


An Introduction to FHIR for Healthcare Interoperability

By IntuitionLabs • 8/26/2025 • 50 min read

[fhir](#)[healthcare interoperability](#)[hl7](#)[health data standard](#)[api](#)[rest](#)[health informatics](#)

EXPLORING FHIR IN HEALTHCARE





FHIR: Fast Healthcare Interoperability Resources – A Comprehensive Report

Introduction

Fast Healthcare Interoperability Resources (FHIR) is a modern health data standard created by Health Level Seven International (HL7) to facilitate **interoperability** – the seamless exchange of clinical information across systems. FHIR was born out of the shortcomings of earlier HL7 standards and the need to leverage web technologies. In 2011, Australian HL7 architect Grahame Grieve proposed a new approach called “Resources for Healthcare” (RFH), which was designed around **web principles** (HTTP, XML, REST) [medinform.jmir.org](https://www.medinform.jmir.org). This evolved into FHIR, incorporating the best of HL7 version 2 and 3 while using modern web standards [medinform.jmir.org](https://www.medinform.jmir.org). The primary **purpose** of FHIR is to reduce integration complexity without sacrificing **information integrity** [medinform.jmir.org](https://www.medinform.jmir.org) [medinform.jmir.org](https://www.medinform.jmir.org). Instead of monolithic messages or documents, FHIR exposes small, self-contained **resources** (like *Patient*, *Observation*, *Medication*) through **standard web APIs**, allowing granular “on-demand” access to data that was not possible with legacy standards [medinform.jmir.org](https://www.medinform.jmir.org) [medinform.jmir.org](https://www.medinform.jmir.org).

FHIR quickly gained momentum in healthcare IT. It is **free, open, and relatively easy to implement**, making it attractive to both vendors and providers digital.nhs.uk. Major technology companies and EHR vendors have endorsed FHIR – for example, in 2018 a coalition of Big Tech firms (including Microsoft, IBM, Google, Amazon) pledged support for FHIR as an emerging standard for **health data exchange** [medinform.jmir.org](https://www.medinform.jmir.org). With the parallel rise of mobile apps and patient-centered care, FHIR emerged at the right time to enable secure data sharing across organizational boundaries. Today, FHIR is widely regarded as the **global industry standard** for healthcare data exchange, forming the backbone of many national interoperability initiatives digital.nhs.uk.

FHIR Technical Architecture and Core Concepts

Modular “Resources”: The fundamental building blocks of FHIR are *resources* – these are small, logically self-contained data models for common health concepts (patients, encounters, lab results, medications, etc.). Each resource has a clear scope and identity (addressable by a unique URL) and carries meaningful clinical data [medinform.jmir.org](https://www.medinform.jmir.org) [medinform.jmir.org](https://www.medinform.jmir.org). Over 150 resource types have been defined (as of FHIR Release 4), organized into categories such as *Administrative* (e.g. Patient, Provider, Organization), *Clinical* (e.g. Condition, Observation, Procedure), *Financial* (e.g. Claim, PaymentNotice), *Infrastructure* (e.g. Bundle, MessageHeader),

and *Workflow* (e.g. Appointment, Task) medinform.jmir.org. By combining resources, complex health information (like a complete encounter or a referral note) can be represented. Each resource also contains a human-readable narrative alongside structured data, ensuring data can be understood both by machines and clinicians nlm.nih.gov nlm.nih.gov.

RESTful APIs: FHIR's design is tightly coupled with web API principles. A FHIR **server** is essentially a RESTful web service that exposes endpoints for each resource type. Clients (e.g. EHRs, apps) interact with these endpoints using standard HTTP verbs – GET to read data, POST to create, PUT/PATCH to update, and DELETE to remove data nlm.nih.gov nlm.nih.gov. Each resource is accessible via a URL (for example, `GET /Patient/123` fetches the patient with ID 123). This RESTful approach aligns with how the web works, lowering the barrier for developers. It allows quick, on-demand queries (retrieving just the needed data) as opposed to receiving entire data dumps. FHIR servers can also support search parameters (e.g. find patients by name or find lab results by date) and batch operations. In addition, the FHIR specification accommodates other paradigms like **Messaging** and **Documents** for cases where those make sense, but the core focus is on RESTful API exchange nlm.nih.gov nlm.nih.gov.

Data Formats (JSON, XML, RDF): FHIR resources can be represented in multiple serialization formats. The most commonly used formats are **JSON** (JavaScript Object Notation) and **XML**, both of which are widely supported in programming and allow nested, hierarchical data. FHIR also defines an RDF/Turtle format, aligning with Semantic Web standards (Resource Description Framework) for advanced use cases in linked data and ontology mapping nlm.nih.gov. In practice, JSON has become the preferred format for most developers due to its conciseness and compatibility with web/mobile technologies rhapsody.health. XML is used in certain enterprise contexts or when integrating with systems that already rely on XML. RDF/Turtle is primarily used in research or specialized semantic interoperability projects. By offering multiple formats, FHIR ensures flexibility: the **information model** is separate from the wire format. Internally, each resource has a defined structure (fields, cardinalities, data types, etc.), and those can be rendered as JSON, XML, or RDF as needed nlm.nih.gov nlm.nih.gov. All formats convey the same content – for example, a Patient resource in JSON or in XML will contain the same demographics, identifiers, and so forth.

Terminologies and Semantic Interoperability: FHIR was built with **semantic interoperability** in mind. Rather than inventing new code systems, FHIR resources make extensive use of existing medical terminologies. For instance, a FHIR Observation for a lab test can carry LOINC codes for the test name and SNOMED CT codes for qualitative results. The FHIR data model includes elements like `Observation.code` or `Condition.code` specifically to hold such standardized codes. This use of **existing vocabularies** (SNOMED, LOINC, ICD, RxNorm, etc.) is a key aspect of FHIR's design nlm.nih.gov. It means that two systems exchanging a FHIR resource can not only parse the structure but also understand the *meaning* of the data, as long as they reference the same coding standards (this is what enables true semantic interoperability). FHIR's **Terminology services** and resources (like *CodeSystem*, *ValueSet*, *ConceptMap*) further support managing and translating codes. However, it is



important to note that while FHIR provides the framework for semantic interoperability, achieving it requires consistent use of codes and mappings – a topic we will revisit in the *Challenges* section.

