

Hierarchical Translation Architecture for a Client-Based International Patient Summary Viewer

Emmanuel Benoist

`emmanuel.benoist@bfh.ch`

Bern University of Applied Sciences

Alexander Kreutz

MIDATA Cooperative

Dominik Steiger

MIDATA Cooperative

Research Article

Keywords:

Posted Date: October 27th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10996485/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Hierarchical Translation Architecture for a Client-Based International Patient Summary Viewer

Emmanuel Benoist^{1*}, Alexander Kreutz² and Dominik Steiger²

^{1*}Department Engineering and Computer Science, Bern University of Applied Sciences, Quellgasse 21, Biel, 2501, Switzerland.

^{2*}MIDATA Cooperative, Pfingstweidstrasse, Zurich, 8005, Switzerland.

*Corresponding author(s). E-mail(s): emmanuel.benoist@bfh.ch;
Contributing authors: alexander.kreutz@midata.coop;
dominik.steiger@midata.coop;

Abstract

Cross-border and cross-institutional exchange of the International Patient Summary (IPS) requires that FHIR-encoded clinical content be presented to patients and clinicians in their preferred language, yet IPS documents are typically authored in the language of the producing system. This article presents a hierarchical translation architecture for a client-based IPS viewer that renders structured FHIR resources in the user's requested language without relying on backend document translation. Rather than treating the IPS as a narrative document to be machine-translated in bulk, the architecture exploits the structured nature of FHIR data, distinguishing static fields, coded terminology, and free-text narrative, and applies a quality-prioritized fallback chain comprising static application internationalization, FHIR multi-language extensions, cached application-level terminology translations, terminology-server lookups, machine translation services, and a transparent original-language fallback. The design is guided by three principles: preservation of the semantic integrity of coded data, prioritization of the highest-quality available translation source, and predictable, transparent fallback behavior. We report an evaluation of terminology-server coverage for LOINC, SNOMED CT, and ICD-10 across six languages (English, German, French, Italian, Dutch, and Estonian), which revealed substantial disparities in translation availability, and we discuss the resulting data-protection and authentication implications of querying external terminology and translation

services from a privacy-constrained client. The results indicate that a structured, client-side translation strategy can reconcile translation quality, semantic correctness, and data protection in distributed, multilingual eHealth ecosystems.

Keywords: International Patient Summary; FHIR; Multilingual Healthcare Information Systems; Terminology Translation; Client-Side Architecture; Cross-Border Health Data Exchange

1 Introduction

Cross-border healthcare, telemedicine, and the increasing mobility of patients require that health data remain understandable regardless of its source language. The International Patient Summary (IPS) [1, 2] provides a concise, standardized clinical summary intended for international exchange. Implementations typically rely on the HL7 Fast Healthcare Interoperability Resources (FHIR) standard [3] and its associated profiles, which enable semantic interoperability across producing systems.

However, FHIR resources combine several categories of content that place different demands on translation. Structured fields, such as status codes, take a small number of fixed values. `CodeableConcept` elements represent clinically meaningful concepts, such as symptoms, diagnoses, or treatments, using language-independent codes drawn from standard terminologies including SNOMED CT [4], ICD-10 [5], and LOINC [6]. Narrative strings are free text authored directly by clinical staff. Finally, FHIR field and resource names are themselves defined in English and must be rendered in the user’s language when displayed.

When such content is displayed in an IPS viewer, the language of the underlying data may not match the user’s preferred language. This mismatch is particularly relevant for a standard explicitly designed to support cross-border treatment, and it can be compounded in meshed national health infrastructures, such as the Swiss Digisanté initiative and similar distributed eHealth ecosystems, where different parts of a single document may originate from different source systems and languages.

This raises the question of how an IPS viewer can display IPS or other FHIR resources in a user’s requested language when the source data is provided in one or more different languages, while preserving translation quality, semantic correctness, and data protection. These properties are critical for medical data: patients and caregivers expect information to be precise and accurate, translation must not alter the clinical meaning of the data, and it must not leak health-related information to unauthorized actors.

