The Office of the National Coordinator for Health Information Technology

What Is HL7® FHIR®?

HL7® FHIR®[1] – Fast Healthcare Interoperability Resources is a next-generation interoperability standard created by the standards development organization Health Level 7 (HL7®). FHIR is designed to enable health data, including clinical and administrative data, to be quickly and efficiently exchanged.

Why was FHIR created?

In 2012, a team of health information technology implementers lead by the inventor of FHIR, Grahame Grieve, asked the question, "What would health information exchange look like if it started now, using modern approaches?" This question factored in the rapidly growing amount of health data and the rise of the “app” economy on smartphones. The team created a draft standard that built on the simplicity of the main method of exchange at the time, HL7 v2 messages, combined with an application programming interface (API) and common World Wide Web technologies including JSON, XML, HTTP, and OAuth. These technologies power all kinds of internet-based data exchange and are used by e-commerce providers and social media companies such as Kayak, Mint, and Google.

By adopting existing standards and concepts already familiar to software developers outside of health care, FHIR reduces the learning curve, makes real time interoperability easier, and enables faster and simpler application creation.

A core goal of FHIR was to create a standard that would lead to high adoption across disparate developer communities. Therefore, it is focused on being easy for software developers to use.

By adopting existing standards and concepts already familiar to software developers outside of health care, FHIR reduces the learning curve, makes real time interoperability easier, and enables faster and simpler application creation.

Current health care systems can co-exist with and use FHIR. Its evolutionary development path from standards based on Messages, such as HL7 v2, and Documents, such as Clinical Document Architecture (CDA), enables continuity with existing provider workflows.

1. http://hl7.org/fhir

Core capabilities of FHIR

FHIR's API[2] is a RESTful, or REpresentational State Transfer, approach to data exchange. REST defines categories of data, or "Resources [3]," to exchange data. The philosophy behind FHIR is to create a set of Resources that, individually or in combination, satisfy most common use cases. The Patient Resource, for example, includes demographic data related to a patient, such as their name, address, and phone number. Resources also improves granular data retrieval, so that a request returns just the relevant data rather than a full record or document that itself must then be searched.

The philosophy behind FHIR is to create a set of Resources that, individually or in combination, satisfy most common use cases.

Once they are modified for specific requirements using FHIR's built-in capabilities, combinations of Resources are brought together in an Implementation Guide to address a specific use case, such as a provider directory or patient-reported outcomes. This structure lends itself well to expansion beyond FHIR's core capabilities.

Like many other components in the standard, FHIR uses modern security standards, including for authentication and encryption. Similarly, among FHIR's privacy capabilities, FHIR can support labeling sensitive information so that only those who have the need and the right can see it.

Health care data is represented by many sets of vocabularies, terminologies, and codes which grow and change over time. As a result, it is important for the data exchanged to be equally understood by the sender and receiver, which is known as "semantic interoperability." FHIR manages the use of this data by including references to code definitions used for data verification and by allowing restrictions on the codes that can be used.

2 See "The FHIR® API" Fact Sheet
3 See "Introduction to FHIR® Resources" Fact Sheet

Interested in learning more about FHIR?

See more of our Fact Sheet series at:
https://www.healthit.gov/topic/standards-technology/standards/fhir-fact-sheets.

ONC's FHIR Fact Sheets are a collaborative effort with HL7 to help educate and demystify FHIR for federal employees. These fact sheets summarize the key technical concepts that make up the foundation of FHIR, how it is developed in an open and public process, and why FHIR adoption has become the focus of the health IT standards world. Full details and developer documentation can be found at HL7's FHIR website.

The Office of the National Coordinator for Health Information Technology

The FHIR® API

HL7® FHIR® includes specifications for an Application Programming Interface, or API, based on established web standards and modern information exchange that has been extended to create a full interoperability solution for health care. The use of common and widespread technologies lowers the barriers to entry for using the standard by making it easier and faster to implement while also opening up development to individuals who do not have health care expertise.

What is an API?