Profiles and Extensions: The base FHIR specification provides a default schema for each resource, but healthcare is diverse and often requires specialization. FHIR addresses this via **Profiles** and **Extensions**. A **profile** is a set of constraints or requirements on a resource for a particular use case or region (for example, the US Core Patient profile mandates certain data fields for U.S. workflows). Profiles can make elements required, add terminology bindings, or restrict cardinalities. **Extensions** allow adding new data fields to a resource when the standard ones are insufficient, while maintaining interoperability (unknown extensions are simply ignored by systems that don't recognize them). This balance between **rigid standardization and flexibility** is a hallmark of FHIR's architecture [kodjin.com](https://www.kodjin.com) [kodjin.com](https://www.kodjin.com). It lets FHIR serve as a general purpose standard (the "80%" common use cases) with the ability to handle the other 20% via site-specific or country-specific extensions. Extensions are governed by globally unique identifiers and definitions, so any system encountering an extension can look up its meaning. Together, profiles and extensions support a consistent core while accommodating variability – though they also introduce complexity in ensuring conformance (discussed later).

In summary, FHIR's technical architecture marries **modern web API design** (REST/HTTP) with **modular clinical data models** (resources), rendered in developer-friendly formats (JSON/XML). This combination enables "plug and play" interoperability: a lightweight, scalable approach where any app or system that knows how to call a REST API and handle JSON can exchange healthcare data. The next sections compare this approach to legacy HL7 standards and explore how FHIR is being implemented in real-world scenarios.

FHIR vs Legacy HL7 Standards (v2, v3, CDA)

FHIR did not emerge in a vacuum – it is the successor to prior HL7 standards that have served healthcare for decades. The most prominent legacy standards are **HL7 Version 2 (v2)** and **HL7 Clinical Document Architecture (CDA)** (which is an XML document standard from HL7 Version 3). Understanding their characteristics highlights FHIR's improvements:



- **HL7 Version 2 (v2):** First released in the late 1980s, HL7 v2 is a text-based messaging standard. It encodes data in delimited segments (e.g. `PID|...` for patient info) sent typically via MLLP or other transport within or between systems. HL7 v2 became **widely adopted in hospitals** for interfacing labs, radiology, ADT systems, etc., and remains in use today. Its strengths are simplicity and ubiquity – it covers many clinical workflows and is relatively lightweight to parse. However, HL7 v2 interfaces tend to be highly customized for each installation, leading to **inconsistencies across implementations** [medinform.jmir.org](https://www.medinform.jmir.org). There was no strict universal data model, so one hospital's "ADT^A01" admit message might not be exactly the same as another's. The standard also lacks a built-in mechanism for fine-grained data queries – it pushes whole messages on events, rather than allowing on-demand retrieval of a single data element. Security and transport were not standardized in v2 (usually left to the implementer), and the encoding (pipe-delimited ASCII) is not web-friendly. Overall, HL7 v2 requires significant mapping effort between systems and has a **steep learning curve** for new interfaces rhapsody.health rhapsody.health. It functions well for point-to-point data pushes within an enterprise, but is less suited to patient-centric, cross-organization data sharing using modern apps.
- **HL7 Version 3 and CDA:** In the 2000s, HL7 attempted a more rigorous approach with Version 3, introducing a formal **Reference Information Model (RIM)** and XML-based encoding. One of the products of HL7 v3 is **CDA (Clinical Document Architecture)**, an XML standard for clinical documents (like discharge summaries or continuity of care records). CDA became famous through implementations like the CCD (Continuity of Care Document) used in U.S. Meaningful Use programs rhapsody.health. CDA provides a structured document with a human-readable narrative and machine-readable entries, which was a step toward interoperability of summaries. However, HL7 v3/CDA suffered from **excessive complexity**. The RIM modeling approach was academically sound but difficult to implement – it resulted in verbose XML and required mastering a large abstract model. Consequently, v3 adoption was limited and not backward-compatible with v2 rhapsody.health. CDA documents, while useful for human communication, still posed challenges for **machine-to-machine data exchange** beyond the document level – extracting discrete data from narrative-heavy CDA can be hard, and many workflows don't align neatly with document boundaries rhapsody.health. Additionally, implementers found it hard to extend or adapt CDA without breaking interoperability, and the learning curve remained steep rhapsody.health.

FHIR's Advances: FHIR was explicitly designed to **address the limitations** of these earlier standards while retaining their good aspects rhapsody.health rhapsody.health. It combines the broad clinical scope of HL7 v2 and v3 with a far more **flexible, web-compatible architecture**. Key differences include:

- FHIR is based on modular **resources** rather than monolithic messages or documents. Each resource is like a small "data packet" (analogous to a v2 segment, but self-contained) that can be independently accessed or updated rhapsody.health rhapsody.health. This granularity was missing in older standards – for example, in HL7 v2 you couldn't easily just fetch one lab result out of a batch message, whereas in FHIR you can GET a specific Observation resource.
- FHIR uses **RESTful APIs and common web technologies** (HTTP, SSL/TLS, OAuth for auth) which are familiar to developers, as opposed to custom MLLP sockets or SOAP-based web services. This significantly lowers the entry barrier and allows use of standard tools and cloud services rhapsody.health.



- FHIR supports data in **JSON or XML** with **mandatory human-readable sections**, making it both machine-parseable and human-inspectable [rhapsody.health](#) [rhapsody.health](#). HL7 v2 messages are somewhat human-readable but cryptic; CDA is human-readable but heavy. FHIR strikes a middle ground with concise JSON and optional narrative.
- FHIR was built with **extensibility and versioning** in mind. It has a robust extension framework so that adding custom fields does not break other systems (extensions are self-described). It also established rules for backward compatibility once the standard reached a normative stage (FHIR R4) [rhapsody.health](#). In contrast, HL7 v3 was a clean break from v2 (no backward compatibility), and v2 itself saw many variations.
- In terms of **workflow**, HL7 v2 is event-driven (you get data when an event triggers a message), whereas FHIR allows both event-driven (subscriptions, messages) *and* query-driven (REST calls) patterns. This makes FHIR more flexible: clients can pull the latest data when needed, rather than relying solely on pushes.
- **Interoperability scope**: HL7 v2 and CDA were largely syntactic interoperability standards – they could get data from A to B, but often required out-of-band agreements to ensure semantic alignment. FHIR, by design, places greater emphasis on standardized terminologies and unambiguous data definitions to achieve **semantic interoperability** [rhapsody.health](#). It also covers “FHIR workflow” (a set of resources like Task, Appointment, etc.) to standardize common healthcare workflows, something outside the scope of HL7 v2 messages.

