German FHIR Profile Governance

Document 5 · Dr Kai U. Heitmann · Release 1.54.1 – 7 September 2026

A profile-corpus analysis with its method disclosed: the artefacts carried by German profile families, their explicitly declared terminology constraints, and what the counts can and cannot be asked to carry.

Written for: FHIR profiling and terminology teams, national programme staff, and anyone who has to judge how far a figure from this corpus can be taken.

This report stands on its own and is cited from Document 3. Readers seeking the architecture argument rather than the profile analysis should begin with Document 2. Section references point to Document 3, and other appendix references point to Document 4.

The report originated as Appendix D-10 and retains that numbering so that cross-references remain stable.

The closing section states the corpus, the counting rules and the limits in full. It does so to make the method legible, not to invite re-execution: the source packages are not published with this report, and reproducing the figures byte for byte from it alone is not possible. What a reader can do is see exactly what was counted and what was deliberately not, and weigh every figure accordingly.

Contents

D-10: The German Electronic Patient Record and the Profile-Consistency Question

The German electronic patient record (ePA) and the surrounding profile families provide a detailed public case through which to examine whether a common persistence or exchange technology is sufficient for interoperability. The case is included because claims made about the ePA can be tested against normative specifications, published implementation guides, governance documents and an identified corpus of German FHIR artefacts. Statements by the responsible bodies are treated as primary evidence of their requirements and positions, not as independent evidence of implementation quality or effectiveness.

Evidence is graded by proposition rather than by publisher, following Document 4, Appendix G. **Evidence grade A** applies to constitutive requirements and directly repeatable readings of identified artefacts: statutes and ordinances, commercial-register records, normative specifications and reproducible counts where the material is fixed. **Evidence grade B** applies here to published working-group positions and official assessments of implementation quality or governance: inspectable primary documents, but self-reporting by involved parties. **Evidence grade C** applies to usage figures reported through the press and to industry-association statements. A normative gematik specification therefore carries grade A for the interface behaviour it requires and no grade for the proposition that the behaviour is achieved in practice. Author-derived totals and counts whose complete input corpus cannot be reconstructed publicly are identified as such rather than inheriting the authority of their publishers.

Results and limits at a glance

- The **corpus** counted on 12 August 2026 contains 2,996 distinct StructureDefinitions, of which 2,317 are resource profiles. Its differentials declare 2,399 bindings, 79.4% of them `required`.
- The same differentials contain 8,073 listed fixed or pattern constraints. These have not been classified by element path or datatype and cannot all be treated as terminology constraints.
- Reuse inside the corpus is local: 721 of 1,110 in-corpus derivations stay inside their package, and where a derivation does cross a package boundary it names an explicit version in half the cases. Version pinning is therefore concentrated at governance boundaries (D-10.7).
- The medication and allergy subset has a lower `required` share among its explicitly declared bindings. This does not establish lower overall constraint strength because inherited bindings and element-level appropriateness were not assessed.
- The method and the source inventory are disclosed in full. Independent reproduction from this document alone is not possible, because the exact repository revisions and a frozen input corpus were not preserved publicly.

D-10.1 Two Storage Paths, and What the FHIR Services Actually Do

The ePA holds data on two paths with different semantics, and the difference matters for every claim made about it.

The document path is IHE XDS.b (transactions ITI-41, ITI-18, ITI-43, ITI-62, ITI-92). In the document-management specification of the current record generation the permitted media types include PDF/A-1 and PDF/A-2, plain text, XML and JSON, `application/fhir+xml`, `application/fhir+json` and HL7 v3 XML. That list is release-specific and is not to be read back onto the earlier generations.

FHIR content can therefore be stored here, but as a registered document payload governed by XDS metadata, alongside the PDFs and the KBV record types (MIOs). The media type settles nothing about what the object is: it

separates neither a single serialised resource from a FHIR document bundle, nor either of them from XML or JSON that is merely syntactically well formed.

The ePA MHD Service (`de.gematik.epa.mhd`, 2 resource profiles) is a read-only FHIR façade over that registry, read off the conformance statement rather than inferred from the initials: the document responder declares mode `server` and, for `DocumentReference`, the interactions `read` and `search-type` and no others, with no `create`, `update` or `delete` anywhere in the statement [1]. **Evidence grade A.**

MHD internationally carries both consumer and provider transactions, so the profile family does not by itself imply read-only. The ePA-specific statement does. What the façade exposes is registry metadata and document retrieval, the search result modelled as a bundle of `DocumentReference`. The clinical content of a registered document does not become an independently queryable resource.

The data-service path was added with ePA 3.0 in 2025. Its FHIR interfaces work on Medication, Patient and AuditEvent resources, and these are not three parallel clinical repositories: the medication resources carry the structured clinical state, the patient resources identify and contextualise the record, and `AuditEvent` records access rather than care. Only the first is a longitudinal clinical dataset.

Here the server holds a canonical state of its own. The specification requires it to validate every incoming resource against its own current profile version, to discard the logical `Resource.id` supplied by the client, which leaves business identifiers in `identifier` untouched, to ignore the client's `meta.profile` and stamp its own instead, and to execute migration rules in ascending version order when the implementation guide is upgraded. The documented example is the 1.3.0 upgrade, in which `MedicationStatement` instances are generated retroactively from stored `MedicationRequest` records [2].

A claim in circulation, tested against the sources. It is asserted in the German debate that the ePA persists FHIR in its incoming format, and that primary systems must consequently support every profile version ever submitted, for decades. For the FHIR data services the specifications state the opposite: the arriving payload is not what is kept. The claim holds only in the weak sense that the document path retains what was submitted, and what is retained there is a document, not a resource. The observation behind the claim is nevertheless sound. D-10.2 examines where the durable version burden actually sits.

The profile volume of the ePA itself is modest, measured in resource profiles. The current release line, ePA 3.1.3, comprises three implementation guides. Read from the artefact index each of them publishes, `de.gematik.epa` 1.3.2 carries 3 resource profiles, `de.gematik.epa.medication` 1.3.5 carries 10 and `de.gematik.epa.mhd` 1.1.3 carries 2, **15 in all** [3]. **Evidence grade A** for the published indexes; the summation across the three is the author's own, gematik publishing no aggregate figure.

Resource profile is meant in the sense of the table note of D-10.2, so the operation parameter profiles of the medication service, the larger group of the two there, are excluded, along with extensions, datatype profiles,

[1] CapabilityStatement of the ePA document responder, `de.gematik.epa.mhd` 1.1.3, the current version of the release line. `gemspec.gematik.de/ig/fhir/epa-mhd/1.1.3/CapabilityStatement-epa-mhd-document-responder.html`, retrieved 2 August 2026.

[2] gemSpec_Aktensystem_ePAfueralle V1.8.2, 14 July 2026, together with the general principles of `de.gematik.epa` 1.3.0 and the migration chapter of `de.gematik.epa.medication` 1.3.2. `gemspec.gematik.de/ig/fhir/epa/1.3.0/general-principles.html` and `…/epa-medication/1.3.2/migration.html`.

[3] Published artefact indexes of the three guides of the current release line, all retrieved 2 August 2026: `gemspec.gematik.de/ig/fhir/epa/1.3.2/artifacts.html`, `gemspec.gematik.de/ig/fhir/epa-medication/1.3.5/artifacts.html` and `gemspec.gematik.de/ig/fhir/epa-mhd/1.1.3/artifacts.html`.

terminology artefacts and conformance resources. The audit record for this figure, and the respect in which its method differs from the counts taken from packages, are in D-10.6.

Figures of "dozens to hundreds of ePA profiles", which occur in the debate, are not supported for these packages. What the figure measures is the profile surface, not the size or the implementation cost of the ePA specification as a whole, which further comprises operations, value sets, code systems, the confidentiality and cryptography requirements and the entire document path.

Transition rules are asymmetric across the German stack. For the ePA FHIR services no general transition period or parallel-version rule was found in the implementation guide for primary systems (gemILF_PS_ePA V3.8.1, 24 March 2026), and the version burden sits with the server. The e-prescription carries explicit and dated rules: a transition window from 1 July 2026 to 14 January 2027 during which versions 1.3.2 and 1.4.x are processed in parallel, selected by `authoredOn`, after which only 1.4.x is accepted. Those rules are set in separate contracts between the framework partners rather than in the specification itself, which is why they do not generalise across the stack.

D-10.2 The Profile Families a German Primary System Must Carry

The word "corpus", once, because everything counted here rests on it. A *corpus* in this document is the identified set of artefacts that a count was run over: every FHIR `StructureDefinition` found in a fixed collection of German source repositories and published package archives, deduplicated by canonical URL. It is not "all German FHIR profiles" and does not claim to be.

It is a set assembled at a stated moment, which can be listed and reassembled. That is what makes a count over it checkable rather than assertable. What the set contains, how it was assembled and what falls outside it are set out under *The corpus* at the end of this report, together with the counting rules that operate on it.

The ePA is one family among many. No single product carries all of them, and which subset applies follows from sector, product scope and statutory role: the hospital interfaces bind hospital systems, the KBV forms bind ambulatory contract care, the notification profiles bind wherever a notifiable finding arises, and the research core dataset binds the university hospitals taking part in it.

One family sits outside statutory health insurance altogether. The Deutsche Gesetzliche Unfallversicherung publishes FHIR profiles under SGB VII, the statutory accident insurance, not under SGB V. It is nonetheless a profiling body of the German health system, its artefacts share the base layer and the terminologies, and a supplier serving hospitals or occupational medicine meets them like any other family. Excluding it because the legal basis differs would measure the statute rather than the profiling burden.

The table below therefore sets out the burden of the market rather than of any one installation. Nine families are shown, each governed independently of the others. A vendor with a broad product range meets further ones beyond them, among others the emergency and nursing datasets and the cancer-registry specifications.

