



# Project Lantern Data Sources and Linking Mechanisms

**Sponsor:** Office of the National Coordinator for Health IT (ONC)  
**Contract No.:** 75P00121F80185  
**Project No.:** 100905.10.100.1000.AA0

The views, opinions and/or findings contained in this report are those of The MITRE Corporation and should not be construed as an official government position, policy, or decision, unless designated by other documentation.

Approved for Public Release; Distribution Unlimited, Case 20-1740

©2023 The MITRE Corporation.  
All rights reserved.

**Bedford, MA**

## **Authors:**

**Matt Mayer**

**Emily Michaud**

**Brianna Mathiowetz**

**February 2023**



## **Executive Summary**

Lantern is an open source tool developed by the Office of the National Coordinator for Health Information Technology (ONC) and the MITRE Corporation that monitors and provides analytics about the availability and adoption of FHIR API service base URLs (endpoints) across healthcare organizations in the United States. It also gathers information about FHIR Capability Statements returned by these endpoints and provides visualizations to show FHIR adoption and patient data availability. Lantern sources most of its data from publicly available endpoint and organization lists, though some of the data is generated from the Lantern application itself. This document details the publicly available data sources and explains the processes used to produce data by the Lantern application.

# Table of Contents

<b>1</b>	<b>Endpoint Data</b> .....	<b>1</b>
1.1	NPPES Endpoint List.....	1
<b>2</b>	<b>Organization Data</b> .....	<b>2</b>
<b>3</b>	<b>Developer Data</b> .....	<b>2</b>
<b>4</b>	<b>Software Product Data</b> .....	<b>3</b>
<b>5</b>	<b>Data Validations</b> .....	<b>3</b>
<b>6</b>	<b>Linking Mechanisms</b> .....	<b>4</b>
6.1	Linking Endpoints to Organizations .....	4
6.1.1	Matching by NPI ID.....	4
6.1.2	Matching by Organization Name .....	5
6.1.2.1	Jaccard Index .....	5
6.1.2.2	Weighted Jaccard Index .....	5
6.2	Linking Endpoints to Developers .....	7
<b>7</b>	<b>Query Intervals</b> .....	<b>7</b>
<b>8</b>	<b>Endpoint Info History Pruning</b> .....	<b>8</b>
<b>Appendix A Abbreviations and Acronyms</b> .....		<b>A-1</b>

# List of Tables

Table 1. NPPES Endpoint List Fields.....	1
Table 2. NPPES NPI List Fields .....	2
Table 3. Fields Parsed from the CHPL Developers List.....	2
Table 4. Fields Parsed from the CHPL Products List.....	3
Table 5. Validation Result Table Format.....	3
Table 6. Base Validations .....	3
Table 7. FHIR R4 Validations .....	4
Table 8. Token Frequency to Token Value Mapping Ordered by Precedence.....	6
Table 9. Appendix Terms and Definitions.....	A-1

# 1 Endpoint Data

The Lantern project uses publicly available endpoint lists to generate an aggregated list of FHIR API endpoints. The majority of lists now come from the Certified Health IT Product List (CHPL) by querying [the /search API](#) for g(10) certified products. Outside of CHPL, there are a few publicly available lists that Lantern has included:

- CareEvolution ([https://fhir.docs.careevolution.com/overview/public\\_endpoints.html](https://fhir.docs.careevolution.com/overview/public_endpoints.html))
- 1upHealth (<https://1up.health/fhir-endpoint-directory>)
- Endpoints with type “FHIR” in the CMS National Plan and Provider Enumeration System (NPPES) National Provider Identifier (NPI) endpoint file, maintained by the Centers for Medicare and Medicaid Services (CMS).

The minimum required information that needs to be included in endpoint lists is the FHIR endpoint base URL. Most lists also include an organization name for each endpoint, and Lantern will also parse zip code information from endpoint lists, if available.