It's worth noting that FHIR does **not completely replace** HL7 v2 and CDA overnight – those standards are deeply entrenched. In fact, as of 2021, HL7 v2 was still used in some form by 95% of U.S. healthcare organizations [medinform.jmir.org](#). Many interfaces (e.g. lab instrument feeds or ADT feeds) run perfectly well on v2 and will not be rewritten immediately. FHIR is often adopted **alongside** legacy standards, bridging gaps for new use cases (especially patient-facing apps, data exchange across organizations, analytics, etc.) [medinform.jmir.org](#). Over time, we expect FHIR's domain to expand as new integrations prefer FHIR, while older standards gradually phase out or are encapsulated by FHIR-based APIs (for example, an integration engine might convert incoming v2 messages into FHIR resources). HL7's position is that FHIR is “**the next generation**” built on lessons from v2 and v3 [medinform.jmir.org](#). In practice, a hybrid environment exists today: FHIR is transforming interoperability for modern health IT, but HL7 v2/CDA remain relevant for certain entrenched workflows. Organizations must plan for mapping and co-existence during this transition period.

Use Cases and Implementations of FHIR

FHIR's flexibility and broad industry support have led to **diverse implementation scenarios**. Below we explore several major use cases, illustrating how FHIR is applied in practice:



- **Integration with Electronic Health Records (EHRs):** One of FHIR's first big use cases is exposing data from EHR systems in a standardized way. Major EHR vendors such as Epic, Cerner, Allscripts, and others have implemented FHIR R4 APIs to allow external applications and partners to securely access EHR data (with proper authorization). For example, a hospital's EHR can provide a FHIR endpoint where a third-party care management app can fetch a patient's allergies or medications in real time (instead of using a custom HL7 v2 feed). In the U.S., this use case has been supercharged by regulation – under the 21st Century Cures Act rules, certified EHRs **must** provide patients with API access to their records via FHIR (by 2022) rhapsody.health. This has essentially made FHIR the default language for EHR data export. A concrete example is Apple's **Health Records** app for iPhone, which uses a FHIR API (Apple's system supports FHIR DSTU2 and R4) to connect to hospital EHRs and download patient records into the Health app support.apple.com. Many health systems signed on to enable this feature, meaning an iPhone user can, after authenticating, pull their clinical data (problems, lab results, immunizations, etc.) from multiple provider portals – all thanks to FHIR standardizing that API. Similarly, EHR vendors have internal "app stores" where third-party SMART on FHIR apps can be launched within the EHR, showing how FHIR enables an **ecosystem** of apps around EHR platforms. This integration improves workflows (e.g., a clinician can invoke a specialized app that uses FHIR to grab needed data from the EHR and perform some advanced function, then write results back via FHIR).
- **Mobile Health and Patient-Facing Apps:** FHIR is a catalyst for digital health innovation on smartphones and web apps. With FHIR, a small startup can develop an app that works across many hospitals/EHRs, rather than writing custom interfaces for each. **SMART on FHIR** (Substitutable Medical Apps, Reusable Technology) is an open platform that leverages FHIR and OAuth2 to let apps run in or alongside EHR systems medinform.jmir.org. For instance, a diabetes management app can use SMART on FHIR to plug into a patient's record, pulling blood glucose lab results and pushing care plan updates. Because FHIR uses standard web tech and data formats, even individual developers can create **mobile health apps** that were infeasible in the old HL7 world. We are seeing apps for medication reminders, chronic disease tracking, wellness and fitness integration, and more, all using FHIR under the hood to talk to provider systems. Mobile apps can also contribute data back – e.g., a patient could collect blood pressure readings at home and a FHIR-enabled app could POST these as Observation resources to the clinician's system. The **SMART on FHIR** framework provides the security and sandbox for these apps (with patients' consent), while FHIR provides the data access. This synergy was explicitly aimed at breaking down the "walled garden" of EHRs medinform.jmir.org. As a result, numerous **patient engagement** apps and provider-facing mobile tools now integrate via FHIR APIs. Regulatory drivers like the Cures Act and CMS's rules (noted below) are also spurring a wave of consumer apps that aggregate health data for patients via FHIR.



- **Public Health Data Exchange:** Public health agencies historically rely on HL7 v2 (for lab reporting, case reports, immunization registry updates, etc.), but FHIR is increasingly being adopted to modernize and broaden public health connectivity. One advantage is that a single **FHIR server** can potentially service multiple data needs simultaneously. For example, instead of a clinic maintaining separate HL7 v2 interfaces for sending immunization updates, disease case reports, and lab results to various public health programs, the clinic could expose a FHIR API from which authorized public health systems query whatever data they need [cdc.gov](https://www.cdc.gov). FHIR's query flexibility means public health can "pull" data on demand (during an outbreak, for instance) in addition to receiving pushes. The CDC notes that unlike v2 messages, FHIR data is **human-readable JSON/XML**, which lowers the technical barrier for analysts and simplifies troubleshooting [cdc.gov](https://www.cdc.gov) [cdc.gov](https://www.cdc.gov). A practical example is the **electronic case reporting (eCR)** during COVID-19: HL7 developed a FHIR-based eCR standard to automatically transmit case data from EHRs to public health authorities. Another area is immunizations – there are FHIR interfaces now for querying and updating immunization registries (an alternative to older v2 VXU/QBP messages). Public health agencies also benefit from FHIR's ability to support new data elements rapidly; if a new condition or data requirement emerges, a FHIR API can often accommodate it via extensions or new resource use, without having to wait for an entire new v2 message profile [cdc.gov](https://www.cdc.gov) [cdc.gov](https://www.cdc.gov). Moreover, U.S. federal rules encouraging the use of FHIR and **USCDI** (core data set) mean that public health can potentially collect a wide range of standardized data through one pipe rather than many – a "self-service buffet" of data as the CDC describes [cdc.gov](https://www.cdc.gov) [cdc.gov](https://www.cdc.gov). That said, transitioning public health systems to FHIR is an ongoing process; it requires new infrastructure and training, but pilot projects (like CDC's **FHIR-based Digital Bridge** for case reporting, or the SANER project for situational awareness) have shown promising results.
- **Population Health, Research, and Analytics:** Another implementation area is using FHIR for **bulk data export** and data analytics integration. HL7 has a defined "Bulk Data" specification (FHIR Bulk Data Access, also known as Flat FHIR or NDJSON format) which allows exporting large cohorts of data (e.g., all patient records for the last month, or all immunization records for a city) in a single job. This is being leveraged in research networks and quality reporting. For example, CMS's **Quality Payment Program** APIs and some research collaboratives use FHIR bulk data to gather de-identified datasets from EHRs. The **NIH** and academic centers are exploring FHIR for integrating clinical data with research databases – FHIR's standard structure makes it easier to write data pipelines as opposed to parsing many custom formats [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov) [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov). On the clinical research side, FHIR is being used to streamline clinical trial data capture (e.g., via the Vulcan project and others mentioned below). *Precision medicine* efforts (like HL7's FHIR Genomics) use FHIR to represent genomic test results in a standard way, facilitating data sharing for research while linking it to clinical context. FHIR's use of web APIs even allows real-time or near-real-time data aggregation for **public health surveillance** or **population health dashboards**. For instance, a health information exchange could use FHIR queries to continuously pull summary metrics (like number of flu cases) from hospital systems in its region.

