KODJIN
by Edenlab

The Ultimate Guide to FHIR

by FHIR experts from Edenlab,
creators of Kodjin Interoperability Suite

November, 2024

Contents:

UNIT I. FHIR: THE BASICS

1. Why We Talk About FHIR: Edenlab's Expertise

Edenlab is a custom software and product development company that provides healthcare data interoperability solutions based on the HL7 FHIR standard. We help healthcare companies with interoperability, compliance, and movement to new standards.

Our expertise: we run national systems into production

36,5M+ patient records

With thoughtful knowledge of the domain and experience in working with complex high-load projects, we created the Kodjin Interoperability Suite —FHIR-based software solution that helps our clients get the maximum advantage out of a project's cost.

The Kodjin Interoperability Suite offers the following:

FHIR Server

A turnkey event-driven FHIR server for healthcare data storage, data exchange, and management.

Terminology Service

An FHIR service to handle all types of healthcare terminology.

FHIR Profiler

An FHIR profiling tool that helps create FHIR profiles in the fastest and easiest way, following an Agile approach and using an intuitive graphical IDE.

Data Mapper

An interoperability FHIR software to map and transfer healthcare data elements with Liquid template language.

Read more about Kodjin Interoperability Suite at Kodjin.com, Kodjin FHIR Server Documentation

UNIT I. FHIR: THE BASICS

2. The History Behind FHIR

Health data is diverse and sensitive information that requires specific storage, exchange, and security guidelines. The healthcare industry requires a data exchange approach that contributes to the quality of patient care and helps explore the population's data to track trends and promptly respond to the needs of society.

Standardization of the healthcare data exchange is necessary to establish rules for fast and secure communication between health systems and bring all global healthcare stakeholders to work together on health data interoperability.

This guide includes all the needed information on leveraging the main principles of healthcare data interoperability.

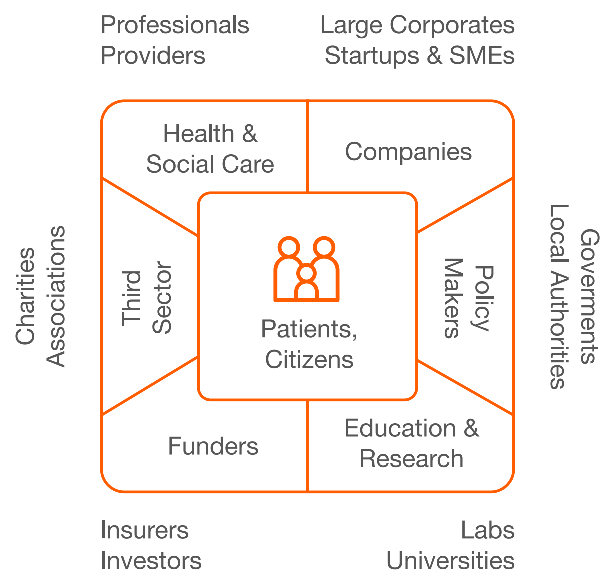

Preconditions for the new healthcare interoperability standards development originated in the early 2000s. At that time, the most widely used healthcare IT standard, namely the HL7 version 2 (v2.x), had been setting rules for the medical data exchange process for over ten years. HL7 v2 is a messaging standard that establishes communication between systems within one enterprise. This failed to address interoperability issues. Even then, the demand for a straightforward exchange of clinical information was already high, so HL7 decided to update the HL7 v2.x to HL7 v.3.

However, the newer version of the standard didn't receive a general appreciation, as it wasn't backward compatible with the HL7 v2. Moreover, HL7 v.3 introduced the concept of RIM (Reference Information Model), an information model for all healthcare information representation. This version of HL7 was meant to cover 100% of use cases, which resulted in an enormous volume of the standard specification. Its implementation required a massive investment in health IT systems modifications. For that reason, the standard was not widely accepted. The adoption of the third version of the HL7 standard was eventually considered to bring little value to healthcare organizations, as it was too costly and complex to implement.

The healthcare system in the US was not designed in the best way to deliver healthcare, which caused healthcare expenditures growth from $74.1 billion in 1970 to $1.4 trillion in 2000.

Total national health expenditures, UA $ Billions, 1970-2020

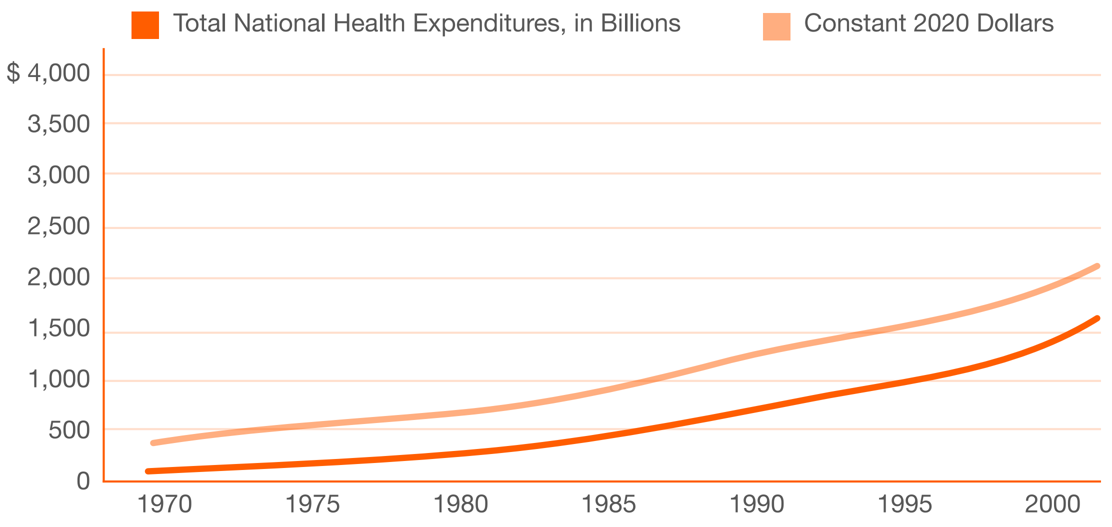

Source: KFF analysis of National Health Expenditure (NHE) data

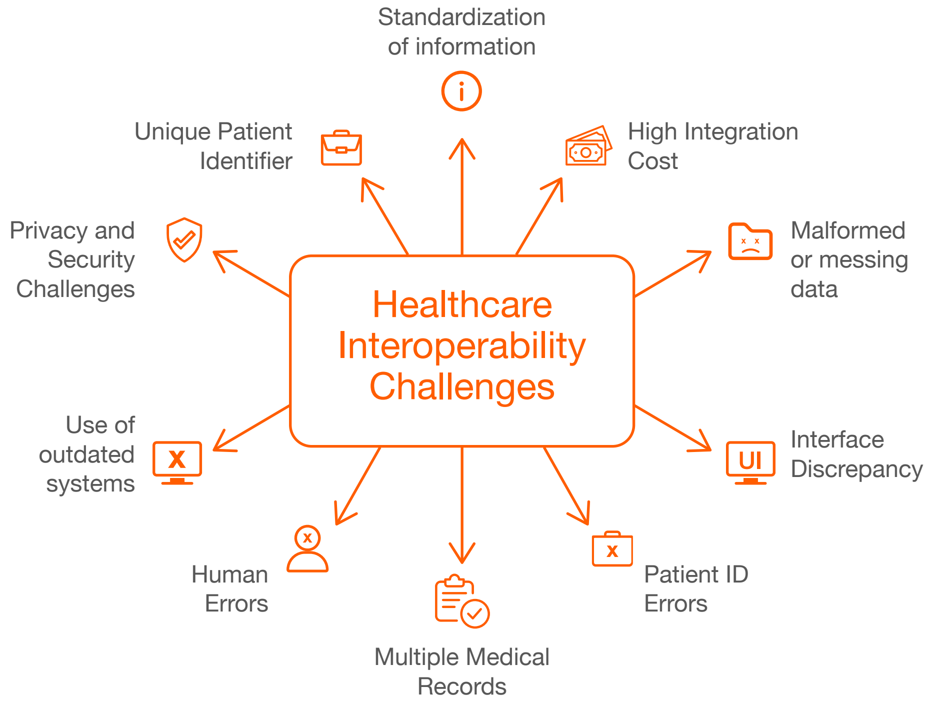

Such an outcome encouraged the U.S. Department of Health and Human Services to ask JASON Advisory Group to investigate what the industry was missing and how to motivate healthcare organizations to leverage new standards. JASON came back with a report claiming that the existing standards were an outdated approach to healthcare data interoperability and suggested developing healthcare IT systems using a comprehensive set of public APIs.

As such, the idea was to create a data interoperability standard with undeniable benefits for clinical IT systems developers and users that would take little time to learn and implement. As a result, HL7 put FHIR on the table.

FHIR regulations

In 2009, President Barack Obama signed the American Recovery and Reinvestment Act of 2009 (ARRA), including the Health Information Technology for Economic and Clinical Health Act (HITECH). The HITECH act aimed to popularize meaningful storage and use of healthcare information technology.

