medication and which differences between the formats have to be observed. The condition is that the clinical perspectives of the parties involved feed into that work.

A joint decision cannot, however, anticipate every later use. A rendering that is clinically appropriate for one purpose can be inadequate for another. The intermediate layer therefore reduces repeated substantive decisions, but it does not replace clinical review on the receiving side. There it has to be judged whether the transformed data is suitable for its actual use.

Four — the risk

A transformation error can go unnoticed. The structure is intact, technical validation passes, the record displays correctly. The content is nevertheless wrong. A dose is transferred with the wrong unit, a code denotes something adjacent rather than the same thing, or a lost qualifier turns an explicitly excluded finding into a present one. Such errors can produce technically valid and seemingly plausible records. An error log alone is therefore not enough to detect them.

These possible consequences give the three preceding questions their weight. They also require the layer's effect on risk to be considered in both directions.

What the intermediate layer can improve. Fewer translations can free more time for reviewing each one. Jointly documented decisions make it traceable why a piece of content is transferred the way it is. Clearly delimited translation steps make faults easier to localise. When a shared rule is corrected, every translation that uses it can benefit.

Which risks it concentrates. An undetected error in the shared description or in a shared rule can affect several recipients. With a direct translation an error stays confined to the connection concerned, provided the same faulty rule is not used elsewhere as well. Shared decisions thereby widen the possible reach of a single error.

The intermediate layer thus offers the chance to examine recurring decisions more thoroughly. At the same time it can extend the consequences of a wrong decision. Whether a net gain in safety results depends on the quality of the reviews and on how detected errors are handled. The expertise required for that has to be funded permanently. Clear responsibilities are needed as well, along with traceable changes and procedures for informing affected recipients and implementing corrections.

No claim is made here about how often such errors occur. What is described are the mechanisms that have to be taken into account in the design. How frequently errors occur, and how the architecture affects that, remains to be investigated empirically.

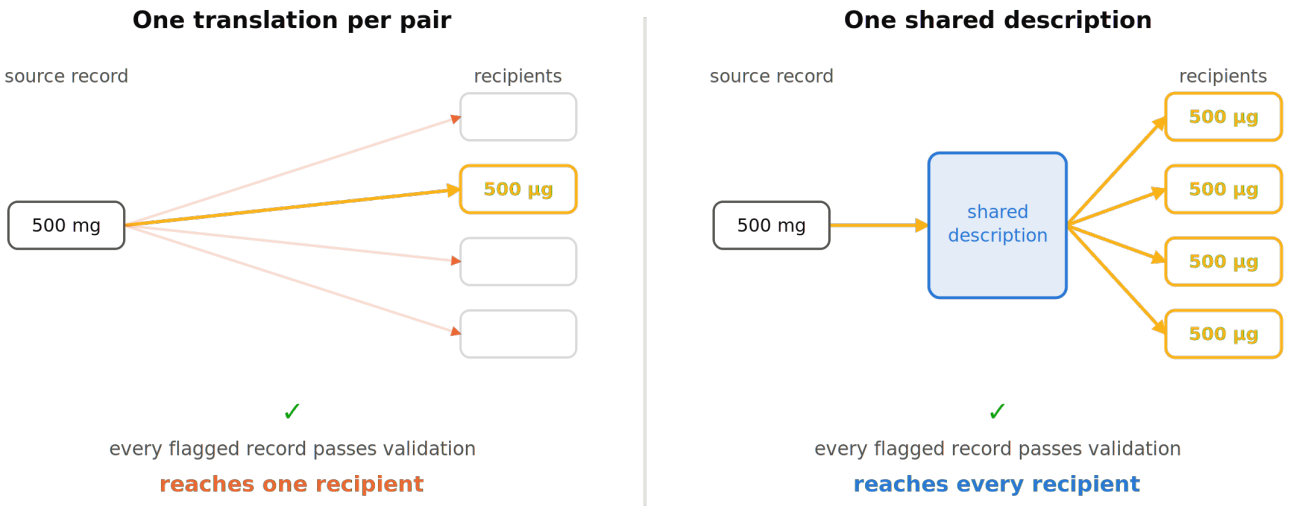

Figure 4 — A dose mis-mapped from milligram to microgram. Both arrangements produce a technically well-formed record that passes validation, so nothing reports a fault. What differs is reach: pairwise the error sits in one pairing, through the shared description it arrives everywhere at once. The figure shows the concentration, not a frequency.