FAMILY	PUBLISHER / STATUTORY BASIS	SCALE (VERIFIED WHERE STATED)	VERSIONING
Deutsche Basisprofile (`de.basisprofil.r4`)	HL7 Deutschland e.V.; **no statutory mandate**	the national base layer; **76 StructureDefinitions at 1.6.0** (counted), of which 19 resource profiles, 29 extensions and 28 datatype profiles	semantic versioning; current stable R4 release 1.6.0 of 29 June 2026
ISiK — hospital interfaces	gematik, § 373 SGB V	Stage 6 (1 July 2026): **177 StructureDefinitions in one package** (counted), of which 156 resource	annual stages, each published around 1 July; the cadence is an observed

FAMILY	PUBLISHER / STATUTORY BASIS	SCALE (VERIFIED WHERE STATED)	VERSIONING
		profiles and 2 Parameters profiles; Stage 5: 127, of which 109 resource profiles	publication pattern, not a governance rule stated as such
ePA	gematik	**15 resource profiles across the three guides of the current release line** (read from the published artefact indexes, 2 August 2026: 3 at `de.gematik.epa` 1.3.2, 10 at `de.gematik.epa.medication` 1.3.5, 2 at `de.gematik.epa.mhd` 1.1.3); the archived material counted for the corpus is older and larger, 14 at `de.gematik.epa` 1.1.0 and 28 at `de.gematik.epa.medication` 1.1.1	tied to ePA release trains
MIOs (`kbv.mio.*`)	KBV, § 355 SGB V; operationally mio42 GmbH	**13 packages counted, from 7 to 214 StructureDefinitions each, 717 in all**; the vaccination record, the medication plan, the digital-health-application record and the teleconsultation record joined the count on 12 August 2026	each record type versioned independently
KBV forms, prescribing and base layer (`kbv.ita.*`, `kbv.itv.*`, `kbv.basis`)	KBV	20 `kbv.ita.*` packages in the registry; **five counted with 270 StructureDefinitions**, plus two `kbv.itv.*` packages with 18 and the KBV base layer `kbv.basis` 1.9.0 with 48	technical annex at v1.74, 16 June 2026
Appointment services (`kvdigital.*`, `de.kvtelematik.*`)	KV Digital GmbH, a subsidiary of the KBV	**five packages counted, 61 StructureDefinitions** after deduplication, of which 44 resource profiles	per package, independently
MII core dataset	Medical Informatics Initiative	24 packages in the registry against 19 modules named on the initiative's own site; **21 modules counted, 557 StructureDefinitions**, of which 442 resource profiles	calendar versioning since 2024
DEMIS — notifiable-disease reporting	Robert Koch Institute, § 14 IfSG	≈ 9 packages, plus five for the genomic-surveillance extension; **three counted, 509 StructureDefinitions**, of which 489 resource profiles	per package, independently

FAMILY	PUBLISHER / STATUTORY BASIS	SCALE (VERIFIED WHERE STATED)	VERSIONING
DGUV — statutory accident insurance	Deutsche Gesetzliche Unfallversicherung e.V., **SGB VII rather than SGB V**	**three packages counted, 88 StructureDefinitions**, of which 70 resource profiles, 17 extensions and 1 datatype profile: `dguv.basis` 1.5.0 with 23, `dguv.enla` 1.0.0 with 45, `dguv.oper` 0.1.0 with 20	per package; `dguv.oper` is a comment release and pre-1.0

Counting convention for the figures marked as counted. They were established on 12 August 2026 against 66 German source repositories and 46 published package archives supplied to the author, in a single deduplication pass, under the counting rules set out at the end of this report. Where a family is given a resource-profile figure, that column is meant. A family total is deduplicated within the family and can therefore fall below the sum of the packages listed for it. For repositories without a published archive the count is of build output rather than of a release.

The method is now checked at ten points rather than one. Where both a published archive and a source repository were held for the same package, both were counted and compared. At Stage 6 of the hospital interfaces, and at nine KBV packages — the base layer, and the record types for the child examination booklet, the maternity record, the transfer note, the hospital discharge letter, the laboratory report, the patient summary, the dental bonus booklet and the dialysis dataset — release and build output yielded **exactly the same canonical URLs, with no difference in either direction**. This does not prove that build output equals a release in general, and it says nothing about the repositories for which no archive exists. It removes the doubt at the ten places where it could be tested, and those ten carry 787 artefacts.

The figures are not commensurable across families, and no ranking should be read out of them. A publisher who expresses a controlled vocabulary as one profile with a bound value set produces a handful of artefacts where a publisher who instantiates it per concept produces hundreds, and both are correct FHIR. Two of the families in this table are of the second kind: the notification packages carry one profile triple per notifiable pathogen, and the child examination booklet carries one observation profile set per examination stage.

Across the whole counted corpus, 2,317 resource profiles resolve to 273 distinct base definitions, and 40% of them profile a single resource type, Observation. A StructureDefinition count therefore records editorial convention as much as clinical scope. This bears on Finding 5 rather than against it: the artefact surface a primary system must carry is governed by how the specifying bodies choose to encode vocabulary, which is a question of profile discipline and not of the choice of persistence standard.

Qualifications on the largest counted family. Of the medical information objects, thirteen packages are now held and counted, the vaccination record among them since 12 August 2026. Seven of the thirteen counted stand at draft, comment, update or Festlegung status rather than at a normative release, so their inventories may still move. The spread is wide: the smallest package holds 7 StructureDefinitions and the largest, the child examination booklet, holds 214, of which 210 are resource profiles. The vaccination record itself is small: 22 StructureDefinitions, of which 15 resource profiles, against the different order of magnitude the family's largest packages suggest.

The hospital interfaces are counted as one package, not as modules. The Stage 6 package holds 177 StructureDefinitions in all, of which 156 are resource profiles, and it is monolithic: the module structure belongs to the separately versioned module packages, which cannot be added to it. The v3 generation occupies a different canonical namespace altogether, and it holds 42 StructureDefinitions of its own; the v4 medication module re-cuts content the stage package already carries, sharing 13 of its 15 canonical URLs with it. For this family the published release and the build output could both be read, and they yielded exactly the same 177 canonical URLs; nine further packages have since been tested the same way, with the same result.

Three properties of this landscape bear on the argument of this document.

First, the shared base layer is not shared in one version. To *pin* a dependency is to name an exact version of it in a package manifest or an artefact reference, rather than leaving it to resolve to whatever the registry currently serves. Reading the dependency declarations of every package manifest in the counted corpus, `de.basisprofil.r4` is pinned at **eight distinct versions by 39 packages**, with a floating `1.5.x` range declared by five further modules of the Medical Informatics Initiative.

Four of the pins can be attributed concretely: 1.4.0 by the hospital-interface basic and vital-signs modules of the v3 generation, 1.5.2 by the e-prescription patient invoice and by the published medication module 4.0.3, 1.5.4 by the record research package, and 1.6.0 by Stage 6 of the hospital interfaces. Two of the eight pins are pre-1.0 releases, held by two medical information objects. The pattern recurs one layer down: `kbv.basis` is pinned at eight distinct versions across the packages that depend on it.

Two publication surfaces of the German base profiles diverge. The implementation-guide page of the German base profiles announces thirteen R4 releases and five for STU3, while the package registry holds 68 versions of `de.basisprofil.r4`; the registry additionally carries patch and pre-release builds that the guide page does not announce. The difference is not mere bookkeeping. Version 1.5.2, which packages of the counted corpus pin, is present in the registry and absent from the guide page [4]. A version count is in any case a measure of release history rather than of scale, which is why the scale column of the table above states artefact counts and not version counts.

A pin attribution moreover holds for a package and not for the programme that publishes it. Within the e-prescription family, `de.gematik.erezept-workflow.r4` 1.6.4 pins 1.5.4 while `de.gematik.erezept-patientenrechnung.r4` 1.1.0 pins 1.5.2. Within the record family, `de.gematik.epa` 1.1.0 and `de.gematik.epa.medication` 1.1.1 pin 1.5.0 while `de.gematik.epa.research` 1.1.2 pins 1.5.4. The divergence therefore runs inside single programmes of a single publisher and not only between publishers.

What this costs is largely not a difference of inventory. The three most recent versions hold near-identical artefact sets: 1.5.4 and 1.6.0 carry the same 76 canonical URLs and differ in four definitions, and 1.5.2 and 1.5.4 share 75, of which nine differ, all nine being vital-sign profiles. Packages pinned at different recent versions therefore see largely the same artefacts under different constraints, among them the address and person-name datatype profiles, which are among the most widely reused definitions in the landscape.

Older pins are a different matter, the inventory having grown from 59 artefacts at 1.0.0 to 76 at 1.6.0, and the growth is not strictly monotonic. The profile for the ten-digit private-insurance identifier is present at 1.4.0, absent at 1.5.2 and present again at 1.5.4, so a later pin does not by itself guarantee that an artefact is resolvable. Whether any package in the corpus actually references it was not examined and is not asserted. **Evidence grade A** for the pin distribution and the artefact sets, both being direct readings of the manifests and the packages.

The KBV base profiles are moreover a parallel development rather than a derivation; the KBV states only that the preparatory work of HL7 Deutschland was "soweit wie möglich beachtet und … integriert", so a primary system carries two base layers at once.

Second, the hospital specification declares its independence explicitly. The ISiK base-module implementation guide devotes a chapter to compatibility with the other German specifications. It records that gematik creates technically independent profiles and guarantees only that instances valid against a gematik profile are also valid against the profiles underlying it, with noted exceptions, and that incompatibilities arising from older base-profile versions inside the KBV profiles are not separately listed. Three reasons are given: technical flexibility for

[4] Implementation-guide page of the German base profiles, `ig.fhir.de/basisprofile-de/`, retrieved 2 August 2026.

connectathon fixes, gematik tooling requirements, and must-support flags that mark the elements relevant to confirmation.

Third, the durable version burden sits with the KBV record types, not with the ePA services. The KBV requirement is that a primary system or patient-facing application must in principle be able to read *all* versions of a record type, and that versions withdrawn after 1 January 2022 must still be displayed. There is no conformance certification with a test number for these; validation packages are published on GitHub for self-service. This is the mechanism the claim examined in D-10.1 was reaching for, and it is more consequential than the version behaviour of the ePA data services, because it has no end date.

One clinical concept is profiled at least five times over. Medication is separately profiled in the ePA medication service, the KBV e-prescription packages, the KBV medication-plan record type, the ISiK medication and medication-safety modules, and the MII core-dataset medication module.

A convergence has begun at exactly this point, and it belongs here for symmetry. The implementation guide for medication-related use cases (`de.fhir.medication`) is issued by HL7 Deutschland together with gematik, and four independently governed families declare a dependency on it: the hospital interfaces at Stage 5 and again at Stage 6, the KBV base package 1.9.0, the KBV e-prescription package 1.4.0, and the medication module of the research core dataset. That is the clearest case of shared authorship in the corpus. It is not called the first case of its kind, because a priority claim of that sort cannot be established from the material held.

The position of the record medication service is partly resolved. At the version held for the corpus count, `de.gematik.epa.medication` 1.1.1 declares no dependency on the shared guide, while the dependency listing published with the current release line, 1.3.5, includes `de.fhir.medication` 1.0.7 [5]. That listing is a resolved dependency tree rather than a direct manifest, carrying two versions of the HL7 terminology package and three of the gematik terminology package. Whether the record medication service declares the shared guide itself or reaches it through another package is therefore not established here, and it is not counted as a fifth declaring family.

The convergence is moreover not yet convergence on one version. The dependents pin `de.fhir.medication` at 1.0.2, 1.0.4 and 1.0.7 and at a floating `1.0.x` range, so the one shared artefact is already carried at three fixed versions and one moving one. The shared layer reproduces, on a smaller scale, the pattern the base layer shows. **Evidence grade A** for the dependency declarations, which are direct readings of the manifests.

D-10.3 The Governance Record, in the Words of the Responsible Bodies

The most instructive material is not external criticism but the record produced by the German interoperability governance itself.

The statutory structure creates the institutional conditions for fragmentation without implying fault by any party. Three determination powers sit in three statutes, and another large specification family arises outside them. § 355 SGB V assigns the record-type specifications to the Kassenärztliche Bundesvereinigung. § 373 SGB V assigns the hospital interfaces to gematik. § 14 IfSG assigns the notification profiles to the Robert Koch Institute. The core dataset of the Medical Informatics Initiative arises from research funding rather than from a statutory determination power, yet introduces another release cycle. This allocation is a plausible institutional explanation for the observed programme boundaries. The derivation graph alone does not prove causality. § 385 SGB V provides a route for cross-cutting standards recommendations and possible binding rules by ordinance; § 386 SGB V establishes a related right to interoperable data.

[5] Implementation guide of the ePA medication service, `de.gematik.epa.medication` 1.3.5, `gemspec.gematik.de/ig/fhir/epa-medication/1.3.5/`, retrieved 2 August 2026.