- **Payer and Health Insurance Use Cases:** Outside of providers, health insurers and payers are also adopting FHIR. The **HL7 Da Vinci** project is an industry consortium creating FHIR implementation guides for common payer-provider data exchanges (e.g., coverage eligibility, prior authorization, care gaps). Additionally, the U.S. CMS **Interoperability and Patient Access Rule** requires health plans to offer patients API access to their claims and coverage data via FHIR. One implementation is CMS's **Blue Button 2.0** API which uses FHIR to give Medicare beneficiaries access to their claims history rhapsody.health. Payers are using FHIR for things like provider directories, formulary information, and sharing clinical data for care management. This is a notable expansion of FHIR beyond just clinical EHR data – it's creating a more **unified health data ecosystem** where claims, clinical, and even patient-generated data can all flow using the same standard.

These examples only scratch the surface. Other notable FHIR-based initiatives include **CDS Hooks** (providing real-time decision support by calling FHIR services), the **Gravity Project** (social determinants of health data standards on FHIR), and various national health information exchanges switching to FHIR. The common theme is that FHIR's **interoperability building blocks** can be assembled to fit many contexts – from a solo practitioner's patient portal to a national health data network.

Challenges and Considerations in FHIR Implementation

While FHIR offers significant improvements, implementing it at scale comes with several **challenges**. Healthcare IT leaders and developers must navigate these to fully realize FHIR's benefits:

- **Versioning and Evolution:** FHIR has undergone multiple version iterations (DSTU1, DSTU2, STU3, R4, R5, etc.), and not all systems upgrade simultaneously. This means there can be **version mismatch** issues – for example, an app built for FHIR DSTU2 (from 2015) might not work against a server that supports the normative FHIR R4 (2019) without adjustments. The good news is that with FHIR **Release 4 (R4)**, many key parts of the standard became "normative," guaranteeing backward compatibility for future versions rhapsody.health. Still, challenges remain in managing versions: vendors and healthcare organizations must periodically update their FHIR implementations to stay current, which can be non-trivial. Tooling is improving (e.g. version translation libraries), but as of 2025 we have a mix of FHIR STU3, R4, and some R5 in use. **Profiling** adds another versioning layer – e.g. the US Core profiles might get updated (e.g., from v3.1.1 to v4.0.0 to align with USCDI data elements), requiring adjustments by implementers. To mitigate issues, clear version negotiation and **capability statements** (a feature of FHIR servers that declares what versions and features they support) are important. In summary, organizations need a **strategy for version control** and backwards compatibility. Many are choosing to stick to one version (R4) for a stable period. Version churn is a natural byproduct of an evolving standard, but it requires governance to ensure all parties in an exchange agree on a common version or can translate between versions.



- **Conformance and Testing:** Ensuring that different systems' FHIR implementations truly interoperate requires rigorous **conformance testing**. The base FHIR spec leaves some optionality (and profiles add constraints), so two systems claiming "FHIR R4 support" might still have subtle differences. To address this, the FHIR community and governments have invested in testing frameworks. HL7 provides official **validation tools** to check if a given resource instance is structurally and terminologically correct fhir.org fhir.org. For example, there is a FHIR **Validator** tool that can validate an Observation resource against the spec and even against a profile (like US Core Observation) fhir.org. Beyond validating data, there are **suite tests** for FHIR APIs: the **Inferno** test tool (developed by ONC) tests an EHR's FHIR server for compliance with the ONC Certification criteria (including authentication, correct support of resource queries, etc.). Other services like AEGIS **Touchstone** and MITRE **Crucible** allow organizations to run their FHIR servers through a battery of test scripts fhir.org fhir.org. Despite these, true plug-and-play often requires coordinating on **implementation guides** and **profiles**. Lack of adherence to profiles or inconsistent use of codes can result in partial interoperability (e.g., two systems might both use FHIR but one doesn't populate a data field that the other expects). Conformance testing is thus a **continuous process**: as FHIR APIs are deployed, they need to be monitored and tested with trading partners. In essence, "FHIR compliance" is not just passing one test – it's an ongoing commitment to meet the agreed standards and to update as standards evolve. The challenge for decision-makers is investing in these testing and **certification processes**. The benefit, however, is reducing downstream troubleshooting and ensuring that when your system says "I support FHIR," it will actually work with another FHIR system out-of-the-box fhir.org.
- **Scalability and Performance:** Exchanging individual resources via REST works well for many scenarios, but healthcare datasets can be huge, raising questions of **scalability**. If, for example, a public health agency needs to pull millions of records (population-level data) from hospital FHIR servers, doing this one resource at a time would be inefficient. To tackle this, the Bulk Data (Flat FHIR) specification allows exports of large data sets in a single job. However, real-world experiences have uncovered **performance limitations** with bulk FHIR exports – such as memory and processing strain on servers, network bottlenecks, etc. A 2023 study on FHIR bulk data performance noted that addressing these issues is critical to achieve the Cures Act interoperability goals pmc.ncbi.nlm.nih.gov. Implementers have to consider server architecture: *How* will the FHIR server handle 100,000 simultaneous requests or a query for all 1 million patients? Caching, indexing, and robust infrastructure (potentially cloud auto-scaling) are needed at the **server side**. On the client side, efficient paging through results and handling of large JSON files are considerations. There's also the matter of **transaction scalability**: Can the FHIR API handle high-throughput data ingestion (like a surge of device readings)? These are not theoretical – consider a national health app where thousands of users upload data concurrently. While FHIR itself is not inherently a bottleneck, the typical JSON/HTTP pattern has overhead compared to say, a specialized message stream. Some in the community are exploring gRPC or GraphQL with FHIR to improve performance, but those are not standard yet. For now, implementers should follow best practices (e.g., use of `_since` parameters to limit data, asynchronous calls for bulk data, etc.) and monitor performance. **Load testing** FHIR APIs is a recommended step before production use. In summary, scaling FHIR to enterprise or national levels is doable (major cloud vendors provide managed FHIR servers that claim high scalability), but you must design for it and possibly accept that very large data moves may still be more efficient via database dumps or specialized pipelines if FHIR overhead is too high. The community is aware of this and working to optimize FHIR for big data.