To address this question, we designed and implemented a hierarchical translation architecture for a client-based IPS viewer. The design is guided by three principles. First, the architecture preserves the semantic integrity of coded data, ensuring that clinical codes and their meanings remain accurate throughout the translation process. Second, it prioritizes the highest-quality available translation source, implementing a fallback chain that favors authoritative sources over generic machine translation.

Third, it provides predictable fallback behavior, with deterministic and transparent handling of the cases in which no translation can be obtained.

The remainder of this article is organized as follows. Section 2 reviews related work on the translation of patient summaries. Section 3 describes the system context in which the translation architecture operates. Section 4 presents the hierarchical translation strategy. Section 5 discusses implementation considerations. Section 6 reports the results of our evaluation, and Section 7 discusses data protection and translation quality. Section 8 concludes the article and outlines future work.

2 Related Work

Solomou et al. [7] presented a mechanism to translate an entire IPS document in the backend, using the EU eHealth Digital Service Infrastructure (eHDSI) OpenNCP as a centralized translation service. This approach treats the IPS as a largely narrative document and applies machine translation at the document level. In contrast, the architecture presented in this article performs translation on the client and leverages structured FHIR semantics to maintain semantic correctness throughout the process. It preserves coding integrity by avoiding unnecessary translation of structured identifiers and clinical codes, which minimizes backend transformation requirements and allows translation to occur efficiently on the client without extensive server-side preprocessing or document manipulation.

3 System Context

3.1 IPS and FHIR

The IPS is defined as a constrained set of FHIR profiles based on the FHIR standard developed by HL7 International [3]. FHIR resources are structured and strongly typed, and often use `CodeableConcept` elements that reference external code systems. Appendix A shows a minimal IPS example: the `AllergyIntolerance` and `Condition` resources encode clinically coded concepts, a penicillin allergy and a hypertension diagnosis, respectively, using SNOMED CT, while the `MedicationStatement` resource combines a coded medication with a free-text narrative dosage instruction (“*50 mg, once daily, oral*”), illustrating the coexistence of coded and narrative content that motivates the translation strategy of Section 4.

FHIR data relevant to translation may contain:

- static code systems (e.g., status fields);
- large international terminologies;
- optional display values;
- optional language extensions.

3.2 Meshed Health Data Systems

In meshed health ecosystems, data are not centrally transformed for presentation but exchanged in standardized form across distributed actors. This architecture increases the importance of on-the-fly, context-aware translation within the viewing application

itself, since no single actor can be relied upon to have already translated the content into every language a downstream user might request.

4 Hierarchical Translation Strategy

When presenting structured IPS data, the following elements require translation:

- static structural elements;
- code-type fields with fixed value sets;
- `CodeableConcept` and `Coding` elements;
- free-text string elements.

The translation process follows a prioritized chain ordered by expected translation quality.

4.1 Level 0: Static UI and Structural Elements

Field names, resource names, and code-type fields with limited enumerations are translated using an application-level i18n file. These sets are limited in size and can be fully maintained for all supported languages, which ensures high-quality, deterministic translations.

4.2 Level 1: FHIR Multi-Language Extensions

FHIR provides extensions allowing multiple language labels within a resource. If such an extension exists for the requested language, it is used with the highest priority. This method guarantees semantic alignment, since the translation is supplied by the data source itself.

4.3 Level 2: Application i18n for Frequent Codes

For codes that are frequently used within the application's use case, translations may be stored directly in the application's i18n file. This is suitable for small, frequently displayed subsets and provides high performance, since lookups on external services are avoided.

4.4 Level 3: Terminology Server Integration

If available, a FHIR terminology server is queried to obtain localized display values. Terminology servers can resolve system-and-code combinations to language-specific displays using the FHIR `$lookup` operation. This approach ensures terminology-consistent translations; however, the availability of translated terminologies varies greatly from language to language. This step ensures high semantic accuracy and is preferable to generic machine translation.

4.5 Level 4: Translation Service

If no display value in the requested language is available, a translation service may be used to translate the concept from another language: if a display value exists in the

source language, the translation service is applied to it directly; if no display value exists but the terminology server supports the relevant code system, the terminology server is first queried for the English display value, which the translation service then translates into the target language.