An API is an entry point, or "interface," that allows a computer program or system to access the features and data of a different program or system. This entry point defines how data must be formatted and the types of interactions supported, such as how data can be searched. To be successfully exchanged and ready for an operation, data must be formatted in the same way. For example, a data field to be used in calculations cannot accept both "1" and "one" as entries, as they cannot be interpreted the same way. APIs can be simple, with rigid definitions of data structure, or complex, with different data structures for a wide array of interactions. The FHIR API mainly involves the access and exchange of data.

Many modern applications, both desktop and mobile, use APIs to retrieve, store, and update data. For example, a smartphone banking app requests a user's data from a financial institution through the institution's API. Each institution's API defines the standards and protocols that allow an external application to access the institution's data. With the proper security procedures and protocols, the app can successfully request the user's data. Ultimately, the user of the banking app can view their data, from accounts at multiple institutions, in one place.

What is REST?

Many applications run on a mobile device or web browser use the information exchange standard REST (**Re**presentational **S**tate **T**ransfer) as the basis for their APIs. REST is a method of exchanging information using the World Wide Web standard transfer protocol HTTP[1] , the underlying internet standard that forms the basis for all website data exchange. "Http" can be seen at the beginning of every webpage's web address, such as "**http:**//www.healthit.gov".

REST means that each request from any client and response from the server contains all the information necessary to service the request. The exchange of data using REST is termed a "RESTful" exchange. A RESTful Resource creates a way to access data through a specific endpoint, using specific structures and formats. More strictly defined structures and formats enable more-precise data exchange. In FHIR, which uses such strictly defined structures, a system can target and retrieve just a single data element rather than receive a document containing a patient's full record.

1. Hypertext Transfer Protocol

To initiate data exchange through an API, the requesting client, such as a smartphone app or server, first sends a request via HTTP for information. If the information is available, the remote server responds with a status code indicating success and returns the information in a standard web format. For example, when a user searches for a location in Google Maps, which uses REST, the web browser is the client. A request in the web browser looks like the following:

https://www.google.com/maps/place/1500+S+Capitol+St+SE,+Washington,+DC+20003

That tells Google a specific location is being searched for in the Maps system. The request is broken down as follows:

https://www.google.com/maps – the receiving system (Google) and application (maps)

/place/ – the Resource (place) that is the end point being queried

1500+S+Capitol+St+SE,+Washington,+DC+20003 – the address the user is interested in

Google's Maps application accepts this request and sends back the data allowing the browser to display the map for that area. While that functionality determines what a user experiences, many RESTful requests are made in the background in response to the initial request, including loading the basic map interface, displaying the many map elements, and showing notations for other businesses in the area.

The FHIR RESTful API

FHIR uses REST as the basis for data exchange in its API. Health care data types such as medications, observations, and patients are represented by their own Resources. Resources can be requested via a RESTful HTTP command like the Google Maps example above, in addition to interactions like searches or requests that can be used to find and retrieve the precise information needed. Servers, like those behind an electronic health records system, are programmed with the types of Resources and interactions they can support. Third-party applications using the FHIR API can be integrated into an EHR and feed information directly into the provider's workflow.

Each request using the FHIR API supplies the Resource and an indicator, command, or parameter that specifies the data needed. The simplest FHIR request returns one Resource of information, like a single Patient. In addition, a request could return a bundle of information, such as that Patient's associated CarePlan and Medications, or a bulk data bundle, like all data on all Patients in an EHR. The request is structured to tell the application what types and how much of the data are needed.

REST does not directly address privacy and security elements such as authentication and authorization, which are addressed separately in the FHIR standard.

By using the REST architectural style, FHIR takes the best of existing health information technology and common internet standards to create a modern method of interoperability. This allows health care systems to implement FHIR without steep learning curves and leading to faster application design.

The Office of the National Coordinator for Health Information Technology

FHIR® Version History and Maturity

HL7® FHIR® has evolved through four releases since its initial presentation in May 2012. It has grown from a true draft standard with 49 Resources to its current 145 and continues to expand. In that time the standard has improved and changed to meet the needs of the health information technology community.

Draft Standard for Trial Use 1 (DSTU1)

FHIR's first publication in September 2013 showed a new way forward for health care data exchange. Draft Standard for Trial Use 1 had 49 resources and focused on two use cases, creating a Personal Health Record on a mobile device and the retrieval of documents, such as encounter or discharge notes, to a mobile device. This initial release sparked the community's interest in expanding FHIR to cover a wider variety of health care and health IT needs.