The FHIR Capability Statements retrieved from these endpoints have the capacity to list software names and versions. However, inclusion of this data is inconsistent and does not clearly map to CHPL. If the list is from CHPL, the one or more software products associated with it are mapped to the endpoints from the list. Furthermore, the FHIR Capability Statements do not have the capacity to link the FHIR endpoint to an organization, so Lantern relies on the organization names and other organization data reported by the FHIR endpoint list data sources to link a FHIR endpoint with an organization. Details regarding the methods used to link endpoints to organizations are included in Section 6.

## 1.1 NPPES Endpoint List

The NPPES NPI Endpoint File is available as a downloadable file in CSV format containing the following information:

**Table 1. NPPES Endpoint List Fields**

Field Name	Field Contents
NPI	NPI Number of the organization using this endpoint
Endpoint Type	Type of endpoint, we only parse rows with the value ‘FHIR’
Endpoint Type Description	Description of endpoint type, endpoint type fields populated with ‘FHIR’ contain the value ‘FHIR URL’ in this field
Endpoint	The base URL of the FHIR endpoint
Affiliation	Indicates if the endpoint is affiliated with another organization
Endpoint Description	Free text description of the endpoint
Affiliation Legal Business Name	Business name of affiliated organization
Use Code	Either Direct, HIE, OTHER, or <no_value>
Use Description	Either DIRECT, Health Information Exchange, Other, or <no_value>
Other Use Description	Free text description of other uses
Content Type	Content type of hosted content, current values are ‘CSV’ and ‘Other’
Content Description	Description of hosted content, current values are ‘CSV’ and ‘Other’
Other Content Description	Free text description of content
Affiliation Address Line One	Affiliated organization address line 1

Affiliation Address Line Two	Affiliated organization address line 2
Affiliation Address City	Affiliated organization address city
Affiliation Address State	Affiliated organization address state
Affiliation Address Country	Affiliated organization address country
Affiliation Address Postal Code	Affiliated organization address zip code

Further information about the contents of the NPPES NPI Endpoint file can be found here: <https://nppes.cms.hhs.gov/webhelp/nppeshelp/HEALTH%20INFORMATION%20EXCHANGE.html>

## 2 Organization Data

Organization data is parsed from the NPPES NPI data. The Lantern project uses the fields listed below to associate organizations with their FHIR endpoints. Lantern only stores NPI entries that represent health organizations, denoted in the NPI files where the “Entity Type Code” field equal is to “2”.

**Table 2. NPPES NPI List Fields**

Field Name	Field Contents
NPI	NPI Number of the organization using this endpoint
Provider_Organization_Name_Legal_Business_Name	Provider Organization Name (Legal Business Name)
Provider_Other_Organization_Name	Provider Other Organization Name
Provider_First_Line_Business_Practice_Location_Address	Provider First Line Business Location Address
Provider_Second_Line_Business_Practice_Location_Address	Provider Second Line Business Location Address
Provider_Business_Practice_Location_Address_City_Name	Provider Business Location Address City Name
Provider_Business_Practice_Location_Address_State_Name	Provider Business Location Address State Name
Provider_Business_Practice_Location_Address_Postal_Code	Provider Business Location Address Postal Code
Healthcare_Provider_Taxonomy_Code_1	Healthcare Provider Taxonomy Code

## 3 Developer Data

Developer data is parsed from the CHPL “/developers” API. Entries represent developers of certified health IT software products. The table below includes the list of fields that Lantern uses; additional fields can be found in [CHPL’s documentation](#).

**Table 3. Fields Parsed from the CHPL Developers List**

Field Name	Field Contents
id	Unique ID used within CHPL to identify developer
developerCode	Additional developer identification number
name	Name of the developer
website	URL of developer’s website
lastModifiedDate	Date which the developer’s entry was last modified
status	Indicates the active status of the developer
addressId	Unique ID of the address entry within CHPL
line1	Developer address line 1
line2	Developer address line 2
city	Developer address city
state	Developer address state

zipcode	Developer address zip code
country	Developer address country

## 4 Software Product Data

Software product data is parsed from the CHPL “/search/v2” API. Software products returned at this route represent certified health IT products that have been registered in the CHPL. The table below includes the list of fields that Lantern uses; additional fields can be found in [CHPL’s documentation](#).