The installed base, and the pace of change

The four questions suggest at first that an organisation can choose the architecture for itself. In fact it meets existing products, installed interfaces and the development plans of many independent parties. What also decides the matter, therefore, is how a shared solution can be introduced into that installed base.

Nobody starts from zero. With MyHealth@EU, further development builds on an exchange that already exists. Suppliers have developed interfaces and rolled them out to their customers. Changes therefore concern not only program code but also deployment, operation and support. Depending on the intervention, they also change how the working day runs for clinicians.

Existing investments and established routines argue at first for stability. Taken by itself that is not an argument for an intermediate layer, since introducing one also causes effort. Its benefit shows only in how it can absorb later changes.

The parties cannot move at the same pace. Suppliers differ in their resources, product cycles and commitments to their customers. A common statutory deadline does not synchronise those development plans automatically. The architecture therefore has to allow a transition in which different version states are temporarily in service.

Transformation can decouple the schedules. When one side moves to a new version, a translation can keep the exchange with an older version on the other side working. That is possible with direct translations too. A shared description, however, offers the possibility of supporting those transitions for several parties under the same rules.

This benefit goes beyond the number of translations. The cost calculation asks whether development and maintenance effort is reduced. Decoupling in time answers a different question: how does the exchange stay possible while the parties convert their systems at different moments?

For that, the versions needed have to be explicitly supported and their translations reviewed. Where specifications are revised faster than the installed systems can follow, handling differing version states becomes a standing task. Whether and to what extent this will happen with European profiles remains to be investigated.

The substantive limit of transformation still holds. Differences in representation can be bridged as far as the meaning is preserved. New clinical content, by contrast, cannot be processed by an older receiving system merely because an intermediate layer is present. Extensions to the shared description and to the receiving systems can be required for that.

Shared specification, shared library and shared service are different operating models. A specification fixes how the translation is to be carried out. A library provides a reusable implementation for it. A service performs the translation for the parties connected to it.

With separate implementations it has to be checked whether they implement the specification consistently. A shared library makes that easier, but it can be deployed in different versions. Rules are therefore needed for supported versions, for updates and for demonstrating correct behaviour. A jointly operated service makes central control of changes easier. There too, controlled version management and substantive review remain necessary.

A shared service concentrates operational responsibility. When several parties use the same processing path, an outage or a security incident can affect several connections at once. How far the consequences reach depends on the concrete architecture and its protective measures. Availability, recovery and the handling of disruptions therefore have to be part of the design from the outset.

How the health data itself is handled has to be settled as well: which data does the service process, which does it store or log, and for how long? Who may access it? Which data protection roles and obligations do the parties hold, and on what legal basis does the processing take place?

This paper does not answer those legal questions. It names them as conditions to be clarified by qualified advice against the concrete arrangement. The decision for a shared description is therefore to be distinguished from the decision about how it is operated. Where it is realised as a shared service, substantive quality, availability and protection of the data have to be planned and funded together.

Three kinds of transformation

The four questions of count, cost, skills and risk apply to three distinct tasks. These are frequently collapsed under the term mapping, but they have to be considered separately:

Trans-Forming moves data from one structure into another. It determines which content belongs in which fields and how the relations between them are preserved.

Trans-Coding assigns codes of one code system their counterparts in another. Here it has to be examined whether their meaning corresponds and which differences remain.

Trans-Lation renders human-readable text from one language into another. The clinical statement and its comprehensibility have to be preserved in the process.

Taken together these tasks are referred to here as **3T**. They interlock, but they require different competences and different review procedures. Their error risks differ as well. The benefit of an intermediate layer therefore has to be judged separately for each of the three.