- Data Mapping and Legacy Integration:** As noted, most healthcare organizations have existing investments in HL7 v2 messages, CDA documents, and other proprietary formats. One of the practical challenges is **mapping and translating** these legacy data flows into FHIR. For example, an ADT admit message (v2) carries patient demographics and encounter info that would map to several FHIR resources (Patient, Encounter, perhaps an Observation for vitals, etc.). Doing this mapping requires expertise and tooling. There have been efforts and tools (some open source) to convert HL7 v2 to FHIR – for instance, mapping guides for ADT to FHIR exist, and some interface engines (Mirth, Rhapsody, etc.) now offer FHIR translators. Still, the task is non-trivial: v2 messages often have site-specific variations, and information in a free-text v2 segment might need coding for a structured FHIR resource. The **systematic literature** indicates relatively few turn-key solutions for v2->FHIR yet medinform.jmir.org, meaning many organizations build custom mappings. Similarly, mapping CDA documents to FHIR bundles is an area of active work (C-CDA on FHIR guides exist to convert U.S. CDA documents to equivalent FHIR Composition and related resources) medinform.jmir.org. Until FHIR is fully ubiquitous, **hybrid environments** will persist, so integrators need capabilities in both worlds. One strategy is deploying an API layer on top of legacy systems: for example, a middleware that listens for HL7 v2 messages internally but then exposes a FHIR API to external consumers, translating on the fly. This way consumers get the modern API without ripping out the old system. For data warehouses and analytics, mapping historical data (often in CSVs or custom schemas) to FHIR can also be a heavy lift – but doing so can unify data and enable easier sharing. The **cost of mapping** is a major consideration when adopting FHIR in an organization with large legacy infrastructure. It's often recommended to use FHIR for new integrations and gradually backfill support for legacy feeds as needed. Over time, as more endpoints natively speak FHIR, the mapping burden will diminish, but during the transition it's a significant challenge. Additionally, staff skilled in both older HL7 standards and FHIR are needed to ensure nothing is lost in translation. This is where consulting HL7's published mappings or participating in implementation communities (like HL7's Confluence pages on mapping v2 to FHIR) can be very helpful.

- Semantic Interoperability and Data Quality:** FHIR can standardize the “container” and transport of data, but true interoperability also demands consistent **data semantics**. Organizations face the challenge of normalizing codes and **terminologies** across systems. For example, one EHR might record a diagnosis with an ICD-10 code, another with a SNOMED CT code – when aggregating via FHIR, ensuring those are understood equivalently is important. FHIR encourages using standard code systems in each resource (and profiles often mandate specific code systems), but legacy data may use local or proprietary codes that have to be mapped to standard ones. Maintaining these mappings (e.g., local lab test codes to LOINC) is an ongoing effort. Furthermore, differences in clinical modeling can pose challenges: Two hospitals might both use FHIR to send lab results, but one uses the *Observation.referenceRange* element and the other doesn't, or they use different units. Such variation can impede seamless data merging. **Governance** is needed to achieve semantic alignment – usually through community-agreed implementation guides and profiles for specific use cases. Projects like *US Core*, *UK Core*, *SNOMED on FHIR* guidelines, and others are aimed at improving semantic consistency. Another aspect is **data quality**: FHIR doesn't inherently solve problems like missing or miscoded data. If an EHR's database has an allergy recorded as free text, exposing it via FHIR still yields a free text string unless the data is cleaned and coded. So organizations often undertake data normalization alongside FHIR API deployments. On the research side, some are leveraging FHIR's RDF format to enhance semantic integration – converting FHIR resources to an RDF graph can allow linking with ontologies and reasoning over data nlm.nih.gov nlm.nih.gov. While promising, these techniques are complex and not yet mainstream. The bottom line is, **semantic interoperability remains a challenge** – FHIR provides the framework (standard fields, ability to carry codes, etc.), but getting different actors to use it in a semantically consistent way requires coordination, agreements (often driven by regulators or professional bodies), and sometimes additional translation layers. We are seeing progress (e.g., mandating certain vocabularies for certain data classes in US and EU), but implementers should not assume that just adopting FHIR guarantees semantic alignment. Careful attention must be paid to coding practices, value set alignment, and consistent use of FHIR fields as intended.



- **Security and Privacy:** Although not explicitly in the question's list of challenges, it's worth mentioning as it is critical. Exposing rich clinical data via web APIs raises concerns about **authentication, authorization, and privacy**. FHIR itself does not mandate a single security method, but it strongly recommends using industry-standard web security (OAuth 2.0, OpenID Connect, TLS/HTTPS, etc.). In practice, most FHIR APIs (especially for patient data) use the SMART on FHIR OAuth profile for user authentication and consent. A challenge for organizations is implementing and maintaining these security layers – e.g., managing OAuth client registrations, refresh tokens, audit logging for API access, and ensuring only authorized data is disclosed. Compliance with privacy laws like HIPAA is essential: any FHIR API that handles Protected Health Information must meet the same safeguards as other health IT systems (encryption in transit, access control, audit trails, etc.) build.fhir.org build.fhir.org. Implementers must recognize that using FHIR doesn't magically make an interface HIPAA-compliant; **security must be designed in**. This includes considerations like: what data do we expose via API (and do we filter any sensitive elements)? How do we verify app developers or third parties consuming the API? What is the consent model? As APIs open up, healthcare organizations face an increased surface area for potential breaches if not managed well. On the flip side, FHIR's compatibility with modern security (OAuth scopes, user-centric auth) can enhance privacy – patients can be given fine-grained control (like an app only gets a *patient/.read* scope, meaning read-only access to that patient's data). But these measures must be implemented properly. It's advisable to follow the FHIR Security Implementation Guide and the **"implementer's safety checklist"** provided by HL7 build.fhir.org build.fhir.org. Regular penetration testing of FHIR endpoints and a solid **API governance** process are becoming standard practice in mature organizations. In summary, robust security is a challenge not unique to FHIR but highlighted by the fact that FHIR makes data exchange easier – we must ensure it's *safely* easier. We will discuss specific regulations (Cures Act, HIPAA) in the next section which also drive home these points.

