



DIBBs Query Connector Query Connector FAQ

Created Date: October 28, 2025

Last Updated: June 4, 2026



Prepared by:

U.S. Department of Health and Human Services



Version History

Version	Author	Reviewed By	Revision Date	Revision Notes
1.0	Matthew Goldberg	Maya Mascarenhas	10/28/2025	
1.1	Matthew Goldberg	Maya Mascarenhas	03/05/2026	
1.2	Matthew Goldberg	Maya Mascarenhas	06/04/2026	



Table of Contents

1	Query Connector overview	4
2	Authentication in Query Connector	6
3	Product model and data ownership	6
4	Capabilities and limitations.....	7
5	Technical overview	7
6	Adopting Query Connector.....	8
7	Glossary.....	9



To contact the Query Connector team, email dibbs@cdc.gov

1 Query Connector overview

1.1 What is Query Connector?

Query Connector (QC) is a free, open-source application offered by CDC that lets public health agencies (PHAs) securely query Fast Healthcare Interoperability Resources® (FHIR®)¹ APIs at healthcare organizations (HCOs) and health information exchanges (HIEs) to retrieve the specific data they need.

Helpful links:

- Demo: <https://dibbs.tools/>
- Documentation: <https://connector.dibbs.tools/docs>
- GitHub: <https://github.com/CDCgov/dibbs-query-connector>

1.2 What challenges does Query Connector solve?

By enabling data exchange between PHAs and healthcare providers without requiring direct connection and onboarding, Query Connector:

1. **Reduces manual burden.** Improves workflows for public health case investigators by replacing the process of collecting patient data by individual searches via portal, phone, email, or fax.
2. **Increases speed of data queries.** Gives public health staff immediate access to patient data relevant to their specific use case.
3. **Protects patient privacy.** Only requests the minimal data necessary and implements role-based permissions to manage who can request and access the data.

Collectively, the reduced manual burden and increased query speed helps public health staff more effectively do case investigations and manage public health responses.

1.3 How is Query Connector deployed?

PHAs deploy and run their own instance of Query Connector on infrastructure they control (on-prem or in their cloud). As a result:

- ✗ PHA owns all data associated with their instance.
- ✗ CDC does **not** have access to PHA data.

¹ HL7®, and FHIR® are the registered trademarks of Health Level Seven International and their use of these trademarks does not constitute an endorsement by HL7.

- ✎ Keeping data in the U.S. is straightforward. Jurisdictions can select a U.S. cloud region or keep the system on-prem.

1.4 How does Query Connector connect to data sources?

To connect QC to a FHIR server, the PHA must first register their instance as a client to obtain authorization. This is typically done either by registering directly with a specific healthcare organization's EHR (e.g., Epic) or through an HIE (such as eHealth Exchange) to enable broader access across multiple providers. Once the organization or exchange authorizes access, the PHA completes the configuration by adding the FHIR connection in QC. Authorization between QC and FHIR endpoints is secured using industry standards such as OAuth, mTLS, and SMART on FHIR.

Connections may be:

- ✎ **Direct** to an HIE, hospital, or any responding organization's FHIR endpoint
- ✎ **Via eHealth Exchange** to connect to FHIR endpoints at HCOs who are current members of their network

1.4.1 Why is there an opportunity for FHIR clients in public health departments?

The opportunity exists because of new EHR certification requirements. Currently the Centers for Medicare and Medicaid (CMS) requires providers that participate in Medicare or Medicaid (which is almost all providers) use a certified EHR. The Office of National Coordinator for Health Information Technology (ONC) lead federal efforts around health IT certification and one key aspect of that is that EHRs must support FHIR-based APIs.

FHIR is a modern HL7 healthcare data standard. It's similar to other HL7 standards like:

- ✎ HL7 V2 used by syndromic surveillance and Electronic Lab Reporting (ELR) systems
- ✎ CDA (Clinical Document Architecture) used by eICR (electronic Initial Case Reports)

1.4.2 What does this mean for public health departments?

Because FHIR support is now mandated by regulation, there are FHIR APIs available from HCOs. As an authorized user, a public health department could connect to these APIs (using a tool like Query Connector) and make very granular queries to access the data they need.

1.5 How long does deployment take? What does the process look like?

After PHA approval, a typical deployment can be completed in about two weeks.

- ✎ **Week 1:** Three 1-hour pairing sessions with the PHA's IT staff and QC engineers.
- ✎ **Week 2:** Remaining tasks are refined based on Week 1, with additional pairing as needed to complete deployment.

PHAs can connect directly to HCOs or HIEs, or connect via a Qualified Health Information Network (QHIN) intermediary such as eHealth Exchange. For direct connections, the PHA should identify an HCO/HIE partner and follow that partner's connection procedures.

The CDC DIBBs team can support PHAs during deployment, subject to internal approvals. For questions and support requests, contact us at dibbs@cdc.gov.

1.6 How much does it cost to run to Query Connector?

Query Connector is free, open-source software. PHA's only pay for the hosting infrastructure—there are no licensing fees, per-user charges, or transaction fees. Because QC isn't doing much heavy computation and most time is spent waiting on network responses, hosting is relatively inexpensive.