A further instrument is in the legislative process, and it is carried here as a draft and not as law. The Act on Data and Digital Innovation in Health Care (*Gesetz für Daten und digitale Innovation im Gesundheitswesen*, GeDIG) was approved by the Federal Cabinet on 15 July 2026, having gone to consultation as a departmental draft in the spring of that year [6]. **Evidence grade A** for the fact and the date of the cabinet decision.

Source basis. The cabinet draft was supplied to the author and read in full on 4 August 2026: 265 pages, dated 14 July 2026, comprising the amending articles, explanatory memorandum and annexes. What follows is read directly from that draft. **Evidence grade A** applies to propositions about its wording. The text remains a draft, and no inference is made that it will be enacted in this form. No finding of this report depends on its enactment.

What the draft says. Amending instruction 97 inserts §§ 386a to 386c SGB V after § 386. § 386a, headed *Interoperabilitätspflicht*, obliges manufacturers within the meaning of § 384 sentence 2 number 3, on the request of those providing care who use their system, to make available without delay and free of charge the personal health data of that provider's patients together with the bindingly determined interfaces, and to enable access to patient data, in each case in the interoperable format. The format follows from the ordinance under § 385, a manufacturer who fails is liable to the provider in damages, and the associations of statutory health insurance physicians are to support providers in pursuing the claim.

§ 386c prohibits circumvention. Among the prohibited acts it names the release of data in a format that prevents complete semantic reconstruction by another admitted information-technology system, and the withholding of data elements that are not bindingly determined under § 385 paragraph 2 sentence 1 number 1 but are necessary to secure the documentation and archiving duties of those providing care. The penalty provision in § 397 attaches to § 386 paragraph 2 sentence 1 and to § 386a paragraph 1 sentence 1, in both cases to a failure to hand over or to make available, and to nothing else in this group.

The summary sheet and the operative text diverge. The wording *im interoperablen Format vorhalten* occurs twice in the 265 pages, once in the summary sheet under "B. Lösung" and once, verbatim, in the general part of the memorandum. In both places the sentence says that those providing care must in future hold their patients' health data in an interoperable format, and that manufacturers must correspondingly make that interoperable data holding possible in the systems in use.

No operative provision carries that duty. § 386 SGB V, its natural place, is not amended: the amending instructions run from § 385 at number 96, to the insertion after § 386 at number 97, to § 387 at number 98. The special part of the memorandum draws the connection in the opposite direction from the summary, stating that exchange without delay presupposes that the data are present in an interoperable format and that manufacturers are *therefore* placed under a duty to create the preconditions.

The most economical explanation is that the summary was written against the departmental draft and not adjusted when the provider-side paragraph was dropped, and that is an inference offered as one. **Evidence grade A** for the presence and the placement of the wording, that being a proposition about the document; no grade for the explanation. The distinction matters in a German procurement, because a summary sheet met first conveys a provider-side holding duty that the operative text does not impose.

The consultation statements retain their value. Three statements of May 2026 cite the same section number, and the question they raise is the one the draft has now answered.

The vendor association bvitg asks that the rules on §§ 386 and 386a permit the exchange of interoperable data without a statutory norm being imposed on internal data-holding models, and asks for consolidated, testable specifications in good time. The hospital association DKG reads the provision as an obligation bearing on

[6] Federal Ministry of Health, press release of 15 July 2026, retrieved 2 August 2026.

interoperable data holding, welcomes it in principle, and asks that it apply only where the necessary formats, interfaces and technical preconditions are bindingly determined, implemented and tried in care, so that those providing care do not answer for incomplete standards. The ZVEI describes the draft as introducing binding interoperability obligations under § 386a and asks for named international standards and a sector roadmap.

On the wording, the three are now superseded by the draft text: the DKG reading of a duty bearing on data holding matches the summary sheet, the bvitg reading of an exchange duty that does not norm internal data-holding models matches the operative text, and the ZVEI request for named standards was not met. **Evidence grade B** for the positions themselves, which are the published statements of the bodies that hold them.

The statutory advisory body has since asked for the duty the operative text does not impose, and it divided the term the debate had until now been using undivided. The Interop Council, the body constituted under § 385 SGB V to advise the competence centre, published a position on the cabinet draft on 11 August 2026 [7]. It sits inside the chain that produces the binding instrument, the associations having spoken to it from outside in a consultation, and that is a statement about position rather than about persuasiveness. It welcomes the manufacturer obligation and then asks for something the operative text does not contain.

What it asks for. "Gesetzlich vorgegeben werden sollte jedoch eine standardisierte, semantisch-inhaltliche Datenhaltung auf logischer Ebene, sodass Daten standardisiert semantisch annotiert vorgehalten werden." That is a statutory duty reaching the way data are held, not only the way they are released. The Council pairs it with a limit in the same paragraph: "Dabei sollen keine gesetzlichen Vorgaben für Datenbanken oder Systemarchitekturen gemacht werden. Die konkrete technische Implementierung bleibt in der Verantwortung der Systemhersteller."

The line it draws is not the line the debate has been drawing. The consultation split on whether an obligation should reach the exchanged data or the model in which they are held. The Council splits the second of those in two: the **logical semantic layer** is to be regulated, the **technical architecture** is not. On that reading a statute would prescribe that a blood pressure is annotated with a published concept and a defined unit, and would prescribe nothing about the store it sits in. **Evidence grade A** for the wording, that being a proposition about the document; **evidence grade B** for the position, which is the published view of the body that holds it.

The demand is narrower than its opening sentence suggests, on two counts stated in the paper itself. The binding does not reach the logical layer at large but a named part of it. Manufacturers are to have "keine semantischen Freiheiten hinsichtlich der Bedeutung oder Repräsentation von verbindlich festgelegten Dateninhalten besitzen, beispielsweise bei Maßeinheiten oder Beobachtungen". **Where nothing has been bindingly determined, nothing is bound.** And the determined part is to begin small and grow: the Council proposes the European Patient Summary, or a German specification derived from it, as the basis for "eines ersten, verpflichtenden Datenkerns", records that "für die praktische Umsetzung wird ein stufenweiser Ausbau befürwortet", and states, because a patient summary alone does not carry every care situation, that "deshalb muss dieser Kern schrittweise erweitert werden". What is asked for is therefore that a named and initially small set of data contents be fixed in meaning, with everything else inside the system left alone.

Whether this contradicts the vendor association cannot be settled from what is held here, and the appearance that it does is misleading. The bvitg statement asks that internal data-holding models not be normed by statute, and treats them as one thing. The Council divides that one thing and disclaims half of it expressly. The two positions are therefore incompatible only under the reading in which bvitg meant the logical semantic layer as well, and compatible under the reading in which it meant the store and the architecture. **The bvitg statement was read in its own wording** (56 pages, dated 17 May 2026; the sentence stands in the executive summary and again under § 386

[7] Interop Council, *Positionierung des Interop Council zu den Interoperabilitätsvorgaben im Rahmen des GeDIG*, Berlin, 11 August 2026, two pages, published on the national interoperability platform. `ina.gematik.de/community-hub/community-formate/interop-council/positionierung-zu-gedig`, retrieved 18 August 2026. Seven full members are named, chaired by Sylvia Thun.

SGB V), and it uses the term without defining it, so the scope of its term cannot be settled from the source and no reading of it is asserted. What can be said without a reading is the smaller and firmer thing: the debate was conducted with an undivided term, and the advisory body is the first participant to divide it. The remaining dispute is whatever survives inside the narrow band of bindingly determined data contents, and it is smaller than the two opening sentences make it look.

Three further points bear on this report. The Council asks that FHIR be "als semantischer und syntaktischer Standard das führende Austauschformat", with DICOM remaining admissible for imaging, which is the first time a § 385 body has named the standard for the ordinance route rather than for a particular programme. It asks that terminologies, codes and value sets be made binding across sectors, naming LOINC and SNOMED CT as examples — while conceding that "eine vollständige Überprüfung im Rahmen von Konformitätsbewertungen nicht in allen Fällen möglich sein sollte", a concession about the reach of the assessment route that D-10.8 records from the other direction. And it asks that data models themselves be standardised in statute in future, "aufgrund der Herausforderung, dass Terminologien von Datenmodellen abhängig sind", naming DICOM and the aviation navigation standard ARINC 424 as precedents.

Two boundaries. This is a position and not law: the Council advises, the competence centre proposes, and the ordinance under § 385 SGB V is where a binding format is set. Nothing in D-10.3 about the operative text of the cabinet draft changes. And **openEHR is not mentioned in the paper, nor is any other persistence standard.** The body best placed to name one does not. That is recorded here as an absence in the document, not as a rejection by the Council, and no inference is drawn from it either way.

What follows for this document, and what does not. Four things, and no more.

First, the draft names no standard at all. The strings FHIR, openEHR, HL7, SNOMED, LOINC, IHE, ISO 13606 and DICOM occur nowhere in its 265 pages, the binding format being delegated to the ordinance route of § 385 SGB V. The German anchoring of FHIR therefore continues to run through § 355 and § 373 SGB V and through what the competence centre recommends, and not through a statutory naming of the standard. That is also where the anchoring could change, an ordinance being a lighter instrument to amend than a statute.

Second, the dispute in the consultation is the distinction this document draws throughout, and it is raised there by the affected parties rather than by this assessment: whether an interoperability obligation reaches the data that are exchanged or the model in which they are held. That the line is contested by those who must implement it is itself the finding, and it holds whichever way the enacted text falls.

Third, an obligation of this kind raises the standing cost of the exposure layer for every architecture, if enacted in this form, and does so symmetrically. It is not evidence for either persistence model. The consequence for procurement is stated at § 1.6.3.

Fourth, § 386c reaches further than the exposure layer alone, and in a direction this document has argued for on other grounds. A prohibition on releasing data in a format that prevents complete semantic reconstruction, and on withholding elements that the binding requirements do not cover but the documentation and archiving duties need, is a statutory formulation of the round-trip loss treated at § 1.6.2. It is model-neutral, it binds whichever persistence model a system uses, and it is the first German instrument to place the completeness of the released meaning, rather than the conformance of the released format, under a legal duty.

The core-profile working group of the Interop Council. Its position paper (v1.0, dated 14 March 2025 in the document-history table and 17 March 2025 on the cover page, published on the national interoperability platform) was written after the work on the ePA medication list and plan, and it states the problem in its own terms. The starting point was the insight that instead of a "weiteren Vervielfältigung hochspezialisierter, aber inkompatibler FHIR-Profile die Notwendigkeit für Use-Case- und sektorenübergreifende Kernprofile vorhanden ist". The paper

characterises the German base profiles as "unverbindliche Festlegungen aus der Community", that is, as a base layer without binding force.

Its first problem statement gives a concrete instance of the resulting incompatibility: `Patient.gender` is mandatory in ISiK and prohibited in the e-prescription profiles, for data-protection reasons, although both are conformant FHIR. Further problem statements name uncoordinated release cycles producing dependency conflicts, and record that no German specification has sufficient staff and financial resources for quality assurance and maintenance. The paper concludes that a national harmonisation "muss jedoch schnellstmöglich in Angriff genommen werden", and recommends coordinated annual release cycles, benchmarked against the American and Canadian practice.