The HITECH sanctioned financial incentives for healthcare providers to encourage certified EHR systems adoption. The initiative has worked since the rate of hospitals that store patient charts in EHR have grown from 9% to 78% in 10 years.

Trends in Hospital & Physician EHR Adoption

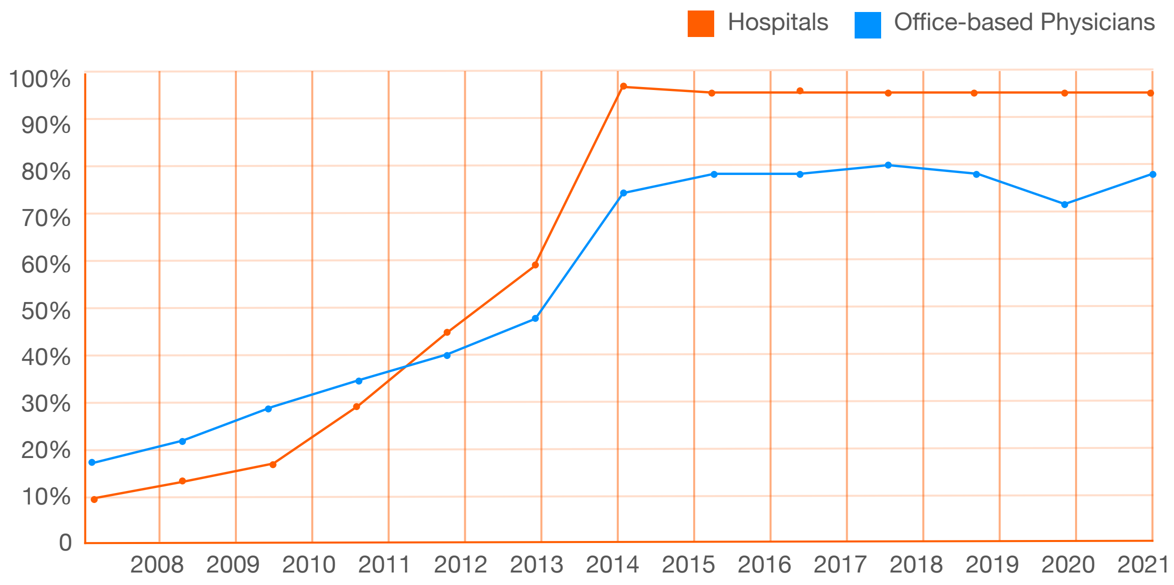

Source: Healthit.gov

To get the financial incentive, a healthcare provider needed to adopt an ONC-certified EHR and meet the requirements of the CMS EHR Meaningful Use (MU) Incentive Payment Program, which obligated providers to:

- Protect the information in EHR;
- Transfer electronic prescriptions;
- Exploit clinical decision support tools;
- Enter the clinical data using the computerized provider order entry (CPOE);
- Ensure patients' access to their EHR information;
- Engage patients in decision-making in the context of their care;
- Create patient care record summaries for a smooth patient transition within different medical facilities;
- Set up interactions with a public health agency and clinical data registry to submit data using a certified EHR.

This has led to an influx of new EHR solutions that met HITECH's requirements but barely covered usability and interoperability. In addition, the lack of standardization on EHR adoption and switching between different EHR product providers made this process exhausting for clinicians.

21st Century Cures Act

Congress passed this Act in December 2016. The Act includes interoperability and patient data access improvement proposals, including:

- the requirement for certified EHRs to support APIs;
- establish export of EHR data to third-party applications on a patient's demand;
- adopt the United States Core Data for Interoperability standard, which describes the core information that any EHR must include;
- make all EHR data computable to make switching IT systems effortless.

The 21 Century Cures Act's goals:

- data exchange transparency;
- unhindered access to medical data via applications for patients;
- ability to choose the IT system that suits the organization's workflow without any data loss or additional investments for providers.

In March 2020, the Office of the National Coordinator for Health Information Technology (ONC) released the 21st Century Cures Act Final Rule. The main goals of the ONC's Final Rule were to:

- provide secure access to healthcare data for patients and healthcare providers to provide care;
- encourage new healthcare IT product development;
- offer more choices in healthcare for patients.

One of the Final Rule's provisions requires free electronic access for patients to all their healthcare data via an application of their choice, free of charge. Another important call of the regulation was the requirement for all ONC-certified health IT systems to adopt standardized APIs. ONC defines APIs as "powerful tools that help support interoperability in healthcare".

The US legislation advanced the usage of FHIR for health IT systems by making the adoption of FHIR mandatory by December 2022. The regulation has enabled all healthcare IT solutions to leverage health data from any EHR, foster interoperability, and support new IT solutions development.

Trusted Exchange Framework and Common Agreement (TEFCA)

In January 2022, responding to the 21st Century Cures Act, ONC, in collaboration with The Sequoia Project, introduced the common principles, terms, and conditions for data exchange between health information networks (HINs) through TEFCA.

Objectives:

- Facilitate secure data access;
- Ensure core data availability;
- Cut costs by eliminating multiple HINs;
- Establish common privacy and security standards.

ONC pushes TEFCA's implementation and the creation of a network of networks by enabling entities to apply for Qualified Health Information Network (QHIN) status for nationwide health information exchange.

TEFCA suggests using FHIR to foster interoperability, trust, and efficiency in healthcare, overcome technical constraints, and enhance workflows for better outcomes.

(HTI-1) Final Rule Certification Program Updates, Algorithm Transparency, and Information Sharing

Regulatory bodies, such as ONC and CMS, actively encourage the adoption of the FHIR standard to advance interoperability through revised regulations.

Enable the following functionality for your health IT system to ensure compliance with updated HTI-1 final rule requirements by

ONC Office of the National Coordinator for Health Information Technology

CMS CENTERS FOR MEDICARE & MEDICAID SERVICES

Year	Changes under the rule:	Who must comply?	What to do to stay compliant?
2024	ONC (g)(10) Standardized Patient API	Developers of Health IT Solutions	Ensure availability of certified APIs for access to patient data via tools based on SMART on FHIR technology.
	Decision Support Interventions (DSI)	Developers of AI-based clinical DSI systems	• Ensure non-discriminatory use of clinical algorithms. • Obtain Food and Drug Administration (FDA) approval for software safety and effectiveness of an AI-based solution. • Provide users with access to predictive DSI performance data. • Ensure compliance with the Decision Support Interventions Certification Requirements.
2025	The United States Core Data for Interoperability: USCDI	Developers of Health IT Solutions	Update to USCDI v3 by using the HL7 FHIR USCore 6.1.0 Implementation Guide.
	Decision Support Interventions (DSI)		Use the eCRNow FHIR App to switch from HL7 CDA to HL7 FHIR reporting methodology to ensure automated data extraction.
2025	ONC (e)(1): "View, Download, and Transmit to a 3rd Party"	EHR Software Developers	Ensure the internet-based method for patient restriction requests to support individuals' right under HIPAA Privacy Rule to request restriction of certain uses and disclosures of their data.
2025 -2027	Insights Reporting Requirements		Comply with insights reporting criteria based on the size of the EHR, ensuring adherence to ONC's mandatory reporting requirements starting in early 2026.

The Office of the National Coordinator for Health Information Technology (ONC) has released the (HTI-1) Final Rule Certification Program Updates, Algorithm Transparency, and Information Sharing. This update introduces several new mandates for healthcare stakeholders to foster FHIR adoption for interoperability within the healthcare sector.

UNIT I. FHIR: THE BASICS

3. What Is FHIR

FHIR is an acronym for Fast Healthcare Interoperability Resources. This standard is a set of rules for fast interoperability principles implementation and secure medical data exchange. The main idea of this standard is to eliminate the weak spots in healthcare data exchange. It is based on a thoughtful analysis of shortcomings in previous versions of the HL7 standard.