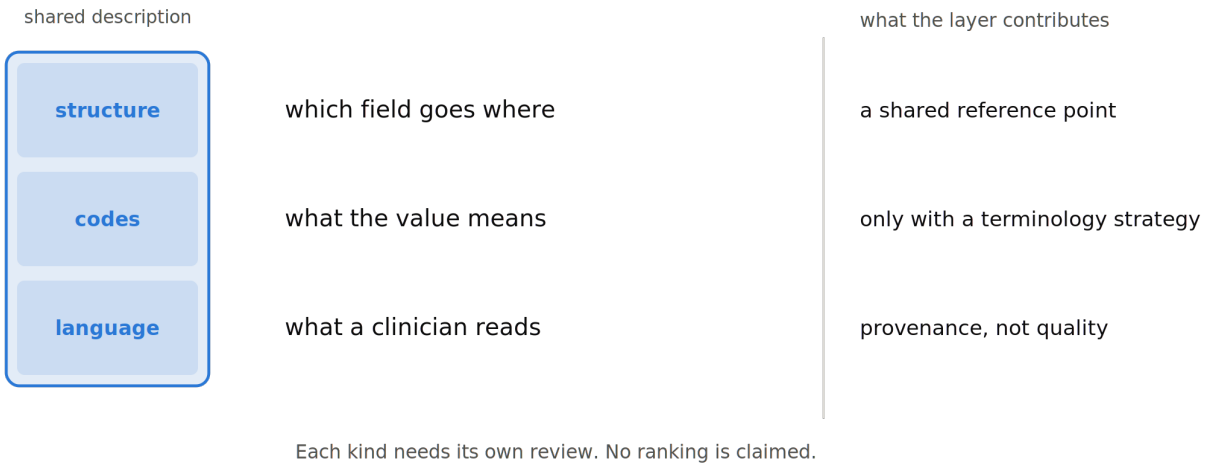

Figure 5 — The three kinds of transformation, stacked inside one shared description, and beside them what the intermediate layer contributes to each. The figure does not rank the three by risk or by value. Each demands its own review.

Together the four questions and the three kinds of transformation form a grid. It shows which tasks, competences and risks have to be taken into account in each case:

CRITERION	TRANS-FORMING	TRANS-CODING	TRANS-LATION
Count	data formats and their versions	code systems, value sets, concept maps and their respective versions	languages, translation directions and bodies of text
Cost	specification, implementation, tests and maintenance	building and reviewing concept maps, terminology infrastructure, and maintenance across version changes	translation, clinical and linguistic review, and updating changed content
Skills	data modelling and development, supplemented by clinical understanding	terminology expertise and clinical judgement	linguistic and clinical competence
Risk	content is assigned wrongly or loses its context, although the structure stays technically valid	a formally valid code conveys a different or less precise meaning	a translated text alters the clinical statement or is misunderstood

Two points follow from this grid.

Each kind of transformation needs its own review. Structural errors can show up in technical validation. Where two formally compatible fields are confused, however, such an error can go unnoticed too. In coding it has to be examined whether the chosen counterpart preserves the clinical meaning in the given context. In language

translation what matters is whether the statement, negations, uncertainties and temporal references are rendered correctly.

The three do differ in where an error becomes visible. A structural error often breaks something that a technical check reports. A coding error produces a formally valid record that no error log objects to. A language error is not in the data at all and appears only when a clinician reads it. That is not a ranking of severity but a statement about which review procedures can take hold in the first place.

A general ranking of risk cannot be derived from it. What decides is the concrete content, its use and the consequences of an error. Review effort has to follow that, not the technical difficulty of implementation alone.

The intermediate layer relieves the three tasks in different ways. For structure it offers a shared reference point through which mappings can be reused. For coding a shared structure is not enough. A terminology strategy is needed as well: which code systems and value sets are used, which mappings apply, and how are missing or only partly matching counterparts handled? A canonical code system like SNOMED-CT can be part of that strategy, but it replaces neither the substantive review nor the maintenance of the mappings.

For language translation the intermediate layer can carry supporting information along, such as the source language, the original text and the provenance of a translation. It can also make the use of agreed multilingual designations easier. The quality of translated free text and its comprehensible display nevertheless have to be assured separately. Keeping the original text accessible supports traceability, but it does not replace a reviewed translation.

The practical consequence: whoever justifies an intermediate layer by its structural benefit still has to plan and fund terminology and language as tasks in their own right. A shared data structure alone does not yet guarantee an exchange in which the clinical meaning is preserved.