The working group on hospital-interface adoption (position paper of 22 February 2024, based on 219 responses) records that "es fehlt bislang an einer abgestimmten, sektorenübergreifenden, iterativen Spezifikation", in the context of listing the hospital interfaces, the record types, the core dataset and the base profiles as mutually incompatible. It further records that a standardised interface alone, without defined use cases and without regard to the logic of the systems being connected, does not enable meaningful exchange. Productive use of the hospital interfaces at that date stood at 1.5% of the hospitals surveyed and 6.7% of certified vendors.

gematik's own assessment, in an interview with the trade press on 9 April 2024, was that the specification had looked "noch zu wenig auf konkrete Business- und Use Cases", that implementations were still immature for the care use case at hand, and that orientation and transparency were lacking. The same interview records the sentence "Wir spüren diesen Frust."

A structural consequence gematik states itself. The national interoperability platform records that because the German specifications depend on one another to a high degree, migration between major FHIR releases can succeed only if all the organisations involved agree on a common migration path with coordinated periods and transition windows. No German family can move to a later FHIR version unilaterally.

What the governance structure looks like once it is counted. The account above is a record of powers and processes. The derivation count of D-10.7 shows the same structure from the other end, in the artefacts. Reuse concentrates inside programmes — 721 of 1,110 in-corpus derivations do not leave their package, and ten hub artefacts carry 504 of them — while the national base layer appears as the structural parent of 62 of 2,317 resource profiles.

That distribution is what three determination powers in three statutes, plus a fourth family outside them, produce when each programme builds the hierarchy it needs on the timetable it has. It is also why cross-family conflicts of the `Patient.gender` kind are not surprising: there is no shared parent whose change would force both families to notice. The count is reported without a demand attached, because more derivation is not automatically better. A deeper hierarchy couples the release cycle of every descendant to its parent, and it is precisely the divergence of release cycles that this appendix records as the governance problem.

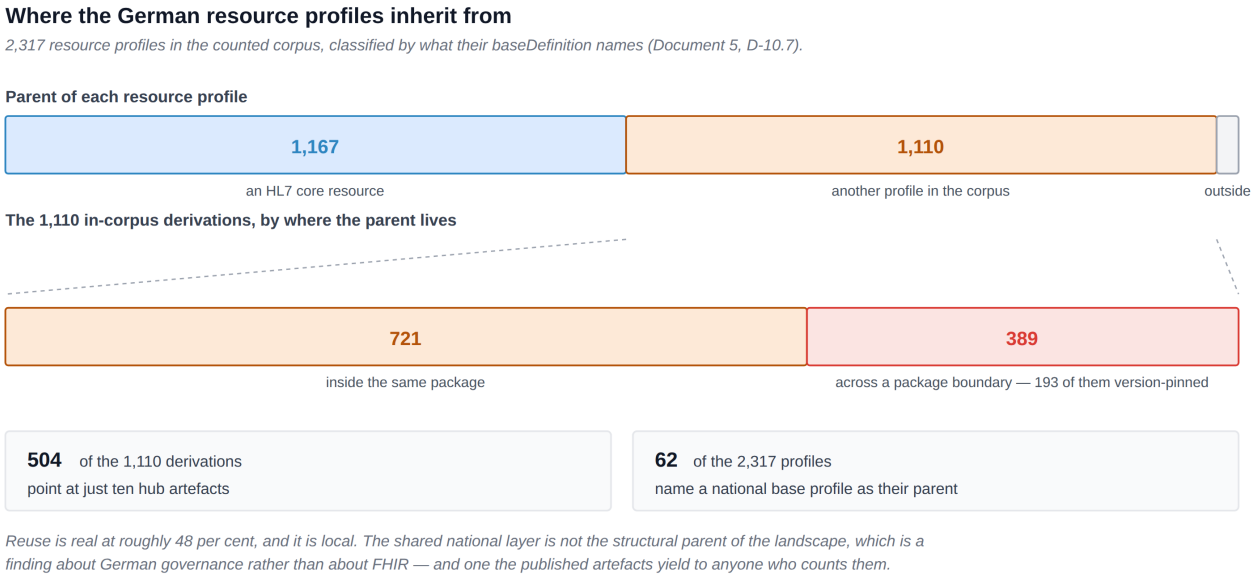

Figure — The derivation graph of the counted corpus, drawn to scale. Reuse is substantial and it is local: the shared national layer is the structural parent of 62 of 2,317 resource profiles (D-10.7).

The `Patient.gender` conflict is not a German failing, and saying so makes the finding smaller in one respect and larger in another. Where two bodies constrain the same base resource for different purposes, their constraints can be jointly unsatisfiable, and conformance testing on either side will not reveal it, because each side is conformant. That is a property of the profiling model and not of any national programme.

It has been computed elsewhere on the two most widely implemented guides in the world. An analysis of **US Core 5.0.1 against the International Patient Summary 1.0.0** finds Patient and Condition potentially infeasible, with `birthDate` and `clinicalStatus` required on one side and optional on the other and value sets bound at differing strengths, and MedicationRequest infeasible outright, because the summary guide carries reported medication as a different base resource and does not require support for `MedicationRequest` at all (Kramer & Moesel 2023; Document 4, D-14).

What is German about the German case is therefore not the conflict but its institutional treatment. Three determination powers in three statutes, no body competent to resolve a collision between them, the conflict recorded in a working-group paper in March 2025 and still open. The structural half is FHIR's and no national body can legislate it away. The unresolved half is governance and is precisely what a national body could fix.

Separating the two is not a defence of German practice. It is what makes the demand addressable, and it is why the corresponding request in Document 6 asks for a register and a responsible body rather than for a change to the standard.

The industry position is recorded in the association statement on the health-data-infrastructure bill of 17 May 2026, which asks that the interoperability provisions be drafted so as to enable the exchange of interoperable data "ohne interne Datenhaltungsmodelle gesetzlich zu normieren", and asks for a minimum lead time of eighteen months between publication of the requirements and the start of any certification obligation. The first of these is directly relevant to the question of this document: the German health-IT industry resists a statutory determination of the *internal* persistence model, in either direction.

D-10.4 The Trajectory, and What It Shows

STAGE	DATE	STORAGE TECHNOLOGY	WHAT WAS NEW
ePA 1.0	1 January 2021	IHE XDS.b, SOAP 1.2	documents; coarse-grained access control, without the differentiated permission model added later; opt-in
ePA 2.0	1 January 2022	still XDS.b	graduated permissions; the KBV record types — as documents in the registry, so FHIR as payload rather than as a queryable service
ePA 2.5 / 2.6	2023–2024	still XDS.b	the final opt-in generation
ePA 3.0, "for all"	15 January 2025 (records provided, model regions); 29 April 2025 (nationwide rollout); 1 October 2025 (mandatory for those providing care)	hybrid: XDS for documents **plus** FHIR data services	the first structured layer
ePA 3.1.3	pilot from summer 2026	FHIR medication service extended	the medication plan; dosages stored in structured, machine-readable form rather than as text
next stage	announced for the second half of 2027	FHIR	laboratory findings as structured data

One entry in that table is routinely compressed in the debate, and the compression misleads. "15 January 2025" is the date on which the sickness funds made the objection-based records available and on which those providing care in the model regions of Hamburg and surroundings, Franconia and parts of North Rhine-Westphalia began to use them. The nationwide rollout began on 29 April 2025, and use by those providing care became obligatory on 1 October 2025 [8]. **Evidence grade A.** Cited alone, the January date suggests a nationwide production deployment that did not exist on that day.

The statutory trajectory runs from the decision of principle in the 2003 health-system modernisation act, through the 2015 e-health act, the 2019 act that obliged the sickness funds to offer a record from 2021, the 2020 patient-data protection act that introduced §§ 341 ff. SGB V, to the 2024 digital act that replaced consent with objection.

The legacy document corpus has never been migrated semantically. The only migration step between record generations, from 2.6 to 3.0, was a document-level transfer with format checks, duplicate detection, exclusion of disallowed formats and conversion of eligible documents into permitted PDF/A formats. The conversion rules per document class were not established for this appendix, and nothing beyond the fact of format normalisation is asserted here. The step had to be triggered by the insured person through the app, and hospitals were told that no action was required of them.

This is a different thing from the service-internal migrations of the structured medication layer described in D-10.1, which do transform stored content. The distinction is between a corpus that was moved and a corpus that is

[8] Federal Ministry of Health, information pages on the record for all, retrieved 2 August 2026.

maintained, and only the second kind of migration has occurred. At no stage were stored documents converted into structured data. Old content was carried forward as it stood.

That is visible in the use figures, provided each is read for what it counts. Some 73 million records had been created by the sickness funds and more than 100 million documents deposited by 13 April 2026. The published breakdown is by document class, findings and reports somewhat above half, physicians' letters next, and 5.5 million dental bonus booklets, and not by media type. The aggregate therefore establishes no PDF share and none is asserted here.

Around 3.6% of the records had an active user at the turn of 2025/2026 according to a survey of the large sickness funds, and the largest fund reported some 8% in April 2026. These measure activation and use on the insured person's side. They do not measure the existence of the record, nor access by those providing care, nor the automatic population of the medication list, and they are insurer self-reports rather than a nationally weighted figure. **Evidence grade C.**

The judgement that the corpus is largely unstructured rests on a description rather than on those figures. The chair of the general practitioners' association called the result an "unsortierte PDF-Sammlung" that is costly to navigate because full-text search is not available. That was the state of the production system at the date of the criticism, and gematik announced in April 2026 that full-text search would be introduced during the year, so the sentence is time-bound and is not carried here as a permanent property.

The one structured artefact in circulation at scale is the electronic medication list, retrieved by those providing care more than 21 million times a week in the same reporting period. The metric counts retrieval operations, not distinct patients, distinct records or demonstrated clinical use, and the list is fed automatically from e-prescription data rather than being filled by those providing care.

What the case supports, and what it does not. It is not evidence against FHIR. It is the clearest national instance of the distinction drawn in § 1.5: each of the German specifications examined here is expressed as a conformant FHIR R4 implementation guide and can define internally valid artefacts, and they are not thereby interoperable with one another. It gives the limit named in § 2.4 — that profile discipline is a precondition — a documented case at the scale of a member state, recorded by the responsible bodies themselves.

A qualification keeps the case in balance. The fragmentation is a governance property rather than a property of the standard: it follows from three determination powers in three statutes, a fourth family standing outside them, and four separately managed release processes.

FHIR then makes the result visible and expensive, because a standard built on a common core plus profiles, extensions and terminology bindings pushes every national and use-case decision into an explicit artefact, and where those artefacts are governed independently, their divergence lands directly on implementers. That reading of the mechanism is this document's own. The eighty-percent principle it rests on is an informal design maxim of the standard rather than a measured law, and it is offered here as an explanation and not as a demonstration. A national governance could set binding national core profiles against it, which is what the working group quoted above asks for.

What the case does establish is narrow and holds firmly: a common exchange or persistence technology is no substitute for a shared semantic contract. It does not establish that the persistence model is without effect on longitudinal querying, validation, migration or secondary use, and no such claim is made here. The finding neither defends nor attacks the German bodies concerned. The bottleneck their own working groups describe is largely invariant against the choice of persistence standard, which is Finding 5 with a national evidence base.

D-10.5 Limits of This Appendix

Some figures in this appendix are counted from material supplied to the author, one is read from the artefact indexes the publisher issues, and some statements are expressly not verified. The three are held apart here rather than covered by one general disclaimer, and the per-count audit record is in D-10.6.

Counted on 12 August 2026, against the corpus and the convention described in the table note of D-10.2: the German base profiles across eight versions, the hospital interfaces at Stages 5 and 6, thirteen medical information objects, five `kbv.ita.*` packages and two `kbv.itv.*` packages, the KBV base layer, five appointment-service packages, 21 modules of the MII core dataset, three DEMIS packages, three accident-insurance packages, and the record packages at the versions supplied.