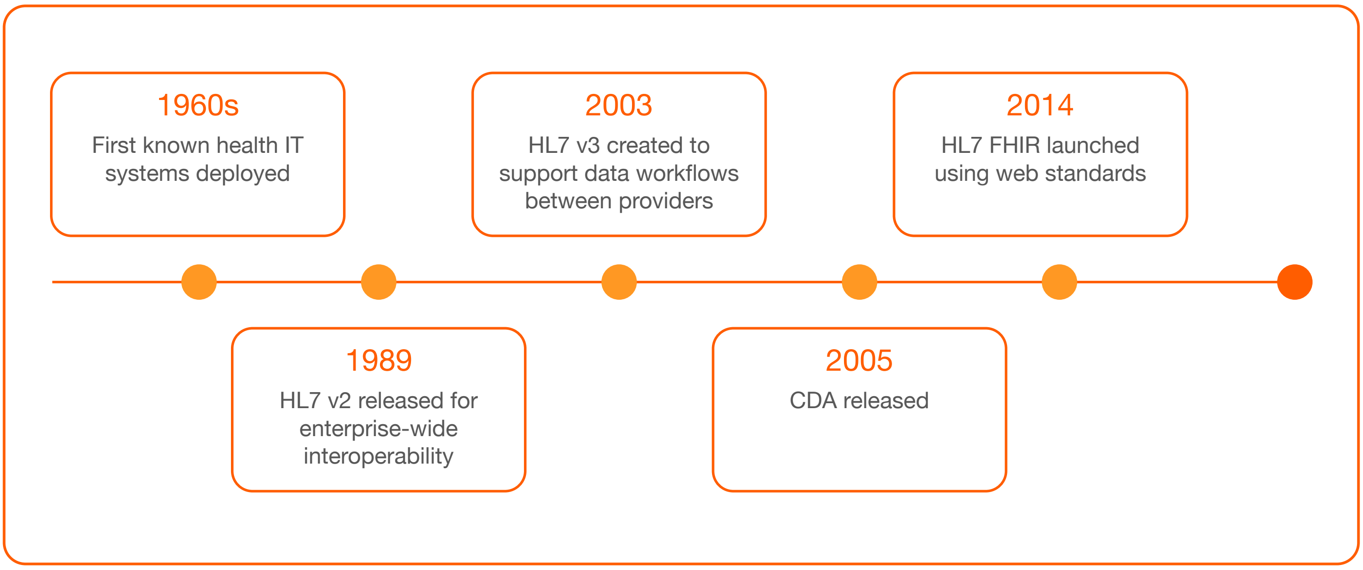

FHIR created a list of principles to stick to when building healthcare data exchange standards known as the FHIR Manifesto to avoid mistakes made during the development of the previous healthcare standards.

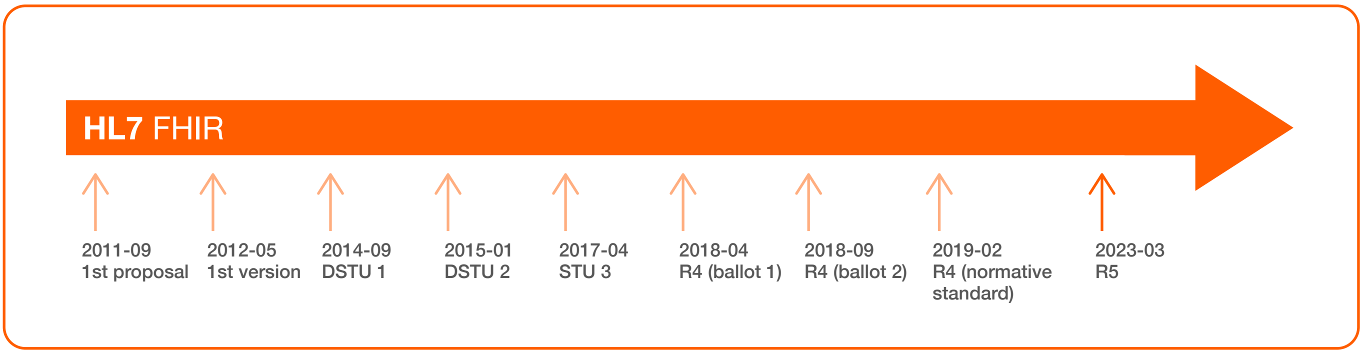

FHIR introduction — the first presentation of the FHIR occurred in May 2012.

FHIR version 0.0.82 — the Draft Standard for Trial Use 1 was published in September 2013. This version demonstrated the new way to forward data exchange in healthcare. It included 49 resources and covered the following use cases:

- a PHRs creation via mobile devices;
- retrieval of clinical documents to a mobile device.

FHIR version 1.0.2 — the Draft Standard for Trial Use 2 was published in 2015. FHIR resources were tested and approved with multiple FHIR-ready computer systems. This version introduces the FHIR Maturity Model, which defines the technical advancement of an FHIR resource, so-called maturity.

FHIR version 3.0.2 — (STU 3), the third version of FHIR released in 2017. The third version included a new framework for workflow and task management. This version included improvements for the following resources: clinical, administrative, financial, clinical decision support, and clinical quality measure.

FHIR version 4.0.1 — FHIR's fourth version was released in 2019. It contained patient, observation, and nine other resources, and RESTful API, XML, and JSON formats were listed as normative content. In 2020, ONC established FHIR R4 as a mandatory health IT certification standard.

FHIR version 5.0.0 — Released in March 2023, the R5 FHIR version includes features such as topic-based subscriptions, updated Medication Definition resources, several new resources, and 157 resource types.

UNIT I. FHIR: THE BASICS

4. Key Points of the FHIR Manifesto

Focuses mainly on implementers

FHIR uses language and tools that are readable and useful to implementers. The standard requires less time to learn than older healthcare standards, so healthcare organizations can adopt it faster so that a user can create healthcare solutions quickly.

80:20 Rule

One of the reasons the HL7 v3 failed was that standard developers tried to cover 100% of use cases, which resulted in a complex implementation mechanism. FHIR developers instead decided to define the core elements of the standard (base resources), which apply to 80% of use cases, while the other 20% of unique scenarios are covered by extending or constraining any FHIR resource.

Leverage existing web-based

The FHIR standard is made over HTTP protocol and adheres to the RESTful approach of API design. These technologies are widely used by the community of implementers so that they can stick to the familiar layer protocol.

Make information readable to humans

All data represented in the context of FHIR is human-readable. Even the coded parts of the information are modeled to be as human-readable as possible. Even if an IT system cannot process data for some reason, the end-user will be able to find the logic in the encoded information of the standard.

Keep FHIR open

FHIR is a standard with a minimal restrictions license. In addition, FHIR provides open-source libraries and relies on the FHIR community's feedback on the specification for further development.

Did you know the Cerner Corporation, Epic System Corporation, and other market leaders have adopted the FHIR API?
These corporations cover 83% of hospitals or 58% of clinicians

Interoperable healthcare data exchange...

COMPONENTS OF THE STANDARDS:

Resources

A resource is a data model. The main purpose of resources is to store and exchange content in various data models, such as patient, procedure, medication, and other resources. Resources are crucial components of the FHIR specification. There are five types of concept models: foundation, base, clinical, financial, and specialized resources.

Terminology

Terminologies are various coding systems, such as LOINC, SNOMED CT, ICD 10, and RxNORM, that assign codes to different clinical concepts. FHIR links to these systems, coded concepts, and definitions in the ValueSet resource.

The relationships between terminology-related structures are shown below:

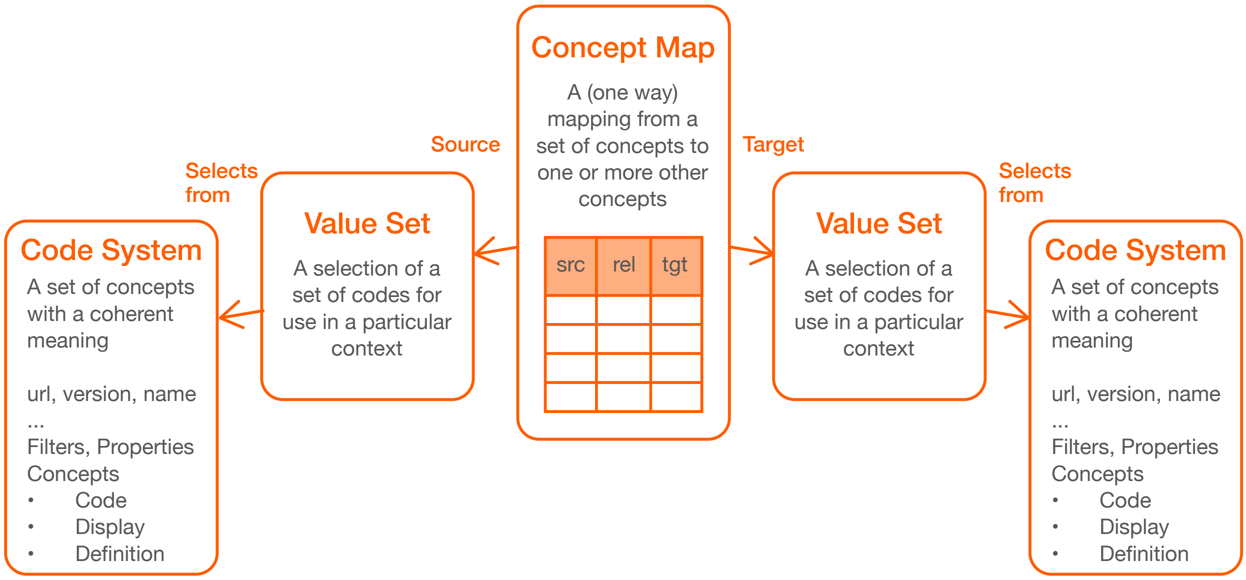

Profiling

Profiling is the mechanism of adapting FHIR resources for a particular use case and is one of the key components of the standard. Profiles define rules on how to restrict or extend resource elements and terminologies, depending on the context.

RESTful API

RESTful API defines a set of rules that tell applications how to connect and communicate with each other based on the REST architectural style.