This method provides broader coverage but lower semantic reliability than terminology-managed translations. Modern browsers such as Google Chrome, Microsoft Edge, and Opera provide client-side translation APIs in which AI-based translation models can be loaded locally, reducing data transfer to external translation providers.

4.6 Level 5: Fallback

If all prior methods fail, the original display value is shown, resulting in content displayed in the wrong language; if no display value exists, only the system and code are shown. This provides the poorest user experience but guarantees at least transparency about the limits of translation.

For non-coded string-type fields, a similar hierarchical chain is used, excluding the terminology-server lookup step:

FHIR language extension $\rightarrow$ application i18n $\rightarrow$ translation service $\rightarrow$ original string

5 Implementation Considerations

The translation engine is implemented as a modular service within the IPS viewer frontend. Both the terminology lookup and the translation service are configurable. The terminology lookup is exposed as a function whose input is a JSON object containing the parameters for a terminology-server call, namely the code, the system, and the target language; the function may return a direct translation, a terminology-server response, or a promise for either. The translation service is likewise exposed as a function whose input is the string to be translated together with the source and target languages; it may return a direct translation or a promise for the translation. All results are cached so that terminology and translation servers are not flooded with repeated requests when a bundle contains many resources with similar content.

6 Results

6.1 Possible Data Exposure

When contacting terminology or translation servers, coded values or display values are frequently transmitted via GET requests over the Internet, so potentially sensitive information can end up in request URLs. URLs are often recorded in server logs, and combined with the requester's IP address or the exact request time, such log entries could in principle be used to re-identify the associated patient. If the terminology and translation servers are operated by the same authority that stores the IPS data, this additional risk may be acceptable within the existing governance model.

6.2 Authentication

Authentication tokens must be handled securely within the client application. Terminology and translation servers often require static API keys, but client applications can be reverse-engineered and any embedded static API key extracted. Services that do not require authentication should therefore be preferred where available; when authentication is unavoidable, a short-lived API key should be supplied through a secure mechanism, such as single sign-on (SSO).

6.3 Availability of Translations

During the evaluation of the implementation, we sought translations for LOINC, SNOMED CT, and ICD-10 in English, German, French, Italian, Dutch, and Estonian. At the time of evaluation, we did not find a freely available terminology server capable of handling all of these languages, and we therefore set up our own FHIR terminology server populated with the available language extensions for SNOMED CT and LOINC. The number of translated codes available differed substantially across languages: as expected, English translations were available for essentially all codes, German coverage was also satisfactory, whereas Estonian had the smallest number of translated codes.

7 Discussion

7.1 Ensuring Data Protection

Data protection remains a critical concern when implementing terminology and translation services in client-based architectures. Several complementary approaches can strengthen the security posture of the IPS viewer. Hosting terminology and translation services within the same trusted domain as the IPS data eliminates the risk of sensitive information being transmitted to external providers. Alternatively, leveraging on-device translation APIs, such as those available in modern browsers, allows translation to occur locally without any network exposure of clinical content. When network-based services are necessary, token scoping and short-lived credentials can minimize the window of exposure should credentials be compromised. Comprehensive logging and audit controls provide accountability and enable detection of suspicious access patterns. Additionally, the system should be designed to avoid transmitting complete clinical context, sending only the minimal information required for a given translation operation.

7.2 Ensuring Translation Quality

Maintaining high translation quality across the system requires a multi-faceted approach. The architecture prioritizes terminology-managed translations as the primary source, since these represent the most semantically accurate and authoritative translations available. To supplement this, maintaining curated i18n dictionaries for common use cases ensures consistent, high-quality translations for frequently encountered terms and concepts. Continuous validation by domain experts, particularly

clinicians and health information specialists, is essential to catch translation errors that could have clinical implications. Finally, monitoring the frequency of fallback scenarios provides valuable insight into system performance and identifies gaps in coverage where additional translation sources or manual curation may be needed.