Global Adoption and Real-World Implementation

FHIR's adoption has accelerated worldwide, driven by both grassroots developer uptake and top-down policy mandates. Here we highlight some notable national and international developments:

- United States (ONC/CMS and Nationwide APIs):** In the U.S., FHIR adoption received a major boost from federal regulation. The **21st Century Cures Act** (2016) and subsequent rules by the Office of the National Coordinator (ONC) and Centers for Medicare & Medicaid Services (CMS) require that electronic health records and health plans implement **"secure, standards-based APIs"** for data access – and in practice this means FHIR RESTful APIs using the HL7 FHIR standard rhapsody.health. In March 2020, ONC's Cures Act Final Rule explicitly named HL7 FHIR Release 4 as the baseline for certified EHRs' APIs, along with the US Core Data for Interoperability (USCDI) set of data elements [federalregister.gov](https://www.federalregister.gov). By the end of 2022, EHR vendors had to update their products to include a patient-facing FHIR R4 API (the "g(10) API" certification criterion) or risk loss of certification rhapsody.health. At the same time, CMS's Interoperability Rule required Medicare, Medicaid, and many private insurers to provide FHIR-based APIs for patients to access claims and coverage data. Together, these policies have made FHIR **mainstream in the U.S. health IT ecosystem** – virtually all major hospital EHR systems (Epic, Cerner, Meditech, Allscripts, etc.) now have FHIR API capability out of the box. Moreover, government agencies and contractors are building on FHIR for broader health information exchange: for instance, the VA and DoD are using FHIR in aspects of their joint health record modernization, CDC is exploring FHIR for public health as noted, and quality reporting programs are shifting to FHIR standards. The Argonaut Project (an industry-led FHIR initiative in the US) laid much groundwork for common FHIR profiles, and now the ONC's sanctioned profiles (US Core) carry that forward [federalregister.gov](https://www.federalregister.gov). The **result** is that the US now has a burgeoning API economy around health data – something that did not exist a few years ago. Patients can use apps of their choice to fetch records from different providers, payers and providers can exchange data more easily for care coordination, and new services (e.g. analytics, telehealth platforms) can integrate using FHIR rather than custom interfaces. The U.S. experience demonstrates how policy and standards together can advance interoperability: the regulation created the incentive (or mandate), and FHIR was ready as a mature standard to fulfill the need.
- United Kingdom (NHS and UK Core):** The UK's National Health Service has also embraced FHIR in its drive for interoperable digital health. NHS England (formerly NHS Digital) has stated that "FHIR is the global industry standard for passing healthcare data between systems" and is central to its API strategy digital.nhs.uk. Early on, the NHS funded the INTEROPen community and "CareConnect" FHIR profiles to ensure common data definitions in England digital.nhs.uk. With the publication of FHIR Release 4, NHS UK moved toward a **unified UK Core** – a single set of FHIR profiles to be adopted across England, Scotland, Wales, and Northern Ireland digital.nhs.uk. The goal is to enable **consistent information flows across all regions**, improving care especially for patients moving between different NHS services or regions digital.nhs.uk. Examples of FHIR use in the NHS include the GP Connect project (using FHIR to share general practitioner records with other providers), the NHS Digital COVID-19 data store (which leveraged FHIR for consolidating data), and various patient-facing services like NHS login and the NHS App which can pull in data via FHIR APIs. The NHS also uses FHIR for specific domains such as scheduling (booking appointments across systems) and has been working on FHIR-based messaging for referrals and transfers of care. In sum, the UK is creating a **national FHIR-based architecture**, where vendors and trusts adhere to the UK Core profiles ensuring semantic alignment. As of 2025, numerous NHS suppliers are updating their products to be FHIR-compatible, and NHS has published an **API Gateway** and catalogue where many APIs use FHIR standards. This mirrors the U.S. approach in some ways but under a single-payer system, meaning the NHS can directly standardize requirements. The UK's adoption also underscores collaboration – the UK is active in HL7's international community to ensure UK needs are reflected in the core FHIR spec and vice versa.



- **Other Countries and Global Efforts:** FHIR is a true international standard, and many countries are following suit with their own initiatives. **Canada** has a FHIR-based Pan-Canadian Patient Summary standard and various provincial projects (e.g., Ontario's use of FHIR for lab results exchange). **Australia**, home of FHIR's creator, was an early adopter; they have the "Au Core" (previously Argonaut AU) profiles and use FHIR in national programs like My Health Record and for COVID vaccination records. **New Zealand** similarly has embraced FHIR for its health information platform. Across Europe, FHIR is being incorporated into the EU's digital health frameworks – for instance, **France** and **Germany** have national FHIR guides (France's insurers use FHIR for billing data exchange; Germany's electronic patient record supports FHIR). A 2024 survey found numerous countries with official FHIR **implementation guides**: not only the US and UK, but also Switzerland, the Nordics (e.g., Finland, Denmark), the Netherlands, Italy, and more [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov). Many of these are coordinated through HL7 International affiliates. The **Global Digital Health Partnership** (an international forum of over 30 countries) has also identified FHIR as a key enabler for cross-border interoperability. In Asia, countries like **India** are exploring FHIR for their national health ID and records system, and **South Korea** (as detailed in a 2021 study) is strategizing FHIR adoption nationwide [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov) [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov). Even low-resource settings find FHIR useful: for example, there are FHIR implementations in Kenya and other countries through the OpenHIE community (for exchanging public health and HIV data). Moreover, all the major global health software (OpenMRS, DHIS2, etc.) are building FHIR APIs to allow integration. A significant indicator of global adoption is that big cloud vendors – Amazon AWS, Google Cloud, Microsoft Azure – now offer **managed FHIR server services** [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov). This makes it straightforward for a health startup or a regional network to spin up a FHIR backend, accelerating adoption around the world. In summary, **global adoption of FHIR is well underway**, with many countries at least experimenting if not mandating it. This fosters a virtuous cycle: as more implementers come on board, the standard matures and tools improve, making it easier for the next wave. There is a recognition that in order for data to truly flow globally (for patient travel, public health crises, or research), a common standard like FHIR is invaluable. International collaborations (such as HL7's "Vulcan" project bridging clinical research or "IPS" – International Patient Summary standard based on FHIR) exemplify how FHIR is used as a lingua franca to exchange health data across borders.