Three aspects that make FHIR different from other health data standards:

Open source implementation library:

FHIR includes many open source implementation libraries that allow developers to create dynamic system designs using premade code and specifications. The open source community is large, and its participants work together to develop cutting-edge software and test it properly.

The net of FHIR:

hospitals, laboratories, insurance companies, pharma companies, universities, health tech startups, and enterprises all over the world widely adopt FHIR, expanding the FHIR ecosystem. The more organizations use FHIR, the higher the level of interoperability, as all systems conform to a single data standard. The US government has also specified the use of FHIR in their regulations, which further boosts FHIR adoption.

FHIR community:

focus on the community is one of the core ideas behind FHIR. Representatives of the FHIR community strongly believe in the need to share their knowledge and expertise to develop effective data interoperability solutions.

UNIT I. FHIR: THE BASICS

5. FHIR Accelerators

The FHIR community is a crucial part of the standard development and remains its greatest accelerator. For that reason, HL7 has designed the HL7 FHIR ACCELERATOR program. Accelerators are separate initiatives of technology vendors and healthcare stakeholders that work with HL7 and whose goal is to assist the healthcare community in creating and adopting quality FHIR Implementation Guides to achieve global health data interoperability.

The Argonaut Project

The Argonaut Project is a private-sector initiative that launched in 2014. Its goal was to establish health data exchange based on internet technologies.

The US Core IG development was guided and tested by the Argonaut Project team, which put its effort into building a solid implementation guide any representative of the healthcare ecosystem could stick to.

The CARIN Alliance

The alliance includes consumers, payers, health IT companies, consumer advocates, certifiers, identity providers, other US organizations, and the Government to address barriers and constraints to ensuring consumer-directed exchange in the US.

Participants of the CARIN Alliance develop IGs that will allow consumers and caregivers to get, use, and share digital health data when, where, and how they need it to achieve their goals.

CodeX

CodeX stands for Common Oncology Data Elements Extensions. The CodeX members test and integrate the mCODE (minimal Common Oncology Data Elements) FHIR Implementation Guide, which is used for cancer data exchange, in order to ensure better cancer care and research.

Source: https://confluence.hl7.org/display/COD

All CodeX use cases are consistent with the idea that data is collected once and re-used for various workflows.

The list of CodeX use cases or mCODE-based FHIR implementation guides includes:

- Cancer Registry Reporting;
- EHR Endpoints for Cancer Clinical Trials;
- Integrated Trial Matching for Cancer Patients and Providers;
- mCODE++ Extraction;
- Radiation Therapy Treatment Data for Cancer;
- Prior Authorization in Oncology.

The Da Vinci Project

HL7 initiated the Da Vinci Project to gather healthcare leaders and tech vendors to create standard mechanisms for adopting FHIR standards and guarantee smooth data exchange between all healthcare stakeholders.

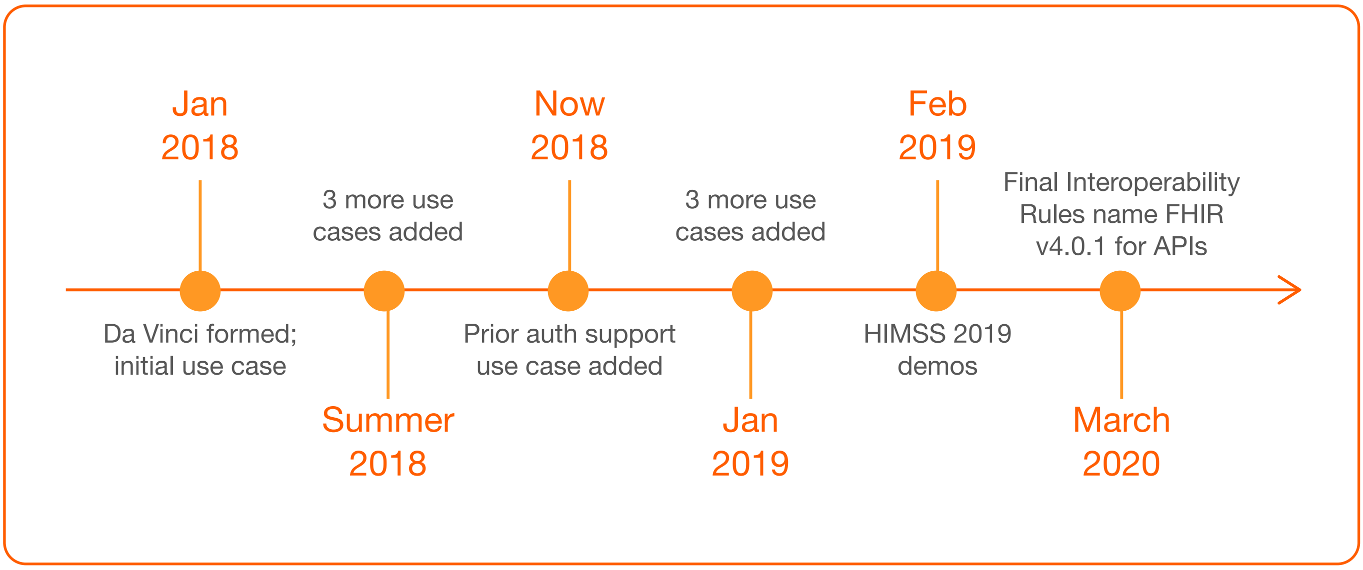

The Da Vinci Project creates solutions based on the collective expertise of industry experts to encourage a redundant shift to value-based care and help stakeholders minimize the cost of new IT solutions leverage. The Office of the National Coordinator for Health Information Technology (ONC) notes the efficiency of the Da Vinci Project's work. The Interoperability and Patient Access Final Rule named HL7 FHIR v4.0.1 as a standard for APIs in healthcare.

FAST

FAST stands for FHIR at Scale Taskforce and is a cross-stakeholder team that began operations in 2017 to address FHIR scalability challenges. The main focus of the FAST team is the infrastructure of implementation guides and industry guidance to ensure the adoption of use cases in a scalable manner.

Among their active projects is the National Healthcare Directory, which aims to build a national endpoint directory based on the FHIR approach and ensure the accuracy of endpoint data.

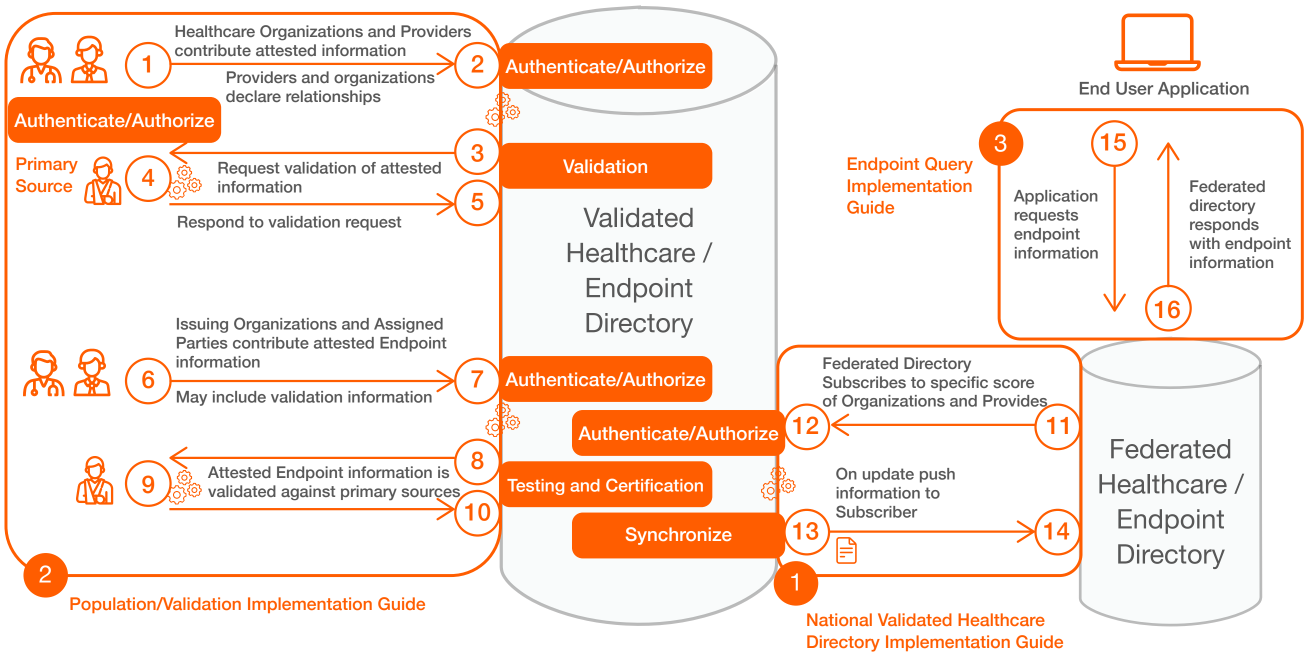

Source: https://confluence.hl7.org/display/FAST/FAST+Projects

Gravity Project

The Gravity Project was launched in 2018 by Social Interventions Research and Evaluation Network (SIREN) with funding from the Robert Wood Johnson Foundation and developed standards associated with social determinants of health (SDOH).

