Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Interest in pursuing the DHIN Data Warehouse RFP exceeded expectations and not all questions have received a response as of 09/21/2026. We will update this document daily until all questions are addressed:

Q: Does the DHIN have an existing Data Warehouse to be migrating from? What is the overall size and what kind of structure model and technology stack?

A: *See current platform overview.*

Q: If DHIN has an existing Data Warehouse, what are the pain points a new system can alleviate?

A: Our data ecosystem spans both third party and internal platforms. We are looking to unify our data platform as much as possible as well as have greater administrative and operational control over the data.

Q: For the requirement ""Code set mapping to industry standards"": Does the DHIN already possess codeset mapping content or shall that be bundled in the proposal? Ex: If the DHIN has these mappings in an existing database, it could be imported to the Data Warehouse to lower cost.

A: DHIN does currently possess codeset mapping content primarily as JSON and a standard vocabulary repository in RDS. Functionality to flag issues or non-compliance would be a benefit.

Q: How many data sources are to be converted / migrated to the new platform? Will data be migrated from the existing Data Warehouse as is or must all data be loaded from source message archives?

A: There are approximately 120 data sources sending clinical and claims data. Clinical data will be loaded from source message archives. Claims data will be a mix of migration and source file loading.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: What is the volume of data to be migrated? How many years?

examples:

FHIR Observation resources = 5.5 billion / 800 GB

FHIR Encounter resources = 184 million / 100 GB

Flat file claims = 200 million / 400 GB

etc...

A: *See current platform overview.*

Q: What are the expected volumes of inbound data, by year? Please provide a count and size in GB or TB.

A: *See current platform overview.*

Q: Please provide the following sizing metrics: Unique MRNs today, expected annual growth

A: 5.5M unique identities with 33M unique source and mrn/patient identifier combinations mapped to them. Annual growth has averaged approximately 100K records. Orphaned data is approximately 1M unique identities from single sources.

Q: Please provide a detailed breakdown of transaction types to be ingested in an ongoing basis (include Count and Storage Size):

- HL7 ADT
- HL7 ORU Results: Lab, Radiology and textual results
- HL7 Medications
- HL7 Immunizations
- HL7 Other types: (list)
- CCD / CCDAs
- CDA - Unstructured documents
- FHIR Resources
- X12 Claims
- Others (Flat files, CSV, etc)

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

ex:

- HL7 ADT = 3,000,000 per month, 2.4 GB
- CCD = 500,000 per month, 625 GB

- A: *See current platform overview.*

Please provide expected data processing and utilization (monthly count):

- Standard FHIR API Queries
- Bulk FHIR Queries (expected executions and patient population size)
- FHIR to CCD Generation
- XCA Unstructured Document Queries
- Reports - please list the anticipated scopes of reports and their frequencies
- Ex:
    - Monthly HEDIS report, 400,000 Patients
    - CKD Quality Measure, all Patients
    - Readmission Report for Beebe, 125,000 Patients

A: *See current platform overview.*

Q: What is the annual budget defined for this project?"

A: Responders are asked to provide their typical annual, all-in costs based on information provided in the current platform overview.

Q: Existing AWS Environment (Sec. 2.1.3): The RFP states a preference for hosting in DHIN's own AWS platform. Could you provide details on the existing AWS environment? Specifically, are there established standards, security policies (e.g., IAM role requirements, encryption standards), and network configurations (VPCs, Direct Connect) that the proposed solution must adhere to?

A: DHIN is HITRUST certified, data is encrypted at rest and in transit, adheres to strict healthcare privacy and security controls, maintains extensive audit logs, and continuously monitors access and traffic.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Current Data Volume (Sec. 2.1): To accurately size the platform, could you provide current and projected data volumes for the key data types mentioned (HL7, FHIR, Claims, Clinical Notes, etc.)? What is the total size of the data that will need to be migrated?

A: *See current platform overview.*

Q: EMPI Integration (Sec. 2.1.5): The RFP requires near real-time integration with the existing Verato EMPI. What specific APIs, interface documentation, and integration patterns are available for this platform? What are the expected transaction volumes and performance requirements for the synchronization?

A: Our EMPI vendor has extensive API documentation for additions, updates, queries, maintenance, and notifications averaging 33 million API calls per month.

Q: Provider Directory (Sec. 2.1.5): The RFP mentions the data warehouse could be ""The"" source of truth for a provider directory or integrate with a 3rd party tool. Does DHIN have a preference between building this capability versus integrating with an existing or new product?

A: DHIN's preference is to be able to leverage existing sources (NPPES, DE State Licensing, and other CMS resources) into our own system.

Q: OMOP Data Model (Sec. 2.1.3): Regarding the requirement to support an OMOP common data model, is there an existing implementation or a specific version you are targeting? What are the primary research use cases envisioned for the OMOP data?