How to tell whether the shared description is good enough

A first test is available early: data is taken from a source format into the shared description and then translated back into the same format. This round trip shows whether the content tested survives that path. What counts as an equivalent result has to be settled in advance. A different order or spelling can be unproblematic, an altered clinical meaning cannot.

The informative value of this test has three limits.

It depends on the test data. A successful result shows only that the content actually tested was preserved. It says nothing about elements that were not included. Purpose-built test cases therefore have to cover rare, optional and complex content as well. This is both a particular challenge and a quality criterion for synthetic data. Real exchange data supplements those tests, but it does not replace them.

It can reward mere passthrough. Where unrecognised content is placed unchanged into a general-purpose container and copied back later, the test can succeed without the shared description mapping that content substantively. Such mechanisms can be sensible for preserving information. For demonstrating model coverage, however, they have to be declared separately. Preserving a value is not the same as being able to interpret it.

It does not demonstrate correct translation into another target format. A value can be assigned wrongly in the shared description and still return unchanged, when the reverse translation inverts the same error. The round trip alone does not detect that. It therefore has to be examined in addition whether the content is mapped correctly in substance within the shared description and arrives correctly in the intended target formats.

From this a fundamental design question arises: what does the shared description have to preserve? If it is to reproduce every peculiarity of every source format in full, it risks becoming unwieldy and hard to maintain. Its scope

therefore has to follow the agreed purposes of exchange. For the content required by those purposes, preservation of meaning has to be demonstrated. Unsupported content and possible losses have to be named explicitly.

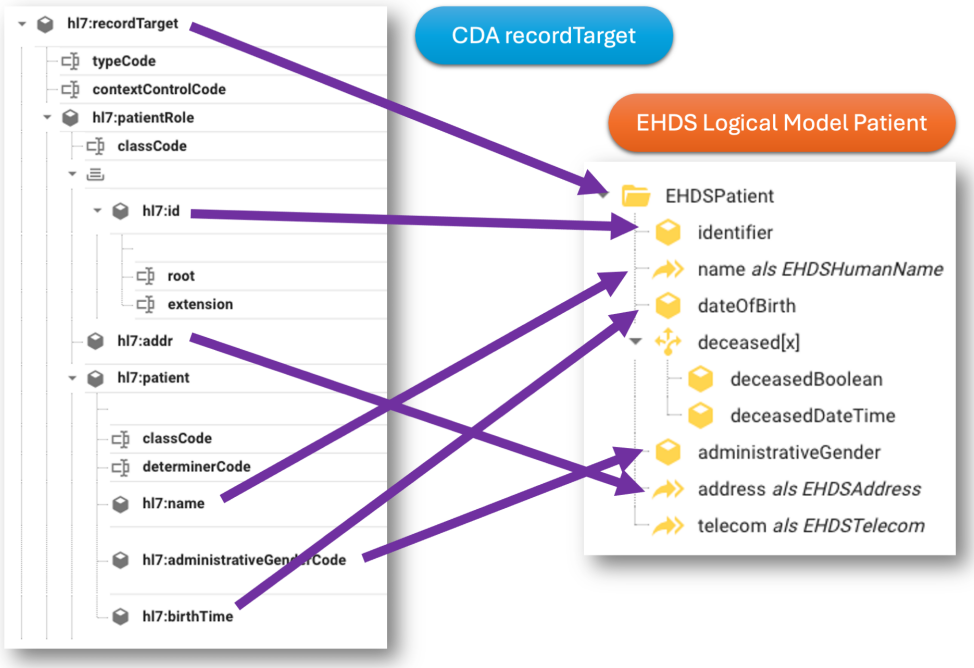

Figure 6 — One inward translation, at the level it is actually decided. The CDA `recordTarget` on the left is mapped into the EHDS logical model for the patient on the right. Some elements correspond one to one, such as `hl7:id` to `identifier`. Others are carried into a datatype of the target model, such as `hl7:addr` into `EHDSAddress`. What the figure cannot show is the decisions taken by leaving an element unmapped, and those are the ones a review has to find.