Regulatory and Policy Implications

21st Century Cures Act and Interoperability Mandates

The **21st Century Cures Act**, passed in the U.S. in 2016, set in motion a series of regulations that have profoundly influenced FHIR's adoption. One of the Act's aims was to drive interoperability and prevent "information blocking" – deliberately hindering the flow of health data. To operationalize this, the U.S. Department of Health and Human Services (HHS) issued rules (in 2020) that heavily feature FHIR:



- The **ONC Cures Act Final Rule** (by the Office of the National Coordinator) requires that certified Health IT (EHR) systems provide a “**Standardized API for patient and population services**” as part of certification. In technical terms, this means EHRs must implement an API that allows software applications to access **all electronically accessible medical record data** for a patient, and do so using **HL7 FHIR R4** at a minimum rhapsody.health. The rule references specific HL7 FHIR profiles (the U.S. Core Implementation Guide, based on FHIR R4) as the **baseline** that systems must support [federalregister.gov](https://www.federalregister.gov). This API requirement replaced earlier, less specific API criteria and has effectively made FHIR R4 support mandatory for EHR vendors in the U.S. by December 2022. It covers both single-patient data access (e.g., an app pulling one patient’s chart) and **population-level data export** (the bulk data API, for analytics and other uses). The ONC rule also ties into the Information Blocking provisions – if a patient requests their data or a third-party app tries to access data via the certified API, a provider/EHR company **blocking** that (unreasonably) could face penalties. Thus, FHIR is at the heart of enabling patients’ access to data, which is a right under Cures Act rules rhapsody.health.
- The **CMS Interoperability and Patient Access Rule** (by the Centers for Medicare & Medicaid Services) complements the ONC rule by imposing requirements on insurers/payers. As of 2021, CMS-regulated payers (Medicare, Medicaid managed care, ACA Exchange plans, etc.) must implement a **Patient Access API** and a **Provider Directory API**, both using FHIR. The Patient Access API allows members to use any third-party app to retrieve their claims, encounter data, and certain clinical data held by the payer (such as lab results the plan has) – again putting the patient in control. The Provider Directory API (FHIR-based) makes health plan provider networks openly accessible for care coordination or patient information. CMS also finalized a requirement for a FHIR-based **payer-to-payer exchange** of clinical data whenever a patient moves between health plans. Together, ONC and CMS rules essentially widen FHIR adoption beyond providers to the **payer sector**, promoting a more integrated view of patient data across the continuum rhapsody.health.

These regulations explicitly or implicitly endorse FHIR as **the standard** for healthcare APIs. The 21st Century Cures Act did not name FHIR in statute, but the logic was that the industry (via HL7 and initiatives like Argonaut) had coalesced around FHIR as the best candidate, and the government moved to accelerate its deployment. **Compliance deadlines** were phased (some were extended due to COVID-19), but by 2023 nearly all certified EHRs and major payers had functioning FHIR APIs in production. The enforcement of these rules (including monitoring for information blocking and possibly future penalties) further cements FHIR’s role. Essentially, it is now *expected* that if a patient asks for their data, you will provide it via a FHIR API (among other formats). For developers, the rules also encourage a **competitive ecosystem** – as the playing field is leveled (every EHR has similar APIs), innovation is happening in the app market using those APIs.

Another aspect of the Cures Act rule is the emphasis on **standards version advancement process** – ONC set up a way to smoothly move to newer versions of FHIR (like R4 to R4B or R5) in the future via administrative procedures, which underscores that the standard will continue evolving under regulatory oversight.

Finally, it’s worth noting that the **information blocking** provision doesn’t allow a provider or IT vendor to refuse data exchange by saying “we don’t support that.” Since FHIR is the required

standard, not supporting it would be untenable. The rules together create a powerful incentive: to avoid regulatory risk and to stay competitive, healthcare organizations must implement FHIR-based interoperability. Thus, the Cures Act is often cited as a model where **policy drives standards adoption**, effectively making FHIR not just a technical choice but a compliance requirement.

HIPAA and Data Privacy

The **Health Insurance Portability and Accountability Act (HIPAA)**, specifically its Privacy and Security Rules, provides the foundational federal protections for health information in the U.S. Any electronic exchange of Protected Health Information (PHI) – regardless of format – falls under these regulations if done by a covered entity (healthcare provider, payer, clearinghouse) or their business associate. FHIR does not change this reality: a FHIR API **must be implemented in a HIPAA-compliant manner** when handling PHI. In other words, adopting FHIR neither exempts one from HIPAA nor inherently fulfills HIPAA – it's **agnostic** to compliance, so the onus is on implementers to enforce privacy and security controls formdr.com.

Key considerations at the intersection of FHIR and HIPAA include:

- Access Control and Authentication:** The HIPAA Security Rule requires controlling access to PHI. For FHIR APIs, this translates to robust authentication (confirming the identity of whoever/whatever is accessing the data) and authorization (ensuring they only get appropriate data). Most FHIR implementations use OAuth 2.0, where the user (patient or clinician) authorizes an app to access data on their behalf, and the app gets an access token with specific scopes. This aligns well with the "minimum necessary" principle of HIPAA, as scopes can limit access (e.g., an app only gets `patient/observation.read` if it only needs to read vitals). However, implementing OAuth and identity management properly is non-trivial. It requires an Identity Provider, and mapping user identities to patient records, etc. Provider-facing APIs might use system authentication (client credentials). All these flows must be secure – using strong encryption (TLS), expiring tokens, etc. Under HIPAA, one must also have policies for workforce access: e.g., if clinicians use an API, ensure they only query patients they have a relationship with (this often is enforced by context or audit rather than FHIR itself).
- Encryption:** HIPAA Security Rule doesn't mandate encryption per se, but it's an addressable implementation specification and industry expectation to encrypt PHI in transit. FHIR APIs should **always be served over TLS/HTTPS** (which provides encryption in transit). Many implementations also choose to encrypt PHI at rest on the FHIR server's database. Modern FHIR server implementations (especially those in cloud environments) support database encryption and other safeguards as part of configuration. The use of JSON and HTTP doesn't inherently pose new encryption issues (HTTPS is standard), but APIs potentially expose more endpoints than traditional HL7 feeds, so the surface area to secure is larger. HIPAA also requires monitoring and protecting the network – so API endpoints need intrusion detection and other network security around them build.fhir.org.