Draft Standard for Trial Use 2 (DSTU2)

FHIR grew in market acceptance with the publication of the Draft Standard for Trial Use 2 in 2015. Efforts including the Argonaut Project developed Implementation Guides (IGs) and other technologies to support FHIR adoption by EHR developers and other health IT entities. The structure of Resources was adjusted to make creating extensions easier, allowing for more use cases to be covered without changes to the core standard. New Resources were also added to support non-clinical data, including claims and benefits processing.

The publication of FHIR DSTU2 included the creation of the FHIR Maturity Model (FMM). When new Resources are created, they are not immediately ready for use in live settings; they must be refined and tested for a variety of uses and settings. The FHIR Maturity Model established a set of levels that progressively measure technical advancement, known as maturity. Resource maturity as defined by the FMM begins with an initial draft and achieves final status with implementation in multiple settings. Since the maturity of the FHIR standard overall is not tied to the maturity of Resources, Resources can move up the maturity ladder between FHIR releases. The FMM, which is also applied to other components of the FHIR standard, defines Resource stability with six levels:

- **FMM0 (Draft)** – The resource is still in early development but has been accepted into the FHIR standard.
- **FMM1** – The Resource has no current technical errors and is believed to address all design goals.
- **FMM2** – The Resource has been tested and approved at a FHIR Connectathon with multiple FHIR-enabled computer systems tested.
- **FMM3** – The Resource passed all quality checks and an HL7 community ballot that determines if it is ready for trial use.
- **FMM4** – The Resource has been tested for functionality for all intended purposes, has been published in a formal HL7 publication, and is operating in at least one prototype system.
- **FMM5** – The Resource is in use in at least five distinct production systems operating in at least two countries.

Substantive changes at the FMM4 or FMM5 levels that would change usage from those already established or would break compatibility with existing implementations would require significant justification to be accepted and to move forward. After FMM5, a Resource reaches "normative"[1] level; at this level, future changes must be backwards compatible so that applications that implement those Resources aren't at risk of being broken as the FHIR standard changes.

Standard for Trial Use 3 (STU 3)

FHIR Standard for Trial Use 3 was released in 2017 with improvements to the core Clinical, Administrative, and Financial Resources, improvements to the Clinical Decision Support and Clinical Quality Measure Resources and a new framework for workflow and task management. Additionally, tools were introduced that made FHIR IG creation and publication to the web easier, faster, and more unified.

Release 4 (R4)

As the first release with normative content, the 2019 release of FHIR Release 4 left behind the Trial Use name. Two key clinical Resources, Patient and Observation, were released as normative, along with the RESTful API, the XML and JSON formats, and nine additional Resources.

In 2020, ONC published the Final Rule for the 21st Century Cures Act, establishing FHIR R4 as the standard required for Health IT Certification.

Looking ahead to Release 5 (R5)

FHIR Release 5 will see increased normative content, with over 30 Resources having been nominated by their HL7 Workgroups to be matured to that status. In addition, the community will continue to develop the supportive specifications to FHIR, such as the authorization framework SMART, Clinical Decision Support Hooks (CDS Hooks), and the Bulk Data Transfer specification, which will help implementers create a complete FHIR-based exchange of health care data.

With the maturing of the FHIR IG tools and templates, better integration with public health, imaging, financial management, genomics and other fields will keep FHIR at the forefront of health IT.

1. Normative, in this context, is defined as content that has been accepted by the American National Standards Institute (ANSI) as an American National Standard.

Interested in learning more about FHIR?

See more of our Fact Sheet series at:
https://www.healthit.gov/topic/standards-technology/standards/fhir-fact-sheets.

ONC's FHIR Fact Sheets are a collaborative effort with HL7 to help educate and demystify FHIR for federal employees. These fact sheets summarize the key technical concepts that make up the foundation of FHIR, how it is developed in an open and public process, and why FHIR adoption has become the focus of the health IT standards world. Full details and developer documentation can be found at HL7's FHIR website.