The round trip is thus one building block of the review. It helps uncover losses of information, but it replaces neither the substantive review of the model nor the review of the translations to the recipients.

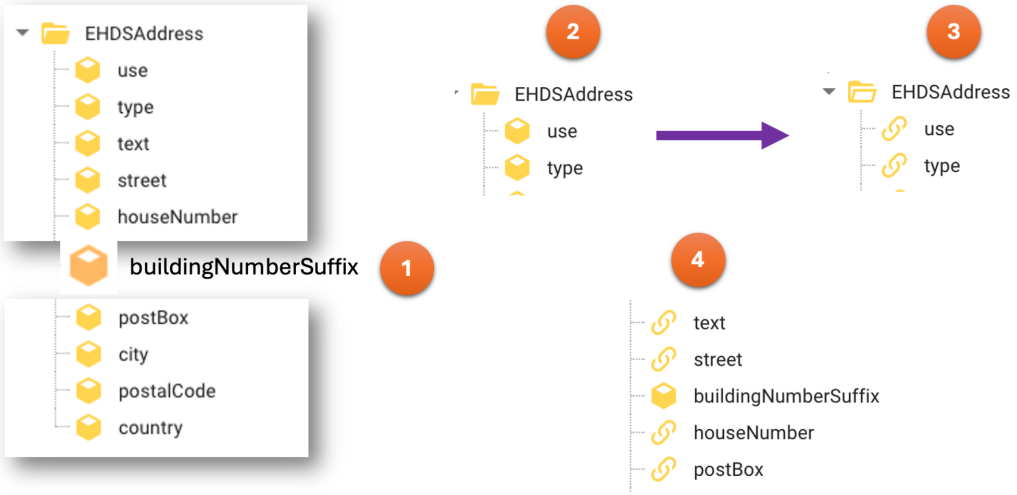

Figure 7 — Extending the shared description, and keeping the extension visible. An element absent from the EHDS model, `buildingNumberSuffix`, is added to the address datatype (1). The model before the change (2) and after it (3) differ in how each element is marked: the inherited ones now carry a link back to the EHDS model, the added one does not (4). The extension is therefore declared rather than silently merged, which is what makes a later reconciliation with a revised EHDS model possible at all.

Where the argument is weak

With few and stable formats the effort for an intermediate layer can exceed its benefit. Two incoming and two outgoing formats require four direct translations. A shared description likewise requires four connections, but in addition the development and maintenance of the model. Where the formats are rarely changed and there are no further requirements for reuse or decoupling in time, direct translations can be cheaper in the long run.

The shared model is also an organisational task. It needs a responsible body, a binding change procedure and a regulated release of new versions. Where those are absent, teams can begin to route around the model with solutions of their own. Then the costs of individual translations arise again alongside the costs of the intermediate layer.

Nor does the model resolve disagreements about meaning by itself. It offers a common place where they can be documented and worked on. A substantial benefit lies in that: differences that could otherwise stay hidden in individual translations become accessible to joint examination. The substantive work does not disappear, but it can be concentrated.

That concentration reaches its limit where a responsible body is absent or where its mandate is not sufficient for a binding decision. In cross-border exchange in particular it therefore has to be settled who may decide contested questions of meaning and how remaining differences are to be handled. Several paths are possible here.

The model can represent differing readings explicitly and mark the conditions under which they apply. That preserves the differences but requires recipients to take them into account. Alternatively a common reading can be agreed for a particular purpose of exchange. It then has to be established how divergent content is handled and whether information is lost in the process. Where a question stays open, that uncertainty has to be recognisable as well. It becomes dangerous above all when an unsettled meaning appears to be agreed.

The intermediate layer can make such differences visible and traceable, provided its documentation and review are binding. It does not, however, create a decision-making authority that is organisationally or legally absent. Before it is introduced it therefore has to be settled who works on contested questions, to whom they are escalated when no agreement is reached, and what applies until the matter is clarified. That includes the decision as to which content can be exchanged under those conditions and which has to be withheld.

Conditions under which the layer is advisable

What argues for an intermediate layer is that several formats on both sides have to be supported over a longer period. Its benefit grows where further formats are added or where different versions stay in service in parallel. A fixed number does not mark a generally valid turning point. What decides is whether the expected reuse justifies the additional effort for the shared model and its maintenance.