Not counted, and therefore asserted nowhere in this appendix:

- the remaining fifteen of the twenty `kbv.ita.*` packages;
- the DEMIS packages beyond the three held, including all five of the genomic-surveillance extension;
- the record packages of the current release line. The material holds 1.1.0, 1.1.1 and 1.1.2, so the resource-profile figure for that line is established by the other method, a reading of the published artefact indexes rather than a classification of package contents (D-10.6).

Where a count exists it is a count of the material as supplied. For repositories without a published archive it is a count of build output rather than of a release, a limitation that ten packages now escape, because for them a published archive was held and agreed with the build output exactly.

Further limits, each attaching to a particular statement:

- The binding dates for Stages 5 and 6 of the hospital interfaces were not established. The vendor portal still showed the stage-5 date as open after the stage-6 packages had shipped.
- Industry-association papers were read only where a ministry mirror made them available, so no position is attributed to the association beyond the statement of 17 May 2026.
- The full text of the 2019 federal audit report on the health card and the telematics infrastructure was not read. The account of it here rests on contemporaneous reporting in the medical press.
- Two secondary figures for records created, in late 2023 and early 2025, rest on a single encyclopaedic source and are therefore not carried in this appendix at all.
- The package `de.fhir.medication` is not itself held. What is read for the convergence finding is the dependency declarations of its dependents and not its contents, and how much of the guide each of them actually uses is not established.
- The date of the core-profile position paper carries two values in the document itself, 14 March 2025 in the document-history table and 17 March 2025 on the cover page. Both are given rather than one being chosen.

The sources retrieved directly from the publishing hosts on 2 August 2026 are the position paper, the ministry page on the record for all, the record-system specification, the conformance statement of the document responder, the artefact indexes and version histories of the three record implementation guides, and the dependency listing of the medication guide. The package registry and the specialist press remain unreachable from this working environment, and no claim rests on them that is not marked.

D-10.6 Audit Record for the Counted Figures

An audit record should identify the input material, date, artefact categories, inclusion and exclusion rules, and result. Strict reproduction additionally requires a frozen or content-addressed copy of that material. The classification rules are deterministic: each artefact is classified through `resourceType`, `kind` and `type` and

deduplicated by canonical URL. The table below records what varies between sources. The later reproducibility section states which inputs can be reconstructed and which preservation information is missing.

All figures were established on 12 August 2026. Columns: **P** resource profiles, **Par** Parameters profiles, **E** extensions, **D** datatype profiles, **L** logical models, **Σ** total StructureDefinitions. **Basis** distinguishes a published package archive from the build output of a source repository. Where a package is held as an archive, the archive is the identifier and the repository column is empty; where a repository publishes no package manifest, the repository name is the only identifier there is, and the package column says so in place of a name. The leading segment of the package identifiers of the Medical Informatics Initiative is elided and is the same for every module of that family. Older releases of packages that also appear here at a later version are not listed in this table but in the lineage tables below.

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
app-transport-framework	de.gematik.fhir.atf	1.4.1	Repo	3	0	1	2	0	6
basisprofil-de-r4	de.basisprofil.r4	1.6.0	Repo	19	0	29	28	0	76
basisprofileonkologie	de.basisprofil.onkologie	1.0.0-ballot	Repo	16	0	9	0	0	25
—	de.gematik.isik	6.0.0	Archive	156	2	10	9	0	177
—	de.gematik.isik-medikation	4.0.3	Archive	8	0	6	1	0	15
—	de.kvtelematik.eterminservice.r4	0.1.1	Archive	8	0	3	0	0	11
—	dguv.basis	1.5.0	Archive	18	0	4	1	0	23
—	dguv.enla	1.0.0	Archive	38	0	7	0	0	45
—	dguv.oper	0.1.0-Kommentierung	Archive	14	0	6	0	0	20
digitalepatientenrechnung	de.gematik.dipag	1.1.0	Repo	13	0	24	0	0	37
ebb	kbv.mio.bb	1.0.0-draft	Repo	20	0	3	0	0	23
einwilligungsmanagement	de.einwilligungsmanagement	2.0.3	Repo	10	0	11	0	0	21
epa-medication-v3-1	de.gematik.epa.medication	1.1.1	Repo	28	33	16	6	20	103
epa-research	de.gematik.epa.research	1.1.2	Repo	3	0	1	1	0	5
epa-v3-1	de.gematik.epa	1.1.0	Repo	14	0	4	0	0	18

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
erezept-errorcodes	de.gematik.fhir.erp-ec	0.0.1	Repo	0	0	0	0	0	0
erezept-patientenrechnung	de.gematik.erezept-patientenrechnung.r4	1.1.0	Repo	4	1	1	0	0	6
erezept-servicerequest	de.gematik.erp-servicerequest	1.2.0	Repo	9	0	8	3	7	27
erezept-workflow	de.gematik.erezept-workflow.r4	1.6.4	Repo	16	6	13	3	2	40
erezept-workflow-eu	de.gematik.erezept.eu	1.1.2	Repo	6	6	2	1	0	15
erp-t-prescription	de.gematik.erp.t-prescription	1.2.0	Repo	4	1	0	0	1	6
fhir-profiles	de.gematik.hddt	1.0.1	Repo	9	0	0	0	0	9
isik-basis-v3	de.gematik.isik-basismodul	3.1.1	Repo	14	0	0	1	0	15
isik-dokumentenaustausch-v3	de.gematik.isik-dokumentenaustausch	3.1.1	Repo	2	0	0	0	0	2
isik-medikation-v3	de.gematik.isik-medikation	3.1.1	Repo	7	0	0	0	0	7
isik-sicherheit-v2	*documentation only*	—	Repo	0	0	0	0	0	0
isik-sicherheit-v3	*documentation only*	—	Repo	0	0	0	0	0	0
isik-stufe-5	de.gematik.isik	5.1.3	Repo	109	0	9	9	0	127
isik-stufe-6	de.gematik.isik	6.0.0	Repo	156	2	10	9	0	177
isik-terminplanung-v2	de.gematik.isik-terminplanung	2.0.6	Repo	6	0	2	0	0	8
isik-terminplanung-v3	de.gematik.isik-terminplanung	3.1.1	Repo	6	0	2	0	0	8

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
isik-vitalparameter-und-koerpermasze-v3	de.gematik.isik-vitalparameter	3.1.1	Repo	10	0	0	0	0	10
isip	de.gematik.isip	1.0.2	Repo	5	0	0	0	0	5
—	kbv.all.st	1.46.0	Archive	0	0	0	0	0	0
—	kbv.all.st-combined	1.47.0	Archive	0	0	0	0	0	0
—	kbv.all.st-rc	1.21.0	Archive	0	0	0	0	0	0
—	kbv.basis	1.9.0	Archive	37	0	4	7	0	48
—	kbv.basis.terminology	1.9.0	Archive	0	0	0	0	0	0
—	kbv.ita.aws	1.2.0	Archive	117	0	52	0	0	169
—	kbv.ita.eau	1.2.1	Archive	10	0	7	0	0	17
—	kbv.ita.erp	1.4.0	Archive	8	0	15	0	0	23
—	kbv.ita.for	1.3.1	Archive	5	0	6	0	0	11
—	kbv.ita.vos	2.2.0-kommentierung2	Archive	26	0	24	0	0	50
—	kbv.itv.evdga	1.2.2	Archive	3	0	0	0	0	3
—	kbv.itv.ks	2.0.0	Archive	7	0	8	0	0	15
—	kbv.mio.ddtk	1.0.0-kommentierung	Archive	16	0	0	0	0	16
—	kbv.mio.diga	1.1.0	Archive	44	0	4	0	0	48
—	kbv.mio.emp	1.0.0-kommentierung.1	Archive	20	0	0	1	0	21
—	kbv.mio.impfpass	1.1.0	Archive	15	0	7	0	0	22
—	kbv.mio.kh-entlassbrief	1.0.0-update	Archive	34	0	7	0	0	41
—	kbv.mio.laborbefund	1.0.0-update	Archive	24	0	14	0	0	38
—	kbv.mio.mutterpass	1.1.0	Archive	99	0	4	0	0	103
—	kbv.mio.patientenkurzakte	1.0.0	Archive	27	0	5	1	0	33

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
—	kbv.mio.tele	1.0.0	Archive	36	0	5	0	0	41
—	kbv.mio.u-heft	1.0.1-festlegungsversion	Archive	210	0	4	0	0	214
—	kbv.mio.ueberleitungsbogen	1.0.0	Archive	90	0	19	1	0	110
—	kbv.mio.zaeb	1.1.0	Archive	6	0	1	0	0	7
—	kvdigital.terminschnittstelle-fuer-dritte	2.8.0	Archive	10	0	6	1	0	17
—	kvdigital.terminschnittstelle-fuer-kven	2.5.0	Archive	11	0	0	1	0	12
—	kvdigital.terminsynchronisation-tvs	1.10.0	Archive	12	0	5	0	0	17
—	kvdigital.vermittlungscode-abrufen-pvs	1.12.0	Archive	3	0	0	1	0	4
mii-modul-bildgebung	…kerndatensatz.bildgebung	2026.0.0	Repo	13	0	14	0	1	28
mii-modul-mikrobiologie	*manifest name* *project*	2027.0.0-alpha.5	Repo	25	0	1	0	3	29
mii-modul-studien	…kerndatensatz.studie	2026.0.2	Repo	7	0	10	0	1	18
mii-kerndatensatz (meta)	…kerndatensatz.meta	2026.0.0	Repo	1	0	1	1	0	3
mii-modul-intensivmedizin	…kerndatensatz.icu	2026.0.2	Repo	93	0	0	0	1	94
mii-modul-biobank	…kerndatensatz.biobank	2026.0.1	Repo	11	0	10	0	1	22
mii-modul-consent	…kerndatensatz.consent	2026.0.1-rc-2	Repo	3	0	0	0	0	3
mii-modul-omics	…kerndatensatz.molgen	2026.0.4	Repo	16	0	5	0	1	22
mii-modul-onkologie	…kerndatensatz.onkologie	2026.0.3-rc.1	Repo	74	0	14	0	3	91
mii-modul-pathologie	…kerndatensatz.patho	2026.0.2	Repo	17	0	0	0	1	18