- **Audit Logging:** HIPAA mandates audit controls – the ability to record and examine activity in systems containing PHI. A FHIR server must therefore log access events: which user/app accessed what data, and when. Many FHIR servers produce an audit log (and even have a FHIR *AuditEvent* resource format for standardized audit data). It's important to configure these logs properly and review them. For example, if an app is scraping large amounts of data, it should be flagged if not expected. Also, if using SMART on FHIR, the authentication server should log consent events (e.g., patient X authorized app Y to access data at time Z). These logs play into both security monitoring and compliance reporting.
- **Data Minimization and Segmentation:** Not all health data is equal under privacy laws – for instance, 42 CFR Part 2 data (substance abuse records) has extra protections, and some state laws protect reproductive health or mental health records more strictly. There is ongoing work (e.g., *FHIR Data Segmentation for Privacy* IG) on how to tag and segment sensitive data in FHIR so that it can be filtered out if needed build.fhir.org build.fhir.org. HIPAA allows patients to request restrictions on certain disclosures; an implementation may need to honor that by not exposing certain data via API unless specific conditions are met. This is a **complex area**: implementing fine-grained privacy controls in an API context. Many systems currently avoid exposing especially sensitive data via APIs at all (or require special scopes). For example, an EHR might suppress behavioral health notes from the FHIR API unless explicitly permitted. As FHIR matures, we expect more nuanced consent and tagging solutions. For now, implementers should be aware of the data they are exposing and any applicable rules beyond HIPAA (e.g., GDPR in Europe has its own consent rules – a number of European FHIR projects incorporate consent mechanisms).
- **Business Associate Agreements (BAAs):** If a third-party cloud host or a third-party app is involved in handling PHI via FHIR, ensuring the right legal agreements are in place is crucial under HIPAA. Healthcare providers using a cloud FHIR service (like Azure API for FHIR or AWS HealthLake) must have a BAA with that cloud provider. Similarly, when patients direct data to apps of their choosing, HIPAA places responsibility on the covered entity up until the point of transfer. Once PHI is with a patient's designated app, HIPAA may no longer apply to that app (if the app is not a covered entity or BA) – instead, FTC or state laws might cover it build.fhir.org. Providers need to educate patients that their data, when moved to a personal app, is under different protection. ONC's rules even require certain language to warn patients about this. Thus, from a policy view, FHIR is an *enabler* for patient data access but also raises consumer privacy issues when data goes outside HIPAA boundaries.

In essence, **FHIR's relationship with HIPAA** is about implementation: FHIR provides a uniform method to exchange data, and HIPAA provides the rules to protect that data. Implementers should follow the comprehensive guidance available in the FHIR Security modules build.fhir.org, employ **best practices** (use OAuth 2.0 with appropriate scopes, TLS, etc.), and ensure that their use of FHIR APIs is covered by their overall HIPAA compliance program. Many EHR vendors have built their FHIR APIs with these considerations from the start, given that they must attest to ONC that they meet security requirements. By doing so, they demonstrate that a well-implemented FHIR API can be **just as secure as any other method** of health data exchange – and potentially more transparent and controllable. As one commentary summed up, FHIR doesn't make you HIPAA-compliant or not by itself; it's **how** you use it that matters formdr.com.



Conclusion

In a relatively short span, FHIR has grown from an experimental idea into a **cornerstone of healthcare interoperability**. Its origin – learning from the past and embracing modern tech – positions it uniquely at the intersection of healthcare and IT. Technically, FHIR's resource-based RESTful architecture brings health data exchange into line with web development practices, opening the door for a new generation of health applications and integrated systems. Strategically, FHIR has gained support from virtually all stakeholders: standards bodies (HL7), governments (through mandates and national programs), vendors (EHR and tech companies implementing and promoting it), and increasingly patients and clinicians (who benefit from data being more accessible when and where needed). This report has detailed how FHIR works, how it compares to earlier standards, and how it's being used – as well as the challenges that remain around implementation details, scaling, semantic alignment, and privacy.

The road ahead for FHIR will involve continued **maturation and optimization**. We can expect further alignment of international efforts (so that a patient summary in one country can be understood in another), more robust tooling (for validation, mapping, etc.), and likely expansions of FHIR into new domains (social care, veterinary health, etc.). Community-driven implementation guides will refine how FHIR is used for specialized workflows like cancer reporting or pandemic response. Meanwhile, legacy systems will gradually either be replaced or wrapped by FHIR layers, simplifying the overall interoperability landscape.

Importantly, FHIR's rise doesn't mean that interoperability is "solved" – it means we finally have a common framework to tackle the remaining problems. Issues like data quality, true semantic interoperability, and ensuring equitable access to technology will require as much attention as before. But with FHIR, stakeholders are at least "speaking the same language," which is a fundamental step forward. Decision-makers in healthcare IT should view FHIR not just as a compliance box to check, but as an **opportunity**: an API ecosystem that can enable new care models, patient engagement tools, and data-driven improvements in outcomes and efficiency. The investment in doing it right (through governance, talent development, and technology upgrades) will pay dividends as healthcare continues to evolve into a more connected, patient-centered information continuum.

In summary, FHIR has transformed interoperability from a niche, background IT concern into an **enabler of innovation** that is visible and impactful at all levels of healthcare. By understanding its foundations, leveraging its capabilities, and remaining vigilant about its challenges, healthcare organizations and developers can harness FHIR to improve data exchange – ultimately advancing the triple aim of better care, improved health, and lower costs through informed, data-rich decision making.

Sources:



- Ayaz, M. et al. *"The Fast Health Interoperability Resources (FHIR) Standard: Systematic Literature Review of Implementations, Applications, Challenges and Opportunities."* JMIR Med Inform 9(7):e21929, 2021 medinform.jmir.org medinform.jmir.org medinform.jmir.org medinform.jmir.org medinform.jmir.org.
- HL7 International. *"FHIR Release 4 – Overview and Specification."* (Accessed 2025) nml.nih.gov nml.nih.gov fhir.org.
- Rhapsody Health. *"FHIR vs. HL7: Key Differences in Healthcare Interoperability."* Blog article, 2020 rhapsody.health rhapsody.health rhapsody.health rhapsody.health.
- Centers for Disease Control and Prevention (CDC). *"Public Health FHIR Playbook."* July 2023 cdc.gov cdc.gov.
- NHS England Digital. *"FHIR APIs – Overview."* Last edited Jan 2025 digital.nhs.uk; *"FHIR UK Core."* Feb 2024 digital.nhs.uk.
- Office of the National Coordinator (ONC), HHS. *Federal Register* 85 FR 25642 (Cures Act Final Rule), 2020 federalregister.gov.
- FHIR Foundation. *"FHIR Conformance Testing."* (fhir.org, accessed 2025) fhir.org fhir.org.
- Apple Inc. *"Technical specifications for Health Records (FHIR API)."* Apple Support, 2022 support.apple.com support.apple.com.
- ONC Health IT. *"FHIR Security and Privacy Module (Implementer Guidance)."* (accessed 2025) build.fhir.org build.fhir.org.
- Byoung-Kee Yi et al. *"Strategy to Adopt and Deploy HL7 FHIR Standard... in Korea."* HIR 27(3):181-190, 2021 pmc.ncbi.nlm.nih.gov pmc.ncbi.nlm.nih.gov.
- Lehne, M. et al. *"Why digital medicine depends on interoperability."* NPJ Digital Medicine 2:79, 2019 pmc.ncbi.nlm.nih.gov.
- HL7 International. *"FHIR v4.0.1: US Core Implementation Guide."* 2020 federalregister.gov.



IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies including Scilex Holding Company (SCLX) and leading CROs across North America.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.



DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.