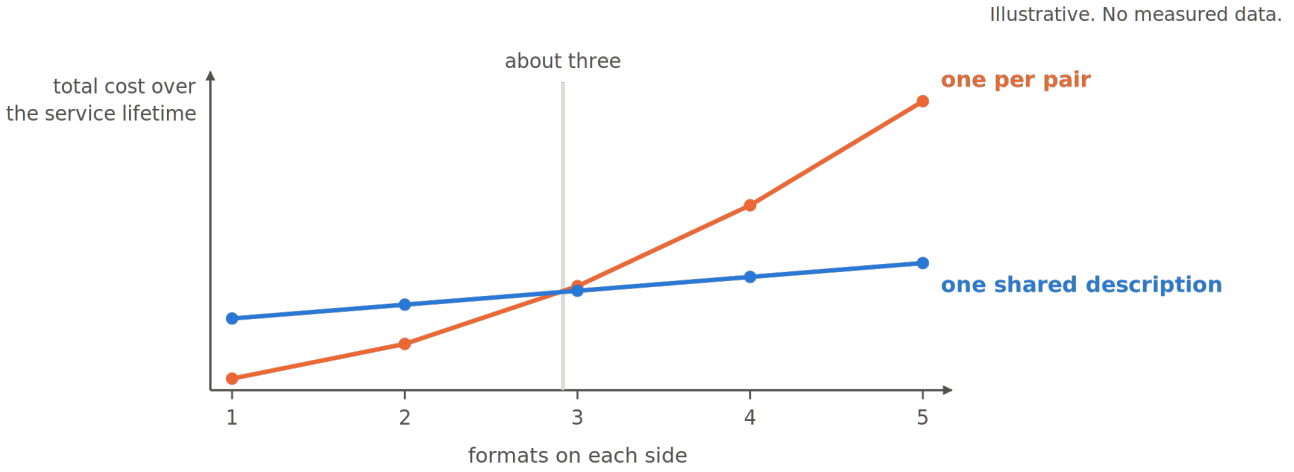

Figure 8 — Where the two arrangements cross. Pairwise translation is cheaper below roughly three formats on each side and rises steeply above it. The shared description starts higher and rises gently. The curves are illustrative: the shape is the claim, the values are not measured.

It is particularly relevant in an operation running over several years, in which parties join or leave and develop their systems at different moments. The intermediate layer can then both reduce repeated translation work and ease transitions between version states.

For its use the following conditions have to be met as well:

Responsibility and substantive review are secured. A named body is responsible for the shared model. Its mandate for settling contested questions of meaning is defined, and it has permanently funded technical and clinical review capacity. Where that is absent, the concentration of decisions cannot be adequately safeguarded. Direct translations, however, require appropriate review as well.

The funding carries across the whole operating life. It covers development, review, operation and maintenance. In doing so it takes account of the fact that costs and benefits can fall to different parties. Whoever provides and maintains translations needs a dependable basis for it, even where the benefit arises mainly elsewhere.

The operating model is chosen explicitly. A distributed implementation allows the parties to operate independently, but it requires binding rules for versions, updates and conformance testing. A shared service makes central control easier and at the same time creates a common operational dependency. With that choice, availability, security and data protection responsibilities have to be settled from the outset.

There is a procedure for unresolved questions of meaning. It is established who works on contested clinical readings, to whom they are escalated when no agreement is reached, and what applies for as long as no agreement exists. That includes rules for marking remaining differences and for handling content that cannot be translated reliably under those conditions.

For a one-off exchange or few stable connections, direct translations can be the more appropriate solution. Where no viable responsibility for the shared model exists, an essential condition of the intermediate layer is not met. The architectural decision should be examined again where additional formats, more frequent version changes or new purposes of exchange alter the starting position.

The case at hand

The criteria so far are general. For the European Health Data Space it now has to be examined what conclusion they support.

What decides it is the variety of systems to be connected. Obligatory cross-border exchange links the Member States through their national contact points. Behind those stand differing national care and IT structures. How many translations are actually needed depends on the formats deployed there and on the shared interfaces already in place. The number of Member States alone does not answer that question. What is decisive is which differing representations have to be brought together at the point of mediation under consideration.