8 Conclusion and Future Work

This article presented a hierarchical, client-based translation architecture for IPS viewers capable of rendering structured FHIR resources in the user’s requested language. Unlike full-document backend translation approaches, this architecture leverages FHIR semantics to maintain semantic correctness, preserves coding integrity to ensure clinical accuracy, and prioritizes authoritative translation sources through a quality-driven fallback chain. It provides deterministic fallback behavior, giving transparent and predictable handling of edge cases, and it integrates cleanly into meshed digital health infrastructures without requiring centralized backend transformation.

Several directions remain open for future work. Empirical evaluation of translation quality through clinical validation would provide valuable evidence of the approach’s effectiveness in real-world settings. Performance benchmarking across different terminology servers and translation services would help optimize system responsiveness. A comprehensive, formal privacy risk analysis would further strengthen the security considerations and governance frameworks discussed in Section 7. Finally, integration with national terminology services would enhance the availability and reliability of terminology-based translations across diverse healthcare systems, in particular for languages such as Estonian, for which coverage was found to be limited in Section 6.

The proposed architecture demonstrates that high-quality multilingual rendering of IPS/FHIR resources can be achieved through a structured, hierarchical, and semantically aware translation strategy implemented at the client level, offering a viable path toward multilingual healthcare information systems in distributed eHealth ecosystems.

Declarations

Funding

The Swiss partners’ contribution to this work was carried out within the Horizon Europe project AIDAVA and was funded by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract No. 22.00093, Switzerland not being an EU Member State and therefore funding its participants in this Horizon Europe project directly rather than through the European Commission.

The work of the other project members was funded by the European Union’s Horizon Europe research and innovation programme under grant agreement No. 101057062.

Competing Interests

The authors declare that they have no competing interests.

Data Availability Statement

The data used for testing the software are originating from different FHIR IPS examples data bases. The most important is the git repository of the FHIR IPS project ¹. We tested also with data provided by the Austrian ELGA project² available in CDA format (i.e. xml) transformed into JSON using the official mapping from HL7 ³. We have also manually manipulated (mainly translated part of it) some of those example files for our testing purpose.

Ethics Approval

Not applicable. This article does not report a study involving human participants; the terminology-coverage evaluation reported in Section 6 concerns publicly available terminology servers and does not involve patient data.

Consent to Participate

Not applicable.

Consent for Publication

Not applicable; the manuscript does not contain any individual person's identifiable data. The patient data shown in Appendix A are entirely fictional and are used solely to illustrate the structure of an IPS document.

Author Contributions

A.K. was responsible for the software implementation of the translation architecture (Software). E.B. and D.S. contributed the design of the hierarchical translation strategy (Conceptualization, Methodology). E.B. led the writing of the manuscript (Writing – Original Draft), and A.K. contributed to parts of the writing (Writing – Review & Editing). All authors read and approved the final manuscript.

Appendix A Example of an IPS File

The following listing shows a minimal, fictional example of an IPS bundle in JSON format, referenced in Section 3 to illustrate the coexistence of coded and narrative content that requires translation.

```
{
  "resourceType": "Bundle",
  "id": "ips-example",
  "meta": {
    "profile": [
      "http://hl7.org/fhir/uv/ips/StructureDefinition/Composition-uv-ips"
```

¹<https://github.com/HL7/fhir-ips-ri>

²https://www.hl7.org/implement/standards/product.brief.cfm?product_id=7

³<https://build.fhir.org/ig/HL7/ccda-on-fhir/en/>

```

    ]
  },
  "type": "document",
  "entry": [
    {
      "fullUrl": "urn:uuid:1",
      "resource": {
        "resourceType": "Composition",
        "id": "ips-composition",
        "status": "final",
        "type": {
          "coding": [
            {
              "system": "http://loinc.org",
              "code": "60591-5",
              "display": "International Patient Summary"
            }
          ]
        }
      }
    },
    {
      "subject": {
        "reference": "urn:uuid:2"
      },
      "date": "2026-03-17T12:00:00+01:00",
      "author": [
        {
          "reference": "Practitioner/author-practitioner"
        }
      ],
      "title": "International Patient Summary for John Doe",
      "section": [
        {
          "title": "Allergies",
          "entry": [
            {
              "reference": "urn:uuid:3"
            }
          ]
        },
        {
          "title": "Medications",
          "entry": [
            {
              "reference": "urn:uuid:4"
            }
          ]
        }
      ]
    }
  ]