The project participants develop, test, and vote for IGs based on specific use cases and coded data elements in the context of diagnosis, screening, planning, and intervention. The technical workstream defines how to record, document, and exchange SDOH data and ensures its safety and security for patient records.

Helios

The complexity of public health causes data access and exchange challenges. The Helios accelerator gathers government, private sector, and philanthropic partners to effectively address these challenges and stick to the FHIR approach when dealing with public health data.
Helios' top priorities:

- make data in public health systems accessible in bulk;
- provide public health critical data needed during emergencies;
- assess and optimize optimal ways for public health to access data in EHRs.

Vulcan

The Vulcan accelerator is focused on developing FHIR standards for optimizing the design and reporting of clinical studies. Its main goal is to establish a partnership between clinical research and clinical care and eliminate duplicates in healthcare data.

The Need for Standards in Clinical Research Data is the Key

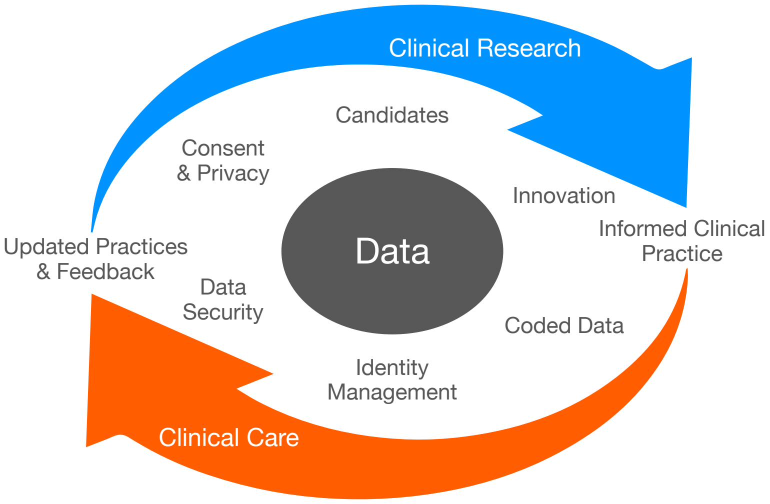

- Health data is used across abroad spectrum of Clinical Research and Clinical Care purposes
- Focus for standards & interoperability has been on Clinical Care, thus Clinical Research has lagged behind leveraging common standards, methods and techniques
- On the positive side, both share many common standards such as consent (to treatment, to research), terminologies and identity management
- Unique Clinical Research standards include candidate identification, clinical trials and phenopackets (to name a few) that rely on health data curated by clinical care processes

Source: https://confluence.hl7.org/display/VA/Vulcan+Accelerator+Home

Vulcan is the newest participant in the HL7 FHIR Accelerator Program. As well as other accelerators, it brings together stakeholders across healthcare to ensure health data interoperability, faster access to quality care, and better treatment outcomes.

UNIT I. FHIR: THE BASICS

6. Benefits of Leveraging the FHIR Standard

Apart from being fully implementers-ready, the FHIR standard provides numerous benefits to all healthcare actors by establishing a fast and secure exchange of medical data between all elements of the healthcare ecosystem.

Not only healthcare providers but also patients can reap the benefits of the widespread FHIR adoption. For example, FHIR-based applications help patients with chronic diseases monitor their conditions and share the dynamics of their health with doctors and other caregivers to receive the most effective treatment.

Unlike previous HL7 data standards, FHIR provides a fast and easy implementation for developers, thus further facilitating its adoption and boosting interoperability. It's also completely free, with an extensive online specification and downloadable tools, which include reference servers and implementation guides.

We hope this guide will have helped you understand FHIR more, and you can see for yourself how the ideas that drive it brought on its widespread adoption.

UNIT II. FHIR TECHNICAL COMPONENTS

1. Resources - the Buildings Blocks of the FHIR Standard

A resource is a set of core modular FHIR components. A resource defines a way to package information. In the context of FHIR, healthcare data must be divided into small units or resources.

General characteristics of any FHIR resource:

- has a resource identifier (URL) by which it can be addressed;
- comprises a specific well-structured set of data elements for each individual resource;
- identifies itself as one of the types of resource defined in this specification;
- a resource updates the version of itself and changes the ID of a resource version once the contents of a resource change.

The best way to demonstrate components of the FHIR resource is to look at its encoded example.

METADATA


```
{
  "resourceType": "Patient",
  "id": "example",
```


NARRATIVE


```
"text": {
   "status": "generated",
   "div": "<div xmlns=\"http://www.w3.org/1999/xhtml\"><p style=\"border: 1px #661aff
solid; background-color: #e6e6ff; padding: 10px;\"><b>Jim </b> male, DoB: 1974-12-25 (
Medical record number: 12345 (use: USUAL, period: 2001-05-06 --&gt;
(ongoing)))</p><hr/><table class=\"grid\"><tr><td style=\"background-color: #f3f5da\"
title=\"Record is active\">Active:</td><td>true</td><td style=\"background-color:
#f3f5da\" title=\"Known status of Patient\">Deceased:</td><td
colspan=\"3\">false</td></tr><tr><td style=\"background-color: #f3f5da\" title=\"Alternate
names (see the one above)\">Alt Names:</td><td colspan=\"3\"><ul><li>Peter James
Chalmers (OFFICIAL)</li><li>Peter James Windsor (MAIDEN)</li></ul></td></tr><tr><td
style=\"background-color: #f3f5da\" title=\"Ways to contact the Patient\">Contact
Details:</td><td colspan=\"3\"><ul><li>-unknown-(HOME)</li><li>ph: (03) 5555
6473(WORK)</li><li>ph: (03) 3410 5613(MOBILE)</li><li>ph: (03) 5555
8834(OLD)</li><li>534 Erewhon St PeasantVille, Rainbow, Vic
3999(HOME)</li></ul></td></tr><tr><td style=\"background-color: #f3f5da\"
title=\"Nominated Contact: Next-of-Kin\">Next-of-Kin:</td><td
colspan=\"3\"><ul><li>Bénédicte du Marché  (female)</li><li>534 Erewhon St PleasantVille
Vic 3999 (HOME)</li><li><a href=\"tel:+33(237)998327\">+33 (237)
998327</a></li><li>Valid Period: 2012 --&gt; (ongoing)</li></ul></td></tr><tr><td
style=\"background-color: #f3f5da\" title=\"Patient Links\">Links:</td><td
colspan=\"3\"><ul><li>Managing Organization: <a
href=\"organization-example-gastro.html\">Organization/1</a>
&quot;Gastroenterology&quot;</li></ul></td></tr></table></div>"
  },
```

```
"identifier": [
  {
    "use": "usual",
    "type": {
      "coding": [
        {
          "system": "http://terminology.hl7.org/CodeSystem/v2-0203",
          "code": "MR"
        }
      ]
    },
    "system": "urn:oid:1.2.36.146.595.217.0.1",
    "value": "12345"
  }
 ],
 "name": [
  {
    "use": "official",
    "family": "Chalmers",
    "given": [
      "Peter",
      "James"
    ]
  }
 ],
 "telecom": [
  {
    "system": "phone",
    "value": "(03) 5555 6473",
    "use": "work",
    "rank": 1
  }
 ],
 "gender": "male",
 "birthDate": {
```


STANDARD DATA


```
  "extension": [
    {
      "url": "http://hl7.org/fhir/StructureDefinition/patient-birthTime",
      "valueDateTime": "1974-12-25T14:35:45-05:00"
    }
  ]
```


EXTENSION


```
 },
 "address": [
  {
    "use": "home",
    "type": "both",
    "text": "534 Erewhon St PeasantVille, Rainbow, Vic  3999",
    "line": [
      "534 Erewhon St"
    ],
    "city": "PleasantVille",
    "district": "Rainbow",
    "state": "Vic",
    "postalCode": "3999",
    "period": {
      "start": "1974-12-25"
    }
  }
 ],
 "managingOrganization": {
  "reference": "Organization/1"
 }
}
```


STANDARD DATA

Metadata:

includes the resource ID, version ID, and the date of the last resource update.

Narrative:

is an HTML representation of a resource's contents that appears in a human-readable form. This part allows displaying the data in the narrative even if the system cannot process the coded standard data.

Extension:

the element that expands resource capabilities and makes it match a particular use case. Implementers can use an extension with use cases that don't enter the common scenarios category (80:20 rule). The extension consists of key-value pairs—one of the simplest ways to display data.

Standard data:

data in this block covers 80% of use cases. For example, the standard data block of a patient resource includes the medical record number, name, gender, date of birth, and other essential attributes of a resource. The raw structure data in this block is easy to read as well as data in the narrative.

The FHIR (version 4.0.1) includes 153 resources gathered in the resource list.