Hosting costs: For most PHAs, QC costs ~\$5,000 per month to host. This covers:

- ☒ Cloud infrastructure (compute, storage, database)
- ☒ All application components
- ☒ Standard redundancy and backups

QC is designed to work with what most public health agencies already have in place. A modest server (CPU/memory) is typically sufficient.

By evaluating your current environment, you can uncover cost savings from leveraging existing assets.

2 Authentication in Query Connector

2.1 How does authentication work?

QC integrates with a PHA's Identity Provider (IdP) via OpenID Connect (OIDC). Staff sign in with existing agency credentials (e.g., Microsoft Entra ID/Azure AD, Keycloak).

No separate QC-only accounts are required.

2.2 How is account administration handled?

All user management occurs in the agency's IdP—the same place the agency already manages accounts. QC recognizes three roles assigned via the IdP:

- ☒ **Super Admin:** Manage user permissions and groups; view audit logs; configure FHIR servers; create, assign, and run queries.
- ☒ **Admin:** Create, assign, and run queries.
- ☒ **Standard:** Run queries only.

Super Admins and Admins can create user groups (e.g., a program's case investigators) and assign queries to individuals or groups, controlling who can query what.

3 Product model and data ownership

3.1 Is QC a SaaS product?

No. QC is not offered as SaaS. PHAs deploy and operate their own instance on infrastructure they control (on-prem or in their cloud).

3.2 What data does Query Connector store?

QC maintains an audit log of user actions, including the parameters of initiated queries (e.g., Personally Identifiable Information (PII)). The audit log is visible only to Super Admins. Query responses (e.g., Protected Health Information (PHI)) are not retained by QC at this time.

Because the PHA hosts QC, any retention or destruction of logs or configuration data follows the jurisdiction's own policies and tooling.

3.3 What happens to data if an agency stops using Query Connector?

No contract is required with CDC to use QC. Since PHAs operate their own instances, they control the environment and data. QC does not retain PHI from query responses; PHAs can archive or delete QC audit logs and configuration according to internal procedures.

4 Capabilities and limitations

4.1 Does Query Connector reconcile or merge records into a single patient view?

QC focuses on running queries and does not retain PHI from responses. Reconciliation/merging occurs in the PHA's existing systems or interface engine after data retrieval.

4.2 Can Query Connector schedule automated queries on its own?

Not currently. PHAs that need scheduled or event-driven queries can trigger the QC API from their own interface engine and manage schedules there.

4.3 Can providers use Query Connector to query public health data?

Potentially. If a PHA exposes FHIR® resources for provider access to public health information such as antimicrobial resistance data, the QC architecture supports this usage.

4.4 Is adopting TEFCA required?

No. Query Connector supports TEFCA adoption but is not dependent on it. As a FHIR client, QC can connect to any authorized FHIR server. Pilots are currently testing both direct connections (e.g., HIE/HCO FHIR endpoints) and connections via eHealth Exchange that rely on a wide range of data sharing agreements.

5 Technical overview

Our [Technical Architecture](#) and [Deployment Guide](#) can help your jurisdiction understand how to implement QC. This guide includes:

- A high-level view of the QC's architecture and how it works
- The technical guidance for deploying QC

6 Adopting Query Connector

6.1 What are the steps to implement Query Connector?

1. **Review options & confirm fit**
Use our self-service documentation to understand connection options, prerequisites, and the recommended deployment approach for your jurisdiction.
2. **Prepare your environment & stakeholders**
Work with your cloud and security teams to complete required approvals and readiness items (accounts, networking, identity/access, and security controls).
3. **Deploy QC**
Follow the deployment guide for your environment to install and configure QC.
4. **Configure connections**
Set up your FHIR connection(s) and any required integration settings following your HCO/HIE or QHIN procedures.
5. **Test & validate**
Complete acceptance testing to confirm QC is ready for end users.
6. **Go live & operate**
Use the maintenance guide for upgrading, troubleshooting, and getting ongoing support.

6.2 What type of support is available?

If your jurisdiction is getting active support, email your point of contact on the DIBBs team. If not, email us directly at dibbs@cdc.gov for any other support needs.

7 Glossary

API: Application programming interface (API) is a set of functions and procedures that allow the creation of applications that access features/data of another operating system, application, or other service.

EHR: electronic health records (EHR) are the electronic version of a patient's medical history that providers maintain over time and may include all of the key administrative clinical data relevant to that patient's care.



FHIR®: Fast Healthcare Interoperability Resources® (FHIR®) is a data standard that allows different healthcare systems to exchange information efficiently. It uses a set of about 150 “resources,” such as patients, observations, and medications, that are based on modern web standards like RESTful APIs.

LOINC: Logical observation identifiers names and codes (LOINC) is a standardized coding system used for identifying health measurements, observations, and services. LOINC codes facilitate the exchange and aggregation of clinical laboratory data.

RCKMS: Reportable conditions knowledge management system (RCKMS) is an authoritative, real-time, digital platform to enhance disease surveillance by providing comprehensive information on public health reporting requirements to public health reporters such as clinicians and laboratories.

SaaS: Software as a Service (SaaS) is a software delivery model that allows users to access software via a subscription, while the provider manages all the underlying infrastructure, security, and updates.

TEFCA: The Trusted Exchange Framework and Common Agreement (TEFCA) aims to establish a universal policy and technical floor for nationwide health data interoperability by simplifying connectivity for public health authorities and agencies to exchange information.