```

```

    },
    {
      "title": "Medical History",
      "entry": [
        {
          "reference": "urn:uuid:5"
        }
      ]
    }
  ]
},
{
  "fullUrl": "urn:uuid:2",
  "resource": {
    "resourceType": "Patient",
    "id": "patient-example",
    "name": [
      {
        "family": "Doe",
        "given": ["John"]
      }
    ],
    "gender": "male",
    "birthDate": "1970-01-01"
  }
},
{
  "fullUrl": "urn:uuid:3",
  "resource": {
    "resourceType": "AllergyIntolerance",
    "id": "allergy-example",
    "clinicalStatus": {
      "coding": [
        {
          "system": "http://terminology.hl7.org/CodeSystem/allergyintoleran",
          "code": "active"
        }
      ]
    }
  },
  "code": {
    "coding": [
      {
        "system": "http://snomed.info/sct",
        "code": "247472004",

```

```

        "display": "Allergy to penicillin"
      }
    ]
  },
  "patient": {
    "reference": "urn:uuid:2"
  }
}
},
{
  "fullUrl": "urn:uuid:4",
  "resource": {
    "resourceType": "MedicationStatement",
    "id": "medication-example",
    "status": "active",
    "medicationCodeableConcept": {
      "coding": [
        {
          "system": "http://www.whocc.no/atc",
          "code": "C07AB02",
          "display": "Atenolol"
        }
      ]
    },
    "subject": {
      "reference": "urn:uuid:2"
    },
    "dosage": [
      {
        "text": "50 mg, once daily, oral"
      }
    ]
  }
}
},
{
  "fullUrl": "urn:uuid:5",
  "resource": {
    "resourceType": "Condition",
    "id": "condition-example",
    "clinicalStatus": {
      "coding": [
        {
          "system": "http://terminology.hl7.org/CodeSystem/condition-clinical",
          "code": "active"
        }
      ]
    }
  }
}
}

```

```

    ]
  },
  "code": {
    "coding": [
      {
        "system": "http://snomed.info/sct",
        "code": "38341003",
        "display": "Hypertension"
      }
    ]
  },
  "subject": {
    "reference": "urn:uuid:2"
  }
}
]
}

```

References

- [1] Hausam, R., D’Amore, J., Cangilioli, G.: International Patient Summary Implementation Guide, Version 2.0.0 edn. HL7, (2025). HL7
- [2] International Organization for Standardization: Health Informatics – Structured Terminology for Clinical Purposes. <https://www.iso.org/standard/84639.html>
- [3] HL7: HL7 FHIR Standard & Reference Guide, Version 5.0.0 edn. (2023). HL7
- [4] SNOMED International: SNOMED CT Systematized Nomenclature of Medicine – Clinical Terms. Accessed: May 19, 2026 (2026). <https://www.snomed.org/>
- [5] World Health Organization: Icd-10: International statistical classification of diseases and related health problems: Tenth revision. Technical report, Geneva, Switzerland (2004). <https://iris.who.int/handle/10665/42980>
- [6] Forrey, A., McDonald, C., DeMoor, G., Huff, S., Leavelle, D., Leland, D., Fiers, T., Charles, L., Griffin, B., Stalling, F., *et al.*: Logical observation identifier names and codes (loinc) database: a public use set of codes and names for electronic reporting of clinical laboratory test results. *Clinical Chemistry* **42**(1), 81–90 (1996)
- [7] Solomou, T., Canciu, I.-C., Christodoulou, M., Savva, P., Yiasemi, C., Antoniou, Z., Constantinou, I., Schizas, C.N., Pattichis, C.S.: Empowering citizens through translated patient summary access and sharing in digital health ecosystems. In: *Digital Health and Informatics Innovations for Sustainable Health Care Systems*.

Studies in Health Technology and Informatics, vol. 316, pp. 497–501. IOS Press, Amsterdam, The Netherlands (2024). <https://doi.org/10.3233/SHTI240457>