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
mii-modul-strukturdaten	kerndatensatz modulstrukturdaten	0.1.0	Repo	3	0	1	0	2	6
mii-modul-symptome	…kerndatensatz.symptom	2026.0.0	Repo	3	0	0	0	1	4
mii-basismodul-labor	…kerndatensatz.laborbefund	2026.0.3	Repo	3	0	2	0	1	6
mii-basismodul-medikation	…kerndatensatz.medikation	2026.0.1	Repo	5	0	2	0	1	8
mii-erweiterungsmodul-dokument	…kerndatensatz.dokument	2026.0.1	Repo	1	0	1	0	1	3
mii-erweiterungsmodul-kardiologie	…kerndatensatz.kardiologie	2026.0.0-alpha.3	Repo	21	0	0	0	0	21
mii-erweiterungsmodul-lungenfunktion	…kerndatensatz.lungenfunktion	2026.0.0	Repo	50	0	0	0	1	51
mii-erweiterungsmodul-mtb	…kerndatensatz.mtb	2025.0.0-ballot	Repo	49	0	11	0	4	64
mii-erweiterungsmodul-pro	…kerndatensatz.pros	2026.5.2	Repo	23	0	2	0	1	26
mii-kerndatensatz-basis	…kerndatensatz.base	2026.0.1	Repo	7	0	1	0	4	12
mii-modul-seltene-erkrankungen	…kerndatensatz.seltene	2026.0.1	Repo	17	0	11	0	0	28
rki	rki.demis.laboratory	5.0.1-alpha.10	Repo	285	0	2	2	0	289
rki-2	rki.demis.disease	3.0.0-rc.9	Repo	187	0	2	0	0	189
rki-3	rki.demis.common	3.0.0-rc.11	Repo	17	2	8	4	0	31
schuleingangsuntersuchung	de.gematik.schuleingangsuntersuchung	1.0.0	Repo	2	0	17	0	0	19
sterbefall	de.gematik.sterbefall	1.0.2	Repo	15	0	3	0	0	18
tim	*no FHIR manifest*	—	Repo	4	0	0	0	0	4

REPOSITORY	PACKAGE	VERSION	BASIS	P	PAR	E	D	L	Σ
vsdm-ersatzbescheinigung	`de.gematik.elektronische-versicherungsbescheinigung`	1.2.0	Repo	11	0	11	0	0	22
vsdm2	`de.gematik.spec-vsdv2`	1.1.2	Repo	7	0	8	1	2	18
vzd-fhir-directory	`de.gematik.fhir.directory`	1.2.0	Repo	12	0	11	1	0	24

Ten packages are held both as a published archive and as a source repository. Nine of them — the KBV base layer and eight KBV record types — are listed above at the archive, because a published release is reproducible where build output is not, and the repository rows they replace are omitted rather than duplicated. The tenth, `de.gematik.isik` 6.0.0, is listed twice on purpose, because its repository row also carries the Stage 5 lineage beside it. In all ten cases the two readings produced **identical canonical URL sets**. Two further archives are packaging variants of the KBV base layer at an older release, and one is the previous major version of the maternity record; all three appear in the lineage tables rather than here.

The column sums exceed the corpus total, because the total is deduplicated across sources by canonical URL while this table is not. The **12 August corpus** contains **2,996 distinct StructureDefinitions**, of which 2,317 are resource profiles, 51 Parameters profiles, 491 extensions, 77 datatype profiles and 60 logical models. Beside them stand **5,423 conformance artefacts that are not StructureDefinitions**, summed per source without cross-source deduplication and counting each archive-and-repository pair once: 2,730 value sets, 1,678 code systems, 309 concept maps, 286 search parameters, 183 capability statements, 163 naming systems and 74 operation definitions.

Five packages carry no StructureDefinition at all. The three `kbv.all.st` variants and the two KBV terminology packages are value-set and code-system deliveries. They are listed with zero rather than omitted, because a reader who holds them and counts differently should be able to see that the difference is not an oversight.

One corpus, one pass, one date. Everything in this appendix is counted in a single pass over a single source set, counted on 12 August 2026. There is no second corpus to reconcile against, and no figure here is carried over from a differently assembled set.

The base-profile lineage. The seven supplied archives and the 1.6.0 repository give the following, on the same convention.

VERSION	P	E	D	Σ
1.0.0	15	23	21	59
1.1.0	16	23	21	60
1.2.0	18	25	23	66
1.3.2	18	27	24	69
1.4.0	18	27	25	70
1.5.2	19	29	27	75
1.5.4	19	29	28	76
1.6.0	19	29	28	76

A second lineage, shorter and newer. The KBV base layer is held at two releases, and the older one in two packaging variants that are identical to each other in StructureDefinition content.

PACKAGE	VERSION	P	E	D	Σ
`kbv.basis`	1.9.0	37	4	7	48
`kbv.basis.ressources.only`	1.7.0	34	6	7	47
`kbv.basis.with.expansion`	1.7.0	34	6	7	47

The maternity record is likewise held at two major versions, `MIO.KBV.Mutterpass` 1.0.0 with 102 StructureDefinitions and `kbv.mio.mutterpass` 1.1.0 with 103, sharing 97 canonical URLs. Neither lineage is folded into the corpus twice: where two releases of one package are present, the corpus keeps whichever the deduplication rule selects, and the sensitivity of that selection is measured in D-10.7.

Auditability and reproducibility are different. For source repositories, the figures count build output on the date given. Build output is not fixed by a package version, and a later checkout may produce a different inventory. The input material is not published with this report, and the working mirror does not retain every exact repository revision. The record therefore makes the derivation auditable — source, stated version where available, category, rule and result — but does not make every count independently re-executable from a frozen public corpus. Only the identified published archives are reproducible in the strict sense, and they now carry a larger share of the corpus than they did: 46 of the sources in the table above are archives.

The record packages are counted by a different method, and it is held apart. The three implementation guides of the current ePA release line are not held as packages here. Their figures are read from the artefact index each guide publishes, under the section headings the guide itself uses — Ressourcenprofile, Parameterprofile für Operationen, Erweiterungen, Datentypen — which correspond to the categories of the convention above. The mapping between the two is the author's, and the three index pages are at the footnote [9].

IMPLEMENTATION GUIDE	PACKAGE	VERSION	RESOURCE PROFILES
ePA Basisfunktionalitäten	`de.gematik.epa`	1.3.2	3
ePA Medication Service	`de.gematik.epa.medication`	1.3.5	10
ePA MHD Service	`de.gematik.epa.mhd`	1.1.3	2
Total			**15**

An index reading and a classification of package contents are not interchangeable. The index reflects what the publisher chose to display and how it chose to group it, while the classification reflects the declared fields of every artefact in the package. The figures are therefore marked as read rather than as counted wherever they appear. One observation belongs with them: the medication service carries fewer resource profiles in the current line than in the archived 1.1.1, which holds 28. What was restructured between the two lines is not established here.

Audit record for the binding and derivation counts (D-10.7). Material: JSON and XML artefacts from the 66 source repositories and 46 package archives, including a base-profile archive located inside the school-entry-examination repository. Classification and deduplication follow D-10.2. Bindings are counted only where declared in an artefact's own `differential`, once per binding. The listed `fixed[x]` and `pattern[x]` constraints are counted separately.

[9] Published artefact indexes of the three guides of the current release line, all retrieved 2 August 2026: `gemspec.gematik.de/ig/fhir/epa/1.3.2/artifacts.html`, `gemspec.gematik.de/ig/fhir/epa-medication/1.3.5/artifacts.html` and `gemspec.gematik.de/ig/fhir/epa-mhd/1.1.3/artifacts.html`.

Where a canonical URL recurs, the first in sorted path order is retained, collapsing 1,608 occurrences. Non-German canonical URLs are excluded. Inherited bindings are not resolved, and family assignment is based on repository path.

Results. 2,996 StructureDefinitions, 2,317 of them resource profiles; 2,399 explicit differential bindings; 8,073 listed fixed or pattern constraints; 94 safety-critical resource profiles. Tie-break sensitivity was tested under five rules, and artefact counts and binding counts varied by about one percent.

Snapshots, in full, because the figure is easy to misread. Under the chosen rule 113 artefacts carry a snapshot. Across the five rules the figure runs from 99 to 176, and it is the one count that is not invariant, because a canonical URL can occur once with a snapshot and once without.

The denominator matters more than the rule. Of the distinct canonical URLs in the corpus, roughly half are held as a published package and roughly half only as a source repository, and a source repository is not expected to carry a snapshot: the IG Publisher writes them into the output package. Counted on the published half alone, 77 artefacts carry a snapshot. For the other half the question was not put, because the corresponding package is not in the corpus.

No proposition in this appendix rests on any of these figures. They are recorded so that the counting rule for bindings, differentials only, has its reason stated and can be challenged. One internal control belongs with them, and it is a statement about the toolchain rather than about any publisher: the corpus holds six archived versions of one base-profile package, of which the five earlier ones carry no snapshot and the latest carries one for every artefact. Whether a package ships snapshots is therefore a publication setting, not a limit of the tooling.

The counting implementation was rewritten and tested against the previous source set before use. It reproduced the earlier artefact, profile, duplicate, snapshot and binding totals within the sensitivity bounds previously reported. Two XML classification defects were found and corrected before the new figures were carried forward: the `context` element of an extension and the `mapping` block of a StructureDefinition had initially been mistaken for the artefact's own type and kind. The comparison also exposed a defect in the previously published derivation-graph figures, corrected in D-10.7.

D-10.7 Binding Strength Across the Counted Corpus

§ 2.5.2 states what the two specifications permit. This section examines the terminology bindings and selected fixed or pattern constraints explicitly declared in the differentials of the counted German artefacts. It does not resolve inherited constraints or measure implementation behaviour.

Method. The corpus is the one described in D-10.2, counted in a single pass on 12 August 2026 under the rules set out at the end of this report. Family assignment is based on repository path and is a presentation convenience rather than part of the artefact classification.

Only bindings declared in the artefact's own `differential` are counted. A profile may inherit bindings from its base definition, and **the great majority of the artefacts held carry no snapshot**, so inherited bindings cannot be resolved from this material and are not included. That is a statement about the material in hand and not about what the responsible bodies publish: roughly half the corpus is held only in source-repository form, where a snapshot is not expected, and the figures with both denominators are in the audit record above.

The tie-break rule, and a test of whether it matters. Where a canonical URL occurs more than once, the first in sorted path order is kept. Alphabetical path order is a deterministic rule and not a subject-matter one: it does not prefer the published release over build output, or the higher version over the lower. That matters more than it did, because the corpus now holds ten packages in both forms and three packages at two releases.

Rather than assert that the rule is harmless, it was measured. The whole count was repeated under five tie-break rules: path ascending, path descending, archive before repository, repository before archive, and the copy carrying the most bindings first. **The artefact counts are invariant across all five: 2,996 StructureDefinitions and 2,317 resource profiles in every run.** The binding counts move within a band of 23 on a total of about 2,400, about one percent: total bindings 2,379 to 2,402, `required` 1,899 to 1,908, `example` 42 to 45, fixed constraints 8,073 to 8,075.

No binding, derivation or classification statement in this appendix changes under any of the five rules. One figure does move, the number of artefacts carrying a snapshot, and it is treated in full in the audit record at D-10.6 rather than summarised here. For everything this appendix does assert, which release of a package the rule happens to select does not change what the corpus contains.

What the profiles declare. Of 2,996 artefacts, 1,338 — 44.7% — carry at least one binding of their own. Those artefacts declare **2,399 bindings** in total, distributed as follows.

STRENGTH	BINDINGS	SHARE
`required`	1,906	79.4%
`extensible`	336	14.0%
`preferred`	112	4.7%
`example`	45	1.9%

Twenty bindings carry a strength without a value set. These are author-derived counts from primary artefacts. The method and source inventory are documented, but independent byte-for-byte reproduction is not currently possible because the frozen input corpus, exact repository revisions and a complete frozen input corpus are unavailable.