**Table 4. Fields Parsed from the CHPL Products List**

Field Name	Field Contents
id	The CHPL ID of the developer who makes this software product
edition	The certification edition of this software product
product	Name of the software product
version	Version of the software product
chplProductNumber	Unique string used by CHPL to identify this software product
certificationStatus	Indicates if the software product is currently active
criteriaMet	List of CHPL criteria which this software product meets
certificationDate	Date that the software product was certified
practiceType	A practice type (either Ambulatory or Inpatient)
developer	The developer of the software product
apiDocumentation	Information about the documentation for the product

## 5 Data Validations

The Lantern system runs validations on the endpoints and stores the results in the **validations** database table. Lantern will run the set of base validations against all endpoints and will run FHIR version-specific validations depending on the version of FHIR advertised in the Capability Statement.

**Table 5. Validation Result Table Format**

Field Name	Field Contents
validation_result_id	Database id referenced by the endpoint in the fhir_endpoints_info table
valid	Indicates whether the actual value matched the expected value
actual	The actual value as reported by the endpoint
comment	Narrative explaining the validation
expected	Value(s) that will result in a passed validation
rule_name	Name of the validation
implementation_guide	Reference to an implementation guide (if any) relevant to the validation
reference	Link to relevant rule or standard that defines the expected value of the validation

**Table 6. Base Validations**

Validation Name	Validation Description
capStatExist	Asserts that a Capability Statement was returned by the endpoint



kindRule	Asserts that the Capability Statement’s <b>kind</b> field has the value “instance”
describeEndpointRule	Asserts that Capability Statement includes a value for either the <b>description</b> , <b>software</b> , or <b>implementation</b> fields
documentValidRule	Asserts that if elements exist in the <b>document</b> field of the Capability Statement, that the documents listed are unique when keyed by the <b>document.profile</b> and <b>document.mode</b> fields
endpointFunctionRule	Asserts that the Capability Statement includes at least one <b>rest</b> , <b>messaging</b> , or <b>document</b> element
messagingEndptRule	Asserts that if the Capability Statement’s <b>kind</b> field has the value “instance”, then the <b>messaging</b> field should not be available
uniqueResourcesRule	Asserts that the list of resources advertised in the Capability Statement’s <b>rest</b> field does not contain duplicate resources

**Table 7. FHIR R4 Validations**

Validation Name	Validation Description
patResourceExists	Asserts that the Capability Statement advertises support of the <b>Patient</b> resource
tlsVersion	Asserts that TLS version 1.2 or higher is used during transmission
otherResourceExists	Asserts that the Capability Statement advertises support for a resource in addition to the <b>Patient</b> resource
smartResponse	Asserts that the SMART Response resource is returned when querying the <b>/.well-known/smart-configuration</b> endpoint
instanceRule	Asserts that if the CapabilityStatement’s <b>kind</b> field has the value “instance” then the <b>instance</b> field should be available
versionsResponseRule	Asserts that the default FHIR version as specified by the <b>\$versions</b> operation should be returned from the server when no version is specified
searchParamsRule	Asserts that the names of search parameters within a resource are unique to said resource

## 6 Linking Mechanisms

### 6.1 Linking Endpoints to Organizations

Endpoints can be linked to organizations in two ways, either by the NPI ID (preferred), or by the organization name. If the endpoint list also includes a zip code for the organization, that is also factored in. Lantern links endpoints and organizations by associating endpoint entries with organization entries and storing the association in the database along with a confidence metric representing the confidence that the association is correct. Details about how the confidence metric is calculated will be discussed further in this section.