The Office of the National Coordinator for Health Information Technology

Introduction to FHIR® Resources

The core of HL7® FHIR® is a set of modular components called "Resources." These form the basic data exchange format and model of FHIR. As of FHIR Release 4, there are 145 Resources defined across health care domains and supporting services, a number that grows with every release.

What are Resources?

In FHIR, health care data is broken down into categories such as patients, laboratory results, and insurance claims, among many others. Each of these categories is represented by a FHIR Resource, which defines the component data elements, constraints on data, and data relationships that together make up an exchangeable patient record. The philosophy behind FHIR is to create a set of Resources that, individually or in combination, satisfy most common use cases.

Each Resource contains data elements necessary for its specific use cases and links to relevant information in other Resources. For example, the Patient Resource[1] contains basic patient demographics, contact information, and links to a clinician or organization stored in a different Resource. Because they are based on modern World Wide Web technologies, Resources use Uniform Resource Locators, or URLs (also generally known as web addresses), to be located within a FHIR system implementation.

A Resource in its raw form does not require that most data elements be assigned a value; when it is customized for real-world use through the Profiling process, certain elements are then required so that the Resource can be functional. For example, a Patient Resource may be Profiled to require that a patient's name, address, and telephone number be supported to enable patient matching.

FHIR Resource Design Choices

A primary motivation behind FHIR's design is to enable interoperability through well-structured data models that use simple and efficient exchange mechanisms. To achieve this, FHIR adopted the following principles:

Reuse – FHIR Resources are designed to meet the general needs of health care to avoid an overcomplicated and redundant Resource set. Extensions and other customizations exist to allow resources to be adapted for specific use cases (the Profiling process). FHIR Resources also link to other Resources so that complex structures can be built.

Performance – Compared to previous standards, FHIR Resources are simpler in their construction, making them better-suited for exchange across a network and more easily understood and implementable by developers.

1. http://hl7.org/fhir/patient.html

Usability – FHIR Resources are designed to be understood by technical experts and non-technical people alike. Even if the details of the XML/JSON format are not understood, non-technical people can view these in a browser or text reader and understand their contents.

Fidelity – FHIR Resources have strict restrictions on intermixing of values with differing data types, like strings and numeric values. They can also be validated by their syntax in addition to defined sets of business rules.

Implementability – A core goal of FHIR was to create a standard that would lead to high adoption across disparate developer communities. Like the entire standard, FHIR Resources are designed to be easily understood and readily exchanged using industry standards, common programming languages, and established data exchange technologies.

Example FHIR Resource – Patient

The example below shows key parts of a Resource – the Resource used, a human readable summary of the data, an extension with its associated data, and the structured data itself.

```
<Patient xmlns="http://hl7.org/fhir">
    <id value="100"/>
    <meta>
        <lastUpdated value="2017-03-04T11:40"/>
    </meta>
    <text>
        <status value="generated"/>
        <div xmlns="http://www.w3.org/1999/xhtml">
            <p>John Doe, Male, DOB: 12/12/1979</p>
        </div>
    </text>
    <extension url="http://hl7.org/fhir/StructureDefinition/patient-religion">
         <valueString value="jedi"/>
    </extension>
    <identifier>
        <system value="http://hl7.org/fhir/sid/us-ssn"/>
        <value value="123-4465-8762"/>
    </identifier>
    <name>
        <family value="Doe"/>
        <given value="John"/>
    </name>
    <gender value="male"/>
</Patient>
```

→ **Resource Identity & Metadata**

→ **Human Readable Summary**

→ **Extension with URL to definition**

→ **Structured Data:**

- MRN
- Name
- Gender
- Birth Date

Provider

Interested in learning more about FHIR?

See more of our Fact Sheet series at:
https://www.healthit.gov/topic/standards-technology/standards/fhir-fact-sheets.

ONC's FHIR Fact Sheets are a collaborative effort with HL7 to help educate and demystify FHIR for federal employees. These fact sheets summarize the key technical concepts that make up the foundation of FHIR, how it is developed in an open and public process, and why FHIR adoption has become the focus of the health IT standards world. Full details and developer documentation can be found at HL7's FHIR website.