FAMILY	REQUIRED	EXTENSIBLE	PREFERRED	EXAMPLE	TOTAL	REQUIRED SHARE
MII core dataset	444	154	34	6	638	69.6%
KBV record types, forms and base layer	499	57	23	35	614	81.3%
DEMIS / RKI	334	58	0	0	392	85.2%
ePA and e-prescriptio n	214	13	42	0	269	79.6%
ISiK hospital interfaces	200	17	4	0	221	90.5%
Deutsche Basisprofil e	70	20	3	0	93	75.3%
Other	72	15	4	1	92	78.3%

FAMILY	REQUIRED	EXTENSIBLE	PREFERRED	EXAMPLE	TOTAL	REQUIRED SHARE
Appointment services	50	0	0	0	50	100%
DGUV accident insurance	23	2	2	3	30	76.7%

The family rule is stated here in full rather than described, so that a later count can be compared row by row and not only in total. The source path is lower-cased and tested against this ordered list of substrings, first match winning: `dguv` → DGUV; `isik` → ISiK; `isip` → Other; `rki` → DEMIS; `basisprofil` → Deutsche Basisprofile; `medizininformatik` or `mii-` → MII; `epa`, `erezept` or `erp-` → ePA and e-prescription; `kvdigital` or `kvtelematik` → Appointment services; `kbv`, `mio`, `ddtk`, `ebb`, `khe`, `lab1x0`, `mp1x0`, `pka`, `uh1x0`, `ulb`, `zb1x0` or `base1x0` → KBV; anything unmatched → Other. Family assignment shapes the presentation and not the corpus totals; the rows sum to the corpus total of 2,399.

The appointment-service row is small and its hundred percent should not be read as a result. Fifty bindings across five packages, all `required`, is what a specification looks like when its coded content is almost entirely administrative — appointment status, cancellation reason, communication channel — and those value sets are closed by their nature. Nothing follows from it about clinical content.

The first result runs against the expectation § 2.5.2 raises. Where the German families bind at all, they bind hard: close to four fifths of all declared bindings are `required`, and `example` — the strength that enforces nothing — accounts for fewer than one binding in fifty. The concern that profiling discipline dissolves at the terminology layer is not borne out by this corpus. That is a finding in favour of the German profiling work, and it belongs here for the same reason the findings against it do.

The second result qualifies the first. A profile can constrain an element without using `binding`, for example through `fixedCode`, `fixedCoding`, `patternCoding`, `fixedUri` or related constraints. Across the corpus, the differentials contain **8,073 listed fixed or pattern constraints and 2,399 explicit bindings**. This shows that binding counts describe only one part of the constraint surface. The current count does not classify the fixed and pattern constraints by element path or datatype, so it cannot establish how many govern coded clinical content, terminology systems or unrelated URI and structural values. Determining whether value fixing is the dominant terminology mechanism requires that additional classification.

Third, the safety-critical subset, where § 1.6.2 asks for the strictest treatment. The corpus holds **94 resource profiles** on `Medication`, `MedicationRequest`, `MedicationStatement`, `MedicationAdministration`, `MedicationDispense` and `AllergyIntolerance` — 35 on Medication, 18 on MedicationStatement, 14 on MedicationRequest, 12 on AllergyIntolerance, 8 on MedicationDispense, 7 on MedicationAdministration. Of these, **50 declare no binding at all.** Taken alone that number would carry a conclusion it does not support: 38 of the 50 fix values instead, 46 slice, and only **two** carry neither a binding nor a fixed value nor a slice — `SystemischeTherapie` and `mii-pr-bildgebung-kontrastmittelgabe`, both thin profiles of a few differential elements. The detection counts a profile as constrained where it carries `fixedString`, `fixedBoolean` or `fixedCanonical`, or a `sliceName` without an explicit slicing block; a narrower detection rule would report more unconstrained profiles without any difference in the artefacts.

Among the bindings explicitly declared in these differentials, the distribution is **79 `required`, 13 `extensible`, 40 `preferred` and 15 `example`**, a `required` share of 53.7% compared with 79.4% across the corpus. On a corpus a fifth larger than the earlier one, that share moved by a tenth of a percentage point. These are author-derived counts from the identified primary artefacts. They do not establish that the subset is less strictly constrained overall:

inherited bindings are unresolved, `extensible` may be appropriate for an open concept space, and each binding would have to be assessed against its element and clinical purpose. No claim is made about clinical consequences.

Fourth, the profile hierarchy — what the corpus reuses of itself. Every `StructureDefinition` names the artefact it constrains in `baseDefinition`, so derivation across the corpus is a directed graph and can be counted rather than asserted. The count below was run on the same artefacts as the rest of this appendix, with the same rules — German namespace, deduplication by canonical URL, first in ascending path order — and it reproduces the corpus totals exactly: 2,996 artefacts, 2,317 resource profiles, 1,608 collapsed duplicates. That agreement is the reason the derivation figures are reported at the same standard as the binding figures.

WHERE A RESOURCE PROFILE'S PARENT SITS	PROFILES	SHARE
An HL7 core resource, constrained directly	1,167	50.4%
Another profile inside this corpus	1,110	47.9%
A German artefact outside the collected corpus	40	1.7%
Total	**2,317**	**100%**

Reuse is real, and it is local. Of the 1,110 in-corpus derivations, **721 stay inside the same package or repository and 389 cross a package boundary**. The corpus has 179 distinct parent artefacts, and the ten most-used of them carry 504 of the 1,110 derivations, so nearly half of all reuse runs through ten **hubs** — the term is used here for an artefact from which many others derive, and it says nothing about the artefact's importance beyond that count. Those hubs are programme-internal rather than national: the most-derived-from artefacts are the DEMIS notification profiles of the Robert Koch Institute — `PathogenDetection` and `Specimen` with 91 descendants each, `LaboratoryReport` with 85 — followed by intensive-care profiles of the Medical Informatics Initiative, `KBV_PR_Base_Organization` with 26, and ISiK counterparts.

The share of direct HL7 profiling rose as the corpus grew. On the earlier and smaller source set, 47.3% of resource profiles constrained an HL7 core resource directly; with 331 further resource profiles counted, the figure is 50.4%. The packages added are KBV forms, record types and appointment services, and they profile core resources directly more often than the average of the corpus they joined. This is a property of what was added, not a change in anyone's practice, and it is recorded because a reader comparing the two releases will otherwise read a trend into it.

The national base-profile layer is not what carries the reuse. Only **62 of 2,317 resource profiles (2.7%) name a `http://fhir.de/StructureDefinition/` artefact as their parent**, and 134 (5.8%) have one anywhere in their ancestry.

That does not establish that the German base profiles are ignored, and the corpus shows why: of the 77 `fhir.de` artefacts in it, 29 are extensions and 28 are datatype profiles, and only 20 are resource profiles. A base layer built mostly from datatypes, identifiers and extensions is consumed by reference from inside a differential, not by derivation, and that consumption does not appear in this graph at all. What the figures do establish is narrower: **the shared national layer is not the structural parent of the German profile landscape**, and the ecosystem's shape is set by a small number of programme hubs instead. The largest single `fhir.de` hub is `observation-de-vitalsign` with 116 descendants counted transitively.

Chains are short. Measured across the whole corpus, 1,873 artefacts derive from something outside it, 917 sit one step below an in-corpus parent, 132 two steps, 69 three steps and 5 four steps. No chain inside the corpus runs deeper than four. Extensions are almost entirely flat: 490 of 491 derive directly from the HL7 `Extension` base, which is expected and means there is no extension reuse layer to speak of.

Two cautions travel with these figures. Derivation is a structural relation and not a semantic one: a profile that names a parent and then redefines most of what it inherited has reused a URL rather than a model, and only reading the differentials would show which happened. And the absence of derivation is not by itself a defect, because a profile constraining an HL7 base resource directly may be doing exactly the right thing for its use case. The count establishes the shape of the ecosystem, not its quality.

Fifth, how the dependencies are versioned — and this is where the reuse question turns into a version-management question. A derivation is a dependency, and a dependency is only as stable as the version binding that carries it. Counted over the same corpus:

DEPENDENCY	TOTAL	VERSION-BOUND	SHARE
`baseDefinition` references	2,996	204	6.8%
`binding.valueSet` references in differentials	2,379	232	9.8%
Dependency statements in package manifests	364	325	89.3%

The two layers divide the work, and the division is deliberate. At package level the German projects pin: 325 of 364 dependency statements name an exact version. At artefact level they mostly do not: 93.2% of parent references and 90.2% of value-set references carry no version. That is not a defect, and reading it as one would misdescribe how FHIR is meant to work. An unversioned canonical is resolved through the package context, so within an implementation guide, its publication run and any validation performed against its declared dependencies, the version of a parent is fixed — by the manifest rather than by the artefact. **Version management exists here, and it lives one layer above the file.**

Where the artefacts do pin, they pin at the boundary. Split the in-corpus derivations by whether they leave their package: of the **721 that stay inside a package, not one** names a version; of the **389 that cross a package boundary, 193 — 49.6% — do**. The KBV base profiles derive from the German base profiles as `…|1.5`, and that pattern repeats wherever one programme builds on another's artefact.

Version pinning in this corpus is therefore not absent from the artefact layer. It is concentrated exactly where two independently governed release cycles meet, and absent exactly where a single release cycle already guarantees consistency. Reading this correctly depends on one counting rule: a `baseDefinition` that names its parent with a version identifies the same parent, and a graph that compares the strings exactly books such derivations as lying outside the corpus, so they never reach the cross-package bucket at all.

Three residues remain, and they are what a governance discussion should be about.

First, the pinning is not complete. 39 of the 364 manifest statements do not name an exact version. Thirty-five use a rolling form, 18 of them `2026.0.x` and the rest the equivalent `x` notation on other lines, so the dependency follows the patch stream of the named package. **Four name `latest`**, which is a different thing again: it does not follow a patch stream but whatever the registry serves at the moment of resolution, and a build repeated a month later can therefore consume different content without any file having changed.

Second, resolution is only as reliable as the context. An artefact consumed outside a guide has no manifest to be resolved through: in a profile or terminology registry, on a validation server that holds several families side by side, in an aggregated corpus such as this one, or in tooling that resolves canonicals against whatever is loaded rather than against a manifest.

Third, and least visible, an unpinned artefact records nothing about the version it was written against. After a manifest is lifted to a new parent release, no reader of the child can tell from the file whether its constraints were

authored against the old parent or the new one, the intent being recoverable only from the repository history. For the 196 cross-package derivations that carry no version, that is the situation.

How the derivation count was produced. A single scan recorded, for every `StructureDefinition` found in JSON and in XML, the canonical URL, `kind`, `type`, `derivation`, `baseDefinition`, the declared bindings and the listed fixed and pattern constraints. The scan finds 4,628 artefact occurrences, and after the German-namespace filter and deduplication it yields 2,996 artefacts and 2,317 resource profiles — the published corpus. A parent counts as *inside the corpus* when its canonical URL, with any version suffix removed, is one of the 2,996 kept URLs; as *HL7 core* when it begins with `http://hl7.org/fhir/StructureDefinition/`; and as *outside* otherwise. Package boundaries are the top-level directory or archive name, so "crossing a package boundary" means the parent was published by a different project in this collection. Ancestry is followed transitively with a cycle guard; no cycles were found.

What this appendix does not establish. Whether an element left unbound in a differential is nevertheless bound by its base definition, which the absence of snapshots from most of the material held prevents from being resolved. Whether a fixed value is clinically equivalent to the binding it replaces. Whether the value sets that the `required` bindings point at are themselves published, versioned and resolvable, which is the question § 2.5.2 puts to procurement and which a count of bindings cannot answer.