#### 6.1.1 Matching by NPI ID

As of the writing of this document, the NPPES Endpoint file and the MEDHOST list (included in CHPL) are the only endpoint list sources that provides both an endpoint FHIR base URL along with a mechanism (an NPI ID) to link the endpoint to a unique organization. Links made between organizations and endpoints using an NPI ID are given a match confidence value of 1, which is higher than any possible confidence value for matches made using the organization name.

## 6.1.2 Matching by Organization Name and Zip Code

In instances where a unique identifier to match an organization to an endpoint is not provided, Lantern uses the organization name and the zip code from the list (if provided), and the primary and secondary organization names provided by the NPPES NPI data set to associate an endpoint with an organization from the NPPES NPI data set. We use a modified version of the Jaccard Index equation to match organizations based on their names and assign a match confidence score. The modified Jaccard Index algorithm assigns weights to each token in the organization names, the weights are inversely proportional to the number of times the token appears across all organization names. If the match confidence is at or above the match confidence threshold of .85, it will be considered as a match. If the confidence interval is above .85 and a zip code was also included for the endpoint and equals the zip code for the NPPES NPI organization, then the confidence value for that endpoint is increased to .99.

### 6.1.2.1 Jaccard Index

The un-weighted Jaccard index is calculated by dividing the number of words that the two names being compared have in common by the total number of words across both names. In the un-weighted Jaccard Index calculation, every token is given the same weight of 1. Mathematically, this is defined as the intersection of two sets, divided by the size of the union of the sets. In this case, each set is the list of words (tokens) that make up an organizations name. For example, if one was comparing the names “Foo Bar” and “Foo Bar Baz”, the sets being compared would be [“Foo”, “Bar”] and [“Foo”, “Bar”, “Baz”] and the resulting Jaccard Index would be  $2/3=.666$  since the two names have tokens “Foo” and “Bar” in common and there are a total of three unique tokens between both names, “Foo”, “Bar” and “Baz”. The higher the Jaccard Index value, the more similar the two names being compared are.

### 6.1.2.2 Weighted Jaccard Index

The Lantern application uses a modified version of the Jaccard Index to measure the similarity of organization names. In this modified version, tokens are weighted proportional to their uniqueness within the set of tokens that appear in all organization names. The consequence of assigning a lower weight to common words is that a higher degree of certainty is assumed when a unique word appears in both sets of organization name tokens. Token values are calculated as follows:

1. Build a list of all tokens that appear in all organization names existing in the NPI data source in addition to every organization name reported by the FHIR endpoints lists. Names will be normalized by converting them to uppercase and removing of any special characters.
2. Build a dictionary of the words that provide the least matching value. This dictionary is composed of the top N most frequent words, combined with words found to have little matching value through trial and error. First, sort the token count list by frequency and determine the top number of words which appear so frequently and do not provide any uniqueness to help aid in matching that their weight in the Jaccard Index calculation should be especially low. Next, determine the top words that hinder match outcomes rather than improve match outcomes by repeatedly running the algorithm and evaluating the difference in the output produced by the matching script to check the validity of match scores of endpoints that were matched to organizations, as well as the validity of endpoints that were *not* matched to organizations. The endpoints that were successfully

matched but had lower scores than expected were compared with their NPPES NPI organization match to see if there are any unnecessary tokens causing a lower matching score, such as abbreviations. Those endpoints that were unsuccessful in matching were compared with the NPPES NPI data set to see if there were any organizations they should be matching with, and if there was, the endpoint organization name was compared with the NPPES NPI organization name to see if there were any unnecessary tokens causing the organizations not to match. These words, along with the most frequent words that do not aid in matching, are added to a dictionary of least valuable words, these words represent the tokens that have the lowest weight. Words to be added to this dictionary are determined by iteratively running the algorithm and increasing the number of included words until the number of quality matches decreases. Determining the number of quality matches is done by examining the difference in output of the matching script between runs and is not yet an optimized process. As of the writing of this paper, we have found the number of these words to be equal to 31, with a fluff dictionary containing the following words:

"EMS", "DR", "PA", "MD", "LLC", "LTD", "PC", "DPM", "LLP", "AND", "OF", "IN", "THE", "MCC", "MMC", "TO", "PLC", "PLLC", "SYSTEM", "SERVICES", "DPMPC", "MDSC", "CORP", "HSHS", "ST", "CARE", "INC", "CLINIC", "GROUP", "CENTERS", "CENTER"

3. Calculate the max, the mean and standard deviation of the token-frequency list, use the chart below, which is ordered by precedence, to assign the value that should be used in the Jaccard Index calculation for each token.

**Table 8. Token Frequency to Token Value Mapping Ordered by Precedence**

Token Frequency	Resulting Token Value
Tokens Included in Dictionary of Least Valuable Tokens	$1 - (\text{Token Frequency} / \text{Max Frequency}) * .2$
$0 < \text{Frequency} < \text{Mean}$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * 2.5$
$\text{Mean} < \text{Frequency} < (\text{Mean} + \text{Standard Deviation} / 3)$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * 1.6$
$(\text{Mean} + \text{Standard Deviation} / 3) < \text{Frequency} < \text{Mean} + \text{Standard Deviation}$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * 1.3$
$\text{Mean} + \text{Standard Deviation} < \text{Frequency} < \text{Mean} + \text{Standard Deviation} * 3$	$1 - (\text{Token Frequency} / \text{Max Frequency})$
$\text{Mean} + \text{Standard Deviation} * 3 < \text{Frequency} < \text{Mean} + \text{Standard Deviation} * 6$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * 0.8$
$\text{Mean} + \text{Standard Deviation} * 6 < \text{Frequency} < \text{Mean} + \text{Standard Deviation} * 9$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * 0.6$
$\text{Mean} + \text{Standard Deviation} * 9 < \text{Frequency}$	$1 - (\text{Token Frequency} / \text{Max Frequency}) * .4$

4. After a value is assigned to a token using the chart above, check to see if the token appears in either the list of organization names gathered from the endpoints or the list of organization names gathered from the NPPES NPI data set, but not the both. A token that may appear frequently in one source, but never in the other, source may have a high frequency weight but will never be valuable in matches. Tokens that meet this criterion will have their weight multiplied by .3.
5. Compute the Jaccard Index between organization names provided by FHIR endpoint lists and organization names provided by the NPPES NPI database substituting the computed token values for the token counts that were used in the unweighted Jaccard Index detailed in Section 4.2.1.1. The computed match confidence is equal to the Jaccard Index

multiplied by .99, since we can only have a match confidence of 1 when organization matching is performed via a unique identifier provided by an endpoint list.

6. Any calculated match confidences at or above the match confidence threshold of .85 will be considered as a match. If the confidence interval is above .85 and a zip code was also included for the endpoint and equals the zip code for the NPPEs NPI organization, then the confidence value for that endpoint is increased to .99.

An additional modification to the Jaccard Index formula in the Jaccard Index calculation used by Lantern is that Lantern considers each individual token in an organization name to be unique, meaning that repeated words will be included when an organization name is converted into a set of tokens. For example, the organization name “Foo Bar Foo” would result in the set [“Foo”, “Bar”, “Foo”].

## 6.2 Linking Endpoints to Developers

Most of the endpoints in Lantern are from endpoint lists in CHPL, so the developer is already associated with an endpoint list. Mapping in this case is simple since it is pulled from the CHPL entry and saved to any endpoint in the developer’s list.

However, there are still lists in Lantern that are not from CHPL where the developer is not included. In this case, Lantern links FHIR endpoints to developers using developer names reported both in the **publisher** field of the Capability Statement and the CHPL **developers** list.