The exchange has to deal with change permanently. MyHealth@EU is the starting point for further implementation. The EHDS Regulation provides for technical specifications for the European exchange format and for their regular updating. The architecture therefore has to support both the transition out of the installed base and later version changes. A shared description can help to reuse adjustments and to connect differing version states with each other for a time.

Common deadlines do not replace coordinated transitions. The Regulation sets binding dates and creates a framework for cooperation. It does not follow from this that all national systems and suppliers roll out their changes simultaneously. Within that framework the technical implementation has to make transitions possible without presupposing a simultaneous switch by all parties. The basis is the EHDS Regulation, and there in particular Articles 15, 23 and 92.

> **The recommendation.** For obligatory cross-border exchange under the European Health Data Space, the intermediate layer is the right shape. Its strongest argument is the combination of reuse and decoupling in time across a long-running operation.

The recommendation rests on the three findings above, and each of them can be checked on its own. At the points of mediation, more differing representations have to be brought together than pairwise translations can carry over the operating life. The specifications are in motion and will stay so. The Regulation sets binding deadlines without synchronising the transition of the systems involved. It does not rest on the cost calculation, which is the weaker of the arguments and the one most easily disputed. How large the economic advantage turns out to be has to be determined against the formats, versions and existing translations actually to be connected.

The recommendation does not by itself justify either an additional layer of mediation at every point or a single central service for the whole of European exchange. And it does not hold unconditionally: clinical review has to be funded permanently, a responsible body with a mandate has to be named, and the operating model has to be chosen explicitly.

What the conclusion leaves open. A recommendation for an intermediate layer does not mean its conditions are already met. Clinical review has to be funded permanently. That holds for direct translations too. With shared rules, however, the requirement carries particular weight, because a single error can affect several recipients.

A responsible body with a clear mandate for settling contested questions of meaning is needed as well. For cases where no agreement is reached, it has to be established how matters proceed. The choice between a distributed implementation and a shared service also has to be made explicitly. With a service, availability, security and data protection responsibilities belong to the decision from the outset. Whoever introduces the intermediate layer and leaves those questions open has chosen an architecture but has not yet secured that it will hold.

When the recommendation is to be reassessed. Were the formats involved and their versions to stay stable over a long period, or were the number of translations actually needed to fall significantly, the additional effort for the shared model could exceed its benefit. Direct translations would then have to be examined again as an alternative.

Were the necessary clinical review capacity not to be secured, on the other hand, what would first be in question is responsible implementation as such. No automatic advantage for direct translations follows from that, since they have to be reviewed substantively as well.

The recommendation is therefore to be measured regularly against the actual development: against the variety of formats, against the effort of changing them and against the review capacity available. Where those foundations change materially, the architectural decision has to be examined again.

What is deliberately not claimed here

An intermediate layer does not automatically improve data quality. It cannot correct faulty source data by itself. The quality of the translation depends on the substantive rules, on their implementation and on their review. Errors in shared rules can affect several recipients.

It does not guarantee safe cross-border exchange. It changes where risks arise, how far errors can reach and who is in a position to detect and correct them.

It is not necessarily cheaper from the outset. Development, coordination and introduction of the shared model cause additional effort. Whether and when that pays off depends on reuse and on later maintenance.

A successful round trip does not demonstrate the correctness of the model. It shows at first only that the content tested was preserved on the way out and back. Passed-through content and mutually cancelling errors can make it succeed as well. Whether the content was mapped correctly in substance, and arrives correctly in other target formats, has to be examined separately.

The legal questions raised by a shared service are not answered here. The paper names the need for clarification. Legal assessment requires a concrete description of the operating model and corresponding professional advice.

The recommendation for the European Health Data Space is a reasoned judgement. It follows from the criteria named and the facts assumed. Both remain open to examination. Where material conditions change, the recommendation is to be reassessed as well.

The shared description is not neutral. It contains decisions about which content is distinguished, merged or not mapped at all. Parties whose own models sit closer to those decisions can face less adaptation effort than others. Such differences and their consequences have to be disclosed and justified during development.