D-10.8 Conformity Assessment — Where Profiling, Testing and Market Access Meet

Everything above describes specifications: who determines them, how many there are, how tightly they bind and how their versions drift apart. It says nothing about the step that decides whether a specification reaches a product. In Germany that step now exists as a formal instrument, and it is the only place in the landscape described here where profiling, testing and market access are joined.

What the instrument is. § 385 SGB V establishes binding interoperability requirements, and § 387 SGB V charges the Competence Centre for Interoperability in Health Care at gematik with checking the compliance of information-technology systems with those requirements and issuing confirmations. The procedure is called conformity assessment (*Konformitätsbewertung*, KOB). The centre supplies test suites, so that testing runs decentrally and alongside implementation rather than as a single examination at the end; systems that pass are confirmed and entered on a public positive list in the interoperability navigator INA. Participation is not a matter of choice for the systems in scope. It was introduced as an obligation together with the electronic patient record for all.

What is in scope today. Not FHIR conformance in general, but a defined artefact set. The current subject is the ePA medication service, with practice, hospital and pharmacy systems assessed on their handling of the medication list. The centre names the electronic patient record at stage 3 and ISiK at stage 5 as the next subjects, and an extension to the medication plan for 2026.

The reported state, with the date it belongs to. As of 11 December 2025 the operator records 141 systems successfully assessed and states that these cover some 98% of the practice-software market. Both figures are self-reported by the body running the procedure and grow with the procedure itself. They are carried here at **evidence grade D**, with the date attached, while the statutory obligation is grade A for what the statute requires and the published description of the procedure is grade B for what the responsible body says it does.

Why it belongs in this report. This report records four specification families determined by four bodies under three statutes, with release cycles that drift apart (D-10.2, D-10.3). Conformity assessment does not repair that. What it does is make conformance to one named family checkable at a defined point in time, and attach a consequence to the result. That is the missing half of every governance argument in this appendix: a specification without an assessment route binds no one in practice, however precisely it is written. Whoever asks what a European conformance regime would have to look like will find here a working national instance of the machinery, complete with its legal basis, its test suites and its published register.

What it does not settle, and the point is the same one this report makes throughout. A passed test suite is a statement about the test. It records that a system answered the cases the suite contains, at the profiling depth the suite reaches, for the artefact set in scope. It does not establish that the exchanged content is clinically usable at the receiving end, and the fourth trap set out below applies to it in full. The distinction is the one § 1.5 of Document 3 draws between conformance and interoperability, and conformity assessment sits squarely on the conformance side of it. Its value is that it makes the conformance side enforceable, which is not nothing and is not the whole thing either.

Interests. ISiK, one of the announced subjects of the procedure, is a gematik specification family developed in close proximity to the German HL7 community, which the author of this report leads. The propositions recorded here are readings of a statute and of the responsible body's own published description, checkable against both, and no proposition about the quality of the procedure is asserted.

Annex: Corpus, Counting Rules and Methodological Limits

This section defines the corpus the figures rest on, the rules applied to it, and the limits that bear on how the counts may be read. The method and the source inventory are auditable. Byte-for-byte reproduction from this document alone is not possible, because the exact revisions of every source repository and a frozen copy of the input corpus were not preserved publicly.

The corpus

The corpus was counted on 12 August 2026. It draws on 66 German source repositories and 46 published package archives, each named in the audit table at D-10.6 with its package, its stated version where one exists, and whether the count uses a published release or build output. Every artefact whose `resourceType` is `StructureDefinition` was considered, in JSON and in XML alike. After the namespace filter and deduplication by canonical URL, the corpus holds 2,996 distinct StructureDefinitions, of which 2,317 are resource profiles.

The corpus is an identified collection of German FHIR artefacts and not a complete inventory of every German FHIR profile. Published archives are fixed releases. Repository build output is the state available on the counting date and can change without a version bump, so figures drawn from it are bound to that date. Both properties are recorded in the audit table rather than treated as noise.

A digest over the corpus is maintained, and its value is not printed here. A SHA-256 over the sorted list of canonical URLs, concatenated without separators and encoded as UTF-8, fixes what the word *corpus* refers to on a given date. It is held in the project record together with the digest of the preceding corpus. Its one purpose is to establish whether a recount is working on the same material, before anyone argues about method. A reader who will not obtain the packages and run the rules gains nothing from the value itself, and no finding in this report depends on it having been checked.

The counting rules

1. **Inclusion.** Every artefact whose `resourceType` is `StructureDefinition`, in JSON and in XML alike.
2. **Namespace filter.** Exclude artefacts whose canonical URL lies outside a German namespace. Sixteen are present, all R5 cross-version backport extensions that the IG Publisher copies into build output.
3. **Classification.** `kind logical` as a logical model; `type Extension` as an extension; `kind resource` with `type Parameters` as a Parameters profile; any other `kind resource` as a resource profile; `complex-type` or `primitive-type` as a datatype profile.

4. **Deduplication.** By canonical URL across the whole corpus, keeping the first occurrence in sorted path order. This also decides which of two releases of one package the corpus keeps. Measured under five orderings the artefact totals are identical and the binding total moves by about one percent; snapshot counts are more sensitive to the choice and are therefore not used as a finding.
5. **Bindings.** Counted from the artefact's own `differential` only, once per binding rather than per element. Inherited bindings are not resolved, because the great majority of the artefacts held carry no snapshot.
6. **Fixed and pattern constraints.** `fixedUri`, `fixedCode`, `fixedCoding`, `fixedCodeableConcept`, `patternUri`, `patternCode`, `patternCoding` and `patternCodeableConcept`, counted separately from bindings.
7. **Family assignment.** By an ordered list of substrings tested against the lower-cased source path, first match winning; the list is written out at D-10.7. This shapes the presentation by family and not the corpus totals.
8. **Derivation.** Compare `baseDefinition` after removing any `|version` suffix: a parent named with a version is the same parent. A package boundary is the boundary between top-level source repositories or package archives. Ancestry is followed transitively, with cycle detection.

Five ways to miscount this corpus

Mixed serialisations. A JSON-only pass loses about a fifth of this corpus, because the notification packages and several record types ship XML. The three notification repositories alone hold 509 XML StructureDefinitions and the child examination booklet 214. An XML reader must also distinguish the StructureDefinition's own `type` and `kind` from the `type` inside a `context` element, and from a `mapping` block that may precede `kind`. A reader that takes the first `type`, or splits the header at the first `mapping`, silently reclassifies extensions as datatype profiles. The corresponding JSON is unambiguous, so a JSON-only test suite catches neither.

A version suffix is not a different artefact. `…/StructureDefinition/x|1.5` and `…/StructureDefinition/x` name the same parent. A graph built by exact string comparison books the first as unresolvable and therefore as lying outside the corpus. The effect is not marginal: it moves 110 derivations out of the reuse count and inverts the finding on cross-package version pinning.

Absence of a binding is not absence of a constraint. A profile can constrain an element through a fixed value, a pattern or slicing without declaring a value-set binding, and in this corpus the listed fixed and pattern constraints outnumber explicit bindings by nearly three to one.

Where a fixed value or a pattern sits on a coded element, it is **tighter** than a binding to a value set containing that code, because it admits one code rather than a set. A low binding count read as low rigour therefore measures the wrong thing.

How many of the 8,073 listed fixed and pattern constraints are of that kind is not established here, because they were not classified by path and datatype. The defensible statement is conditional: fixing is stricter where it applies to terminology, and the share to which it applies is unknown from this count.

A high `required` share is not by itself a quality score. `extensible` carries a real conformance obligation: a code from the value set wherever that set contains a suitable representation of the concept, and another code only where it does not. For a deliberately open concept space it may be the appropriate strength. A `required` share measures how much of the explicitly bound content is closed to codes outside the specified value sets. Whether that is appropriate must be assessed element by element, so the distribution is not a quality grade.

Build output is not a frozen archive. Repository build output can change without a package-version change. It must not be treated as equivalent to a published archive unless the two artefact sets have been compared directly.

What these figures will not carry

Conformance to a profile is not interoperability, and no count in this report can be read as evidence that it is. The precedent is D'Amore et al. (2014) on C-CDA, appraised with its discounts at Appendix D-11.1 of the Evidence Annex.

The consequence for this corpus runs in both directions. Profiles are the instrument that closes exactly the optionality that study measured, so a corpus with this much binding and fixing is doing the thing the earlier technology could not — which is an argument for profile depth and not for the standard as such. But the same finding warns against reading conformance to those profiles as evidence that the data mean the same thing at both ends. What a validator checks is what a validator checks.

These limits govern how the reported figures may be read. They do not alter the central findings: the size of the corpus, the predominance of `required` among explicitly declared bindings, the local structure of profile reuse, and the concentration of version pinning at package boundaries.

Method, and the use of a large language model. The literature search, the collation of sources, the drafting, and the consistency and readability checking across versions of this set were carried out with the assistance of a large language model. The counts reported here were produced by scripts written the same way, re-run under five tie-break rules, and recorded artefact by artefact in the audit tables at D-10.6. The editorial pass that set paragraph and sentence structure was made after the findings were fixed and did not change them. Three limits of the method belong on the record.

The reading boundary. The peer-reviewed publications were read in full. For much of the remaining literature a summary extract produced by a large language model formed the working basis, and it was that extract rather than the publication that was worked from. The extract also covers a number of peer-reviewed publications, so the boundary does not run along the peer-review line and cannot be described by that shorthand.

Appendix G of Document 4 records, source by source, which basis applies. The extract accompanies the set as Literature Summaries.md, so that a reader can see what the working basis actually said and can check the grades it proposes against the grades carried here.

The boundary is a record in one direction only. Every entry in that appendix other than "Full text" rests on something that can be checked independently of the author. The extract exists and names its sources, the artefact sits at the cited location, a conference presentation has no full text to read. The full-text entries rest on the author's account: they record that the complete publication rather than a summary was the working basis, not that the study was replicated or its data checked.

A further reservation belongs beside it. A number of ledger rows carry a basis of assessment that was reconstructed analytically from how the source is used in this text, rather than reported by the author. Correcting those against his own recollection is an outstanding task. Both asymmetries are stated rather than smoothed over.

A model does not weigh evidence, but it did propose. For the sources on which the summary extract proposes an evidence grade, the grade this document carries is in each case the grade proposed. Whether that agreement reflects influence or independent arrival at the same judgement cannot be determined after the fact, and no claim is made either way.

At version 3.16.0 every grade carried by a load-bearing source was re-derived from the rubric now stated at the head of Appendix G, without reference to what the extract had proposed. The outcome is recorded in the development log. **That re-derivation is not an independent second assessment and must not be read as one.** It was performed by the same kind of process that produced the proposals under question. It can therefore test the

grades against a written rule, but it cannot establish that the rule was applied by a mind uninfluenced by the earlier suggestion.

An independent second assessment by a qualified person remains outstanding, and is named here as an open item rather than quietly omitted. What can be asserted is narrower than the earlier formulation of this note. The grading rule, its application across the whole body of sources, the assessment and the conclusions are the author's, as is responsibility for any error that remains.

Transparency. The author is CEO of HL7 Deutschland, a Senior Expert and Community Event Manager at HL7 Europe, and a member of the ART-DECOR Expert Group. Those roles bear on a set of documents that recommends a FHIR-native default for new infrastructure, and they are stated for that reason. ***The assessment presented here is his personal professional view as an interoperability expert and does not represent an official position of HL7 Deutschland or HL7 Europe.***

The premise the recommendation rests on is stated in the open in Document 2 and in Document 3, so that it can be examined rather than inferred.