UNIT II. FHIR TECHNICAL COMPONENTS

2. FHIR RESTful API

How an API works

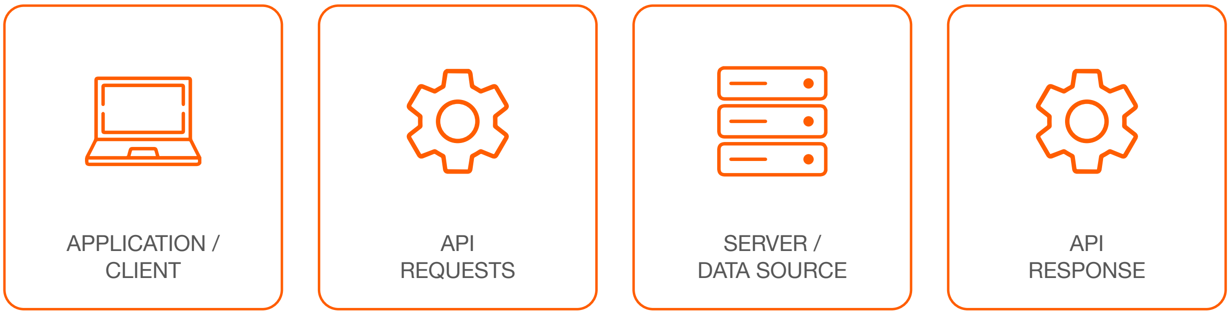

Representational State Transfer

REST is a simple way of sending and receiving data between client and server and it doesn't have very many standards defined. You can send and receive data as JSON, XML or even plain text. It's light weighted compared to SOAP.

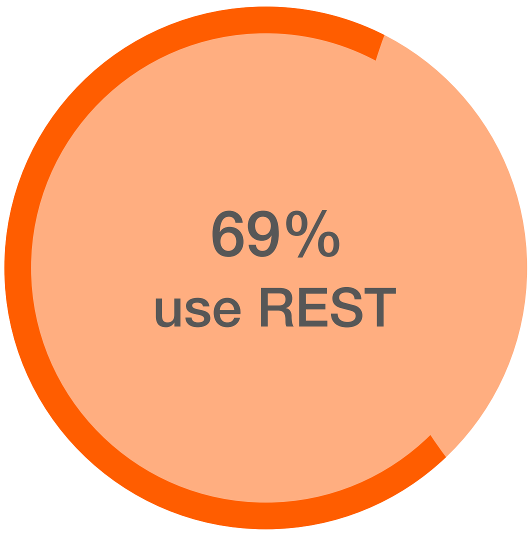

REST architecture style uses HTTP data access requests. RESTful APIs are ones that use the REST exchange standard. FHIR RESTful API uses a simple set of operations for manipulating resources:

CREATE:

used to create a resource, assign it to a server, and get the resource ID from that server. HTTP command for Create operation: POST [base]/[type] {?_format=[mime-type]}

URL example: http://kodjin.example.com/Patient

READ:

used for pulling out an existing resource using its ID. HTTP command for Read operation: GET [base]/[type]/[id] {?_format=[mime-type]}

URL example: http://kodjin.example.com/Patient/94

UPDATE:

this interaction creates a new version of a resource. HTTP command for Update operation: PUT [base]/[type]/[id] {?_format=[mime-type]}

URL example: http://kodjin.example.com/Patient

DELETE:

is used to remove a resource (as an aspect of a search result—the "deleted" data remains on the server). HTTP command for Delete operation: DELETE [base]/[type]/[id]

URL example: http://kodjin.example.com/Patient/94

SEARCH:

establishes criteria for searching resources. HTTP command for Search operation:
GET[base]/[type]{?[parameters]{&_format=[mime-type]}}

URL example: http://kodjin.example.com/Patient/94

HISTORY:

retrieves the history of all versions of a particular resource or all resources. HTTP command for History operation:
GET [base]/[type]/[id]/_history{?[parameters]&_format=[mime-type]}

URL example: http://kodjin.example.com/Patient/aaa145-56/_history

REST architecture style request defines the types and volume of data the application needs. Leveraging proven web-based technologies such as RESTful API ensures a high level of adoption and data interoperability.

UNIT II. FHIR TECHNICAL COMPONENTS

3. FHIR Implementation Guides

HL7 defines an implementation guide (IG) as a set of rules about how FHIR resources are used (or should be used) for a particular context, with associated documentation to support and clarify the usage.

Implementation guides consist of two types of resource references:

Contents:

sets of logical statements implementers should conform to. Contents are conformance resources, the purpose of which is to specify which resource elements can and can't be used, the terminology used in specific elements, features of the FHIR RESTful API, and the way they map to local requirements. Conformance resources refer to the Capability Statement, which defines the required capabilities of an FHIR application (server) for a particular version of FHIR.

Example:

shows the meaning of the IG. It describes the meaning of the contents and how it affects a particular use case.

The FHIR implementation guides include detailed guidelines on how to use FHIR resources.

A guide is a complete set of base resources, frameworks, and APIs that help to deal with specific scenarios. The base FHIR resources cover a specific use case with extensions and constraints added. The process of extending and restricting a resource is called profiling. A profile is the essence of any FHIR implementation guide.

FHIR specification provides an Implementation Guide Registry, where you can search for the existing implementation guides.

Implementation Guides Registry

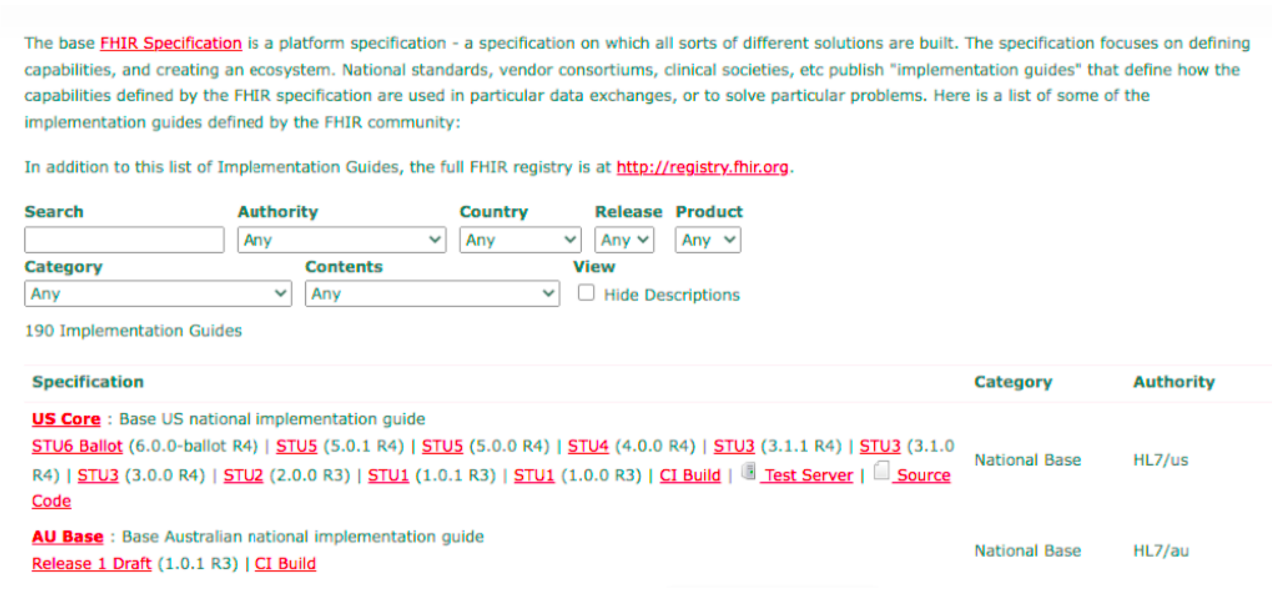
The base FHIR Specification is a platform specification - a specification on which all sorts of different solutions are built. The specification focuses on defining capabilities, and creating an ecosystem. National standards, vendor consortiums, clinical societies, etc publish "implementation guides" that define how the capabilities defined by the FHIR specification are used in particular data exchanges, or to solve particular problems. Here is a list of some of the implementation guides defined by the FHIR community:

In addition to this list of Implementation Guides, the full FHIR registry is at http://registry.fhir.org.

Search | Authority: Any | Country: Any | Release: Any | Product: Any

Category: Any | Contents: Any | View: Hide Descriptions

190 Implementation Guides

Specification	Category	Authority
US Core : Base US national implementation guide STU6 Ballot (6.0.0-ballot R4) | STU5 (5.0.1 R4) | STU5 (5.0.0 R4) | STU4 (4.0.0 R4) | STU3 (3.1.1 R4) | STU3 (3.1.0 R4) | STU3 (3.0.0 R4) | STU2 (2.0.0 R3) | STU1 (1.0.1 R3) | STU1 (1.0.0 R3) | CI Build | Test Server | Source Code	National Base	HL7/us
AU Base : Base Australian national implementation guide Release 1 Draft (1.0.1 R3) | CI Build	National Base	HL7/au

FHIR also provides a free tool to write implementation guides, FHIR IG Publisher. This tool suggests using the master source for the implementation guide from a public GitHub repository.