When a capability statement is received, the following matching steps are performed:

1. Normalize both the reported publisher from the Capability Statement and all of the CHPL developer names by converting all names to lowercase and removing any of the following words:

"inc.", "inc", "llc", "corp.", "corp", "corporation", "lmt", "lmt.", "limited", "corporation."

Finish the normalization process by removing any trailing punctuation.

2. Iterate over the entire list of normalized developer names from the CHPL developers list. If the normalized developer name is a substring of the publisher name or vice versa, then the developer is considered to be a match.

This approach to name mapping is simpler than our methodology that links endpoints to organizations by name because there are far fewer developers than organizations to match. In addition, the **publisher** field of Capability Statements is less variable than the organization names in endpoint lists published by developers.

## 7 Query Intervals

Lantern queries its list of known FHIR endpoints once every 23 hours. Setting the query interval to once every 23 hours means that over time Lantern will have queried each endpoint at every different hour of the day. During each query Lantern records data from each endpoints’ Capability Statement in addition to the HTTP response code and response time associated with the request made to the endpoint.

## 8 Endpoint Info History Pruning

After every query interval, once the **Capability Querier** service has finished querying all endpoints and updating both the **fhir\_endpoint\_info** database table and subsequently the **fhir\_endpoint\_info\_history** database table, the history pruning algorithm is run. The pruning algorithm will iterate over all of the **fhir\_endpoint\_info\_history** entries for each distinct FHIR endpoint URL that have **entered\_at** dates that are older than the time determined by subtracting the environment variable `LANTERN_PRUNING_THRESHOLD` from the current time, and also have **entered\_at** dates that are newer than the current time minus the pruning threshold doubled. Having a lower limit of the pruning threshold time doubled ensures that the algorithm does not repeat pruning checks on the same entries after every query interval, but that it also does not miss any entries that have not yet been pruned. The pruning threshold, which set to one month by default, ensures that there is always data newer than the `LANTERN_PRUNING_THRESHOLD` that is not pruned, since an entry has to be older than the threshold in order to be considered for pruning.

The pruning algorithm will remove any consecutive duplicate entries in the **fhir\_endpoint\_info\_history** table. A **fhir\_endpoint\_info\_history** entry is considered a duplicate if there is an older consecutive entry that has the same stored information for the endpoint's TLS version, MIME types, and SMART response, and if the newer entry's stored Capability Statement only differs by fields included in a list of ignored fields, such as the **date** field. If a **fhir\_endpoint\_info\_history** entry is found to be a duplicate of an older consecutive entry, it is deleted from the table, and this continues until only the oldest of the consecutive duplicated entries remains. This pruning strategy is advantageous in that there will always be a duration of at least `LANTERN_PRUNING_THRESHOLD` worth of queries in the history table for each endpoint, therefore Lantern can inspect `LANTERN_PRUNING_THRESHOLD` worth of data to see how every endpoint responded within the threshold while still saving storage space by removing duplicate data or data which only differs in the values reported for fields in the ignored fields set. Keeping all entries containing any unique data allows Lantern to keep track of how each endpoint has changed over long periods of time.

## Appendix A Abbreviations and Acronyms

The list of abbreviations/acronyms includes all abbreviations, initialisms, and acronyms listed in the document.

**Table 9. Appendix Terms and Definitions**

<b>Term</b>	<b>Definition</b>
<b>API</b>	Application Programming Interface
<b>CHPL</b>	Certified Health IT Product List
<b>CMS</b>	Centers for Medicare and Medicaid Services
<b>CSV</b>	Comma-Separated Values
<b>FHIR</b>	Fast Healthcare Interoperability Resources
<b>HTTP</b>	HyperText Transfer Protocol
<b>NPI</b>	National Provider Identifier
<b>NPPES</b>	National Plan and Provider Enumeration System
<b>ONC</b>	Office of the National Coordinator for Health IT
<b>URL</b>	Uniform Resource Locator