One of the most prominent examples of the implementation guide is the US Core Implementation Guide.

The purpose of the US Core Implementation Guide was to establish healthcare interoperability in the United States and allow third-party applications to connect to EHRs.

The US Core IG is massive, covers many use cases, and includes lots of helpful documentation. This IG can be the one to refer to when building an implementation guide.

UNIT II. FHIR TECHNICAL COMPONENTS

4. FHIR Profiling

Why do we need profiles?

Profiling is a necessary procedure for those who need to adapt core FHIR specifications to a specific context. FHIR resources define 80% of use cases, and profiling helps define rules for 20% of use cases, the context of which goes beyond FHIR core specifications.

The main goal of profiling is to make FHIR work in the context of any particular use case. Profiling performed on resources is defined in the Structure Definition resource, which contains rules about how the elements in a resource are used and where extensions are used in a resource.

FHIR Profile components:

- a profiling goal, profile requirements, an existing resource that seems most relevant for certain use cases;
- snapshot: all the data from a resource and a profile in a fully calculated form of the structure;
- a differential: constraints and extensions details.

Profiling steps:

1	2	3
Define the constraints	Change the ValueSet	Add the extension

FHIR Conformance

The Conformance describes the resources and API features of the FHIR necessary for profiling.

There are four types of Conformance resources:

Operations resources define the behavior of the FHIR server. The group includes:

- **CapabilityStatement** - defines acceptable API calls for different types of resources;
- **OperationDefinition** - defines what CRUD operations are used for interaction with resources;
- **SearchParameter** - defines all FHIR-supported search parameters.

Terminology is a resource-related vocabulary.

- **CodeSystem** - defines a meaning of a set of codes;
- **ValueSet** - defines values accepted for a system of a coded element;
- **ConceptMap** - defines relationships between concepts in different code systems.

Content

- **StructureDefinition** - allows users to create resources and define profiles and constraints. All FHIR resources are StructuredDefinition resources;
- **DataElement** - a block of StructureDefinition resource that includes characteristics of a piece of data, such as identifier, name, data type, permitted value set, and more;
- **Extension** - another StructureDefinition element composed of an URL with the extension definition and a value.

Publishing

ImplementationGuide - includes a set of resources needed for publishing (ValueSet, StructureDefinition, OperationDefinition, and other resources).

Resource validation against a profile

Validation is an important step of FHIR specification that checks a resource's compliance with a specific profile to ensure the accurate representation of healthcare reporting.

Validation types:

- Structure ensures that the content of a resource is described by the specification;
- Cardinality checks that the minimum and maximum cardinality values for all properties are correct;
- Slicing creates multiple instances with defined restrictions for elements that may appear more than once;
- Binding FHIR binds ValueSet to the element, and binding allows users to check if the codes provided in coding data types are valid;
- FHIRPath a language used for FHIR resource navigation and data extraction (for example, interlock the data from a resource, extract it, and make it into the text to create a narrative).

Validation methods:

- Using schemas allows validation of resources represented in XML and JSON;
- FHIR Validator supports XML files only. Validator is a Java jar provided as part of the specification, which allows users to validate examples during the publishing process;
- Using $validate operation $validate operation is used to ask the FHIR server if a resource conforms to a profile.

Usually, profiling involves manually editing JSON files in the text editor, Visual Studio, or any other integrated development environment. It is a time-consuming process, which makes it difficult to modify the structure of a profile and validate the code. Powerful FHIR profiling tools will simplify the profiling process and remove the risk of human error.

Profiling tools

- Kodjin FHIR Profiler automates definition of the value set, resource snapshot generated from differential descriptions, FHIR profile validations, and other aspects of profiling;
- Forge by Firely a tool for building FHIR profiles, extensions, and data models;
- Trifolia-on-FHIR (ToF) FHIR resource editor that uses an FHIR server natively as its backend.

On our blog, you can find more information regarding FHIR profiling tools and recommendations on how to choose and use them.

Profiling resources

- Profiling FHIR detailed FHIR profiling documentation provided by HL7. It includes all the needed information to create a profile;
- Guidance on writing FHIR profiles a complete profiling guide made by healthcare experts (Integrating the Healthcare Enterprise members);
- Simplifier.Net a registry of FHIR resources.

A profile is a resource that makes FHIR flexible. Any profile must be followed with an implementation guide in order to specify terminology and attributes that can be used in a particular context.

UNIT II. FHIR TECHNICAL COMPONENTS

5. FHIR Servers

FHIR servers support the processing, validation, and storage of healthcare data in accordance with the FHIR specification. In simple terms, FHIR is the standard, while the FHIR server is its implementation. By leveraging FHIR servers, healthcare organizations can avoid data silos and ensure prompt access to the needed data.

Why do healthcare organizations need the FHIR server?

For a long time, accessing healthcare information was quite a complex process. The conflict between different data formats and IT systems caused many inconveniences to all healthcare actors. While healthcare data exchange regulations attempted to solve this issue, oftentimes, they only made this process more difficult.

Today, governments and most healthcare organizations realize the importance of data interoperability and exert all possible efforts to achieve it. Therefore, every healthcare stakeholder should consider using an FHIR server to transform their IT system into a modern standard-ready healthcare solution.

An FHIR server allows participants of the healthcare ecosystem to benefit the most from the FHIR standard and use FHIR resources to the maximum.

FHIR API methods:

- iServeInstance: for performing GET, PUT, and DELETE interaction of a resource;
- iServeTyp: for obtaining type information or metadata about resources;
- iServeSystem: for exposing or enabling system behaviors.

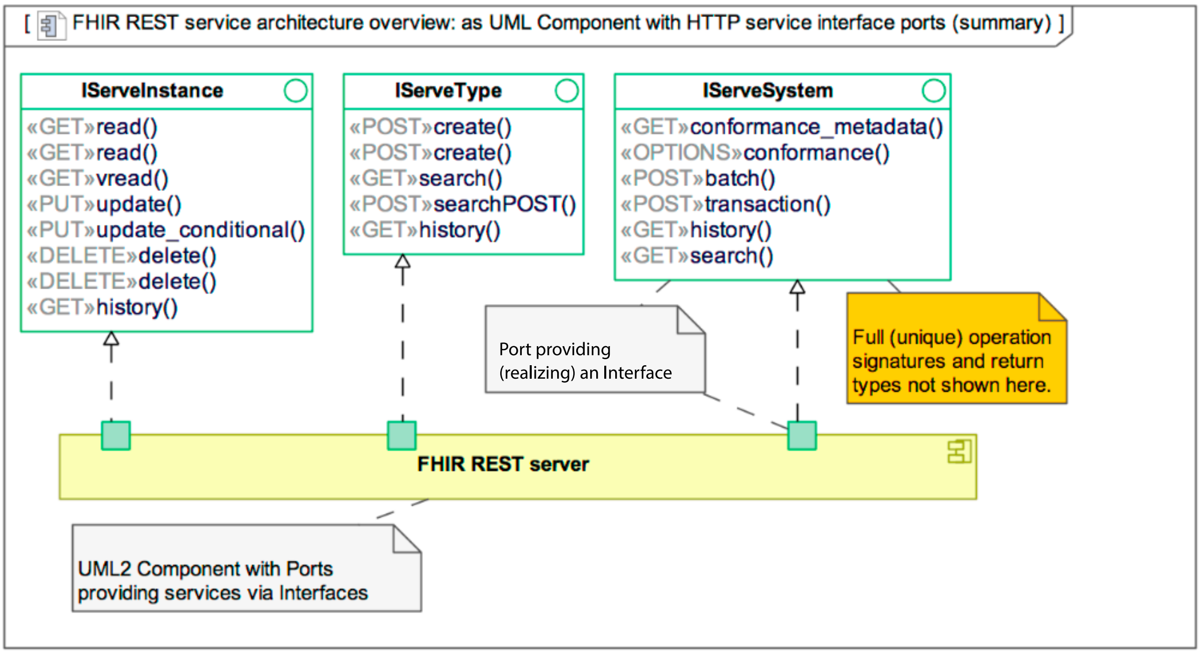

Source: https://confluence.hl7.org/display/FAST/FAST+Projects

Functionality of the FHIR server:

- Supports the following data formats: XML, JSON, TTL;
- Storage for resources;
- Define behaviors/capabilities between the FHIR server and FHIR client;
- Support search and specify search parameters;
- Validation (ensures conformity of FHIR resources to base resource requirements);
- Transactions: a set of requests to FHIR server packed in a batch;
- $versions operation (allows user to define custom operations).

Key features to look for in an FHIR server:

- Microservice architecture: allows scalability, which is crucial for high-load projects;
- Asynchronous event-driven approach: allows all API requests;
- Compliant with FHIR R4 with backward compatibility;
- Low-code approach for smooth software development.

Kodjin FHIR Server

The Kodjin FHIR server is loaded with all the mentioned features and transforms the ideas underlying the FHIR standard into the interoperability-achieving mechanism. It uses a low-code approach and allows users to get the most advantages out of FHIR specification, even within high-load systems. The Kodjin FHIR server is a custom FHIR solution created in response to healthcare data management needs. This server is tested by 36,5M+ users and allows for solving national-level business cases.

DOWNLOAD WHITEPAPER

HIPAA COMPLIANT

ONC Certified HIT
2015 EDITION CURES UPDATE

UNIT II. FHIR TECHNICAL COMPONENTS

6. FHIR Integration Process

There is no generic mechanism for FHIR paves adoption. It always depends on the use case, so the development of the FHIR implementation guide should be based on a thoughtful analysis of system requirements. However, there are common steps any project will follow when implementing FHIR.

Explore interoperability rules

Despite being a relatively new standard, FHIR features lots of valuable resources describing the importance of leveraging interoperability standards and what specifications work for each use case. Before starting with FHIR, make sure you explore the following resources:

- HL7 FHIR website;
- ONC's Cures Act Final Rule;
- Getting Started with FHIR;
- Guidance for FHIR IG Creation;
- The official source for the HL7 FHIR Specification.

Create an integration plan

Information in the above-mentioned sources helps define which elements of FHIR will be useful for a particular use case. Create a summary of a system's needs and describe workflows, define triggers of data exchange, and input and output data elements required for information exchange.

This research ensures smooth integration and effective use of FHIR in the context of core data needs within a system.

There are two common ways to integrate FHIR API:

Options	FHIR Facade	FHIR Server
Advantages	Allows a system to present an FHIR interface to the outside world, even though it doesn't store data in an FHIR format.	Offers full FHIR compatibility and is ideal for organizations that can adopt a more modern, FHIR-centric infrastructure.
Workflow	The FHIR facade receives FHIR requests via an API, translates them into a different internal structure or format, processes them, and then converts the results back to FHIR for external use.	Kodjin uses an event-driven approach to maximize performance and efficiency. The server utilizes a powerful message broker, Apache Kafka, which ensures seamless communication between microservices. Kafka handles task queues, ensuring data creation and search processes are interconnected and efficiently managed.

FHIR facade on top or a source

The FHIR facade is used on top of the existing backend. It allows addressing the source using the FHIR API.

A Facade is deployed in cases where setting up a FHIR Server is not possible, for instance, when:

- A laboratory has established workflows and requires the capability to integrate with clinics via FHIR;
- A clinic needs to integrate with a patient portal;
- Multiple healthcare providers need to integrate in order to deliver a unified service to patients;
- A clinic requires integration with an insurance company.

A FHIR Facade offers a FHIR API for external systems without storing the data in FHIR format. It receives FHIR requests through the API and then transforms and processes the data.

Native FHIR server

In this type of integration, FHIR acts as the central design component of the system. All components of the native FHIR server are compatible with FHIR by default, meaning that from its inception, a native FHIR server was created to support FHIR specifications — resources, profiles, and extensions.

FHIR Server is an out-of-the-box solution for those who require FHIR compliance. It provides both an FHIR API and a repository for storing data in FHIR format.

An unmanaged server requires a customer to provision everything themselves, such as setting up the server and database, scaling it, creating and managing indexes, etc. Such an enterprise-level server would require a separate dedicated team to maintain it. Whereas, with a managed solution, a service provider sets the server up and deals with the technical side of things.

FHIR server performance

For large systems and high-load projects, performance is the most important aspect when choosing an FHIR server. The daily HL7 message in an extensive system can reach up to 8M messages per day. One message involves 15 FHIR resources, so an organization needs to process 120,000 resources per minute of EHR streaming.

For high performance, an FHIR server should have the capacity to create hundreds of thousands of resources per minute. FHIR server performance often depends on the type of infrastructure it's built on, cloud or on-premises. Each type has its benefits and specific use cases. Below, we'll examine both of them to help you determine which is better suited to your needs.

*Note: **contact us** via the website form to explore Kodjin FHIR Server Performance. We created this solution specifically for enterprise-level projects.*

Cloud providers

There are three major cloud providers: AWS, Google Cloud, and Microsoft Azure. Public clouds provide computing resources on demand. The clouds allow scaling storage and resources if needed, have excellent encryption, and require payments only for resources that are actually used. The big cloud providers are usually tested with dozens of millions of patient records and have proven high-level performance.

On-premises

The on-premises infrastructure virtualizes resources the same way as a cloud. However, it requires an extensive team of database administrators to conduct additional work to achieve horizontal scale, which makes this approach more costly.

The Kodjin server can be used in the cloud and on-premises. The microservice architecture of the server ensures extraordinary scalability and allows it to manage data for systems with over 36,5M+ patient records, even when used with on-premises data centers.

Proper preparation for FHIR implementation helps evaluate the costs and define the goals of a project.

Adopt and test FHIR specifications within a system

An FHIR adoption roadmap based on thoughtful system analysis will help you switch the data processing workflow into FHIR-first mode. Use implementation guides to introduce the new data processing rules to your system.

Testing is an essential part of FHIR specification implementation to ensure that information about patients is accessible across all applications and devices. In addition, testing is vital to bring out FHIR implementation issues and fix them swiftly.

FHIR implementation requires deep analysis of data exchange mechanisms within your system to define steps to make it match FHIR requirements. Our team can examine your system and create the roadmap for smooth FHIR implementation.

UNIT II. FHIR TECHNICAL COMPONENTS

7. Our Solutions and Services

Edenlab provides cutting-edge solutions for achieving interoperability, including backend development services, FHIR-based tools for healthcare organizations, and FHIR facade building.

KODJIN FHIR SERVER

A turnkey event-driven FHIR server for healthcare data storage, data exchange, and management

- Microservice architecture;
- Dynamic profiling;
- Declarative search framework;
- Full terminology management throughout a REST API;
- SMART on FHIR out of the box;
- ONC-certified;
- HIPAA-compliant;
- Tested by high-load systems with up to 36,5M+ patient records.

KODJIN TERMINOLOGY SERVICE

An FHIR service to handle all types of healthcare terminology

- Can process complex terminologies;
- Supports standard and proprietary terminology systems;
- Low-code approach that helps time and cost efficiency;
- Built with FHIR at its core.

DATA MAPPER

An interoperability FHIR software to map and transfer healthcare data elements with liquid template language

- High-speed and robust performance tool;
- Designed to work with proprietary healthcare data;
- Mapping data in real time;
- Out-of-the-box support for conversion of HL7 V.2 and CDA to FHIR.

KODJIN FHIR PROFILER

A free FHIR software to create FHIR profiles faster, by using an intuitive graphical IDE

- Syntax and structure control;
- ValueSet bindings;
- FHIRPath constraints;
- Auto-generates profile snapshot from the differential declaration;
- FHIR profile validation;
- Visual profile tree render.

FHIR-BASED CUSTOM DEVELOPMENT

FHIR FACADE

An FHIR interface that helps your data become FHIR-compliant

- Cost-efficient;
- Supports all types of data structures and terminologies;
- Easily integrated with new medtech software and third-party platforms.

These tools already support large healthcare projects, such as the National Ukrainian eHealth System with over 36,5M patient records. eHealth system provides secure storage and barrier-free access to verified data about all participants in the Ukrainian healthcare ecosystem (providers, physicians, patients, etc.) To date, this project remains one of the biggest FHIR implementations in the world.

LINKS

1. The official web page of Edenlab
2. HL7 Version 2
3. HL7 Version 3
4. Total national health expenditures, US $ Billions, 1970-2020
5. American Recovery and Reinvestment Act of 2009 (ARRA)
6. Health Information Technology for Economic and Clinical Health Act (HITECH)
7. Trends in Hospital & Physician EHR Adoption
8. Certification of Health IT
9. CMS EHR Meaningful Use (MU) Incentive Payment Program
10. 21st Century Cures Act
11. The Da Vinci Project
12. Interoperability and Patient Access Final Rule
13. HL7 FHIR v4.0.1.
14. Percent of hospitals and clinicians that adopted or could adopt (but did not) 2015 Edition certified API technology enabled with FHIR, 2019
15. Base Resource Definitions
16. Resource Index
17. RESTful API
18. Implementation Guide
19. Conformance resources
20. Capability Statement
21. Implementation Guide Registry
22. FHIR IG Publisher
23. FHIR public GitHub repository
24. US Core Implementation Guide
25. StructureDefinition
26. Kodjin profiling tools
27. Profiling FHIR
28. Guidance on writing FHIR profiles
29. Integrating the Healthcare Enterprise
30. Simplifier.Net
31. Why Choose an event-driven architecture for FHIR servers
32. The Kodjin FHIR server
33. HL7 FHIR website
34. ONC's Cures Act Final Rule
35. Getting Started with FHIR
36. Guidance for FHIR IG Creation
37. The official source for the HL7 FHIR Specification
38. Why HL7 Message Throughput Is a Critical Metric for Interface Engines
39. AWS
40. Google Cloud
41. Microsoft Azure
42. Kodjin FHIR facade
43. Kodjin terminology service
44. Kodjin FHIR Profiler
45. National eHealth system