A: DHIN does not currently have an OMOP CDM instance. We would target version 5.5. We would like to be able to utilize ATLAS and other OHDSI tools for the identification of cohorts and analysis of data.

Q: Supplemental Datasets (Sec. 2.1.1): Could you provide more detail on the ""Supplemental datasets"" such as the Social Vulnerability Index and risk scores? What are the sources, formats, and frequency of these data feeds?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Supplemental datasets include geocoding, SDOH indices, and risk scores from different potential sources such as Johns Hopkins or Solventum.

Q: Data Quality (Sec. 2.1.2): Are there currently defined data quality metrics or known data quality issues with the existing source systems that the new platform will need to address?

A: DHIN is in the early stages of implementing the PIQI framework for data quality analysis. Data validation and integrity notifications during ingestion would be a benefit.

Q: Historical Data (Sec. 7): The RFP notes that DHIN has been collecting clinical data since 2007 and claims data since 2013. Is the expectation that all historical data will be migrated to the new platform, or only a subset?

A: All historical data will be migrated.

Q: ""Live"" USCDI Versions (Sec. 2.2): The RFP mentions compliance with ""two versions of the currently 'live' or published version"" of USCDI. Could you please clarify what DHIN considers the current ""live"" versions for the purpose of this RFP?

A: DHIN would like to be within two versions of the Standards Version Advancement Process (SVAP) approved releases. The Approved standards for 2026 is USCDI v6.

Q: RHTP Funding (Page 1 & Sec. 2): The RHTP funding requires project completion and invoicing by September 30, 2027. Are there any other specific constraints, reporting requirements, or ""flow down provisions"" associated with this federal funding that bidders should be aware of?

A: As of 9/17/26 we do not have the final contract with the state and are not aware of any “flow down provisions.”

Q: DHIN Resources (Sec. 4.3): The work plan requires identifying DHIN resources. Could you provide an overview of the key DHIN personnel who will be dedicated to this project and their expected level of involvement?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: The project will be overseen by an executive sponsor, project manager, and supported by a full time data engineer with significant time dedicated from a technical architect, data quality analyst, integration engineers, and an infrastructure engineers.

Q: Pricing Methodology (Sec. 4.4): While a flat fee is preferred, the RFP asks for calling out different configurations (e.g., AWS vs. Snowflake). For evaluation purposes, is there a baseline configuration or set of assumptions we should use for the primary proposal?

A: There are no baseline configurations or set of assumptions, the intent is to address managed vs self hosted, or different back end options for a FHIR CDR.

Q: MSA & SOW (Sec. 4.1): The RFP requests a proposed MSA and SOW. Given the tight timeline, would DHIN be open to negotiating the MSA with the selected finalist in parallel with the finalization of the SOW to expedite the process?

A: Yes

Q: M&O Scope (Sec. 2.2): While Section 2.2 provides details about M&O activities and functions, there is no period of evaluation mentioned. How should vendors structure ongoing M&O activities and contracting terms ? Is DHIN seeking a fixed monthly/annual recurring fee, a tiered subscription model based on data volume/users, or a core base retainer with time-and-materials (T&M) for ad-hoc requests? Or is DHIN expecting a completely new procurement / SOW for the M&O work ?

A: DHIN is flexible on the model, but would expect a fixed monthly/annual fee for routine patches, updates, and support for the platform and a SOW for work that may be needed outside of that scope.

Q: Infrastructure & Hosting Cost Allocation (Sec. 2.1.3 & Sec. 4.4): The RFP states a preference for hosting within DHIN's AWS environment while also allowing alternatives such as Snowflake. For ongoing operations, will all cloud infrastructure, storage, egress, and third-party software license costs be paid directly by DHIN under its own cloud agreements, or should vendors bundle infrastructure pass-through costs into the ongoing M&O fee?"

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: If hosted within DHIN's AWS environment, DHIN will cover cloud infrastructure, storage, egress, and third-party software license costs not included in the submitted proposal. If hosted externally those costs should be included in the proposal.

Q: Is there a list of State approved contractors that would be acceptable to DHIN to act as subcontractors per Section 8.4. Alternatively, how would a bidder obtain consent to use a selected subcontractor prior to RFP submission to avoid disqualification?

A: We do not have a list of approved subcontractors.

Q: Please confirm the proposed data warehouse will be downstream of the State's existing HIE solution and interface directly with the HIE and not integrate directly with health providers. To refine pricing, would you please clarify the distinct systems that would be require integration directly with the new data warehouse

A: DHIN is the Delaware's existing HIE solution. The clinical data repository with FHIR data model will be the end point of health care providers sending data to DHIN. The integration engine is Mirth and will be managed by DHIN, and we currently have approximately 120 data senders using multiple EMR systems.

Q: How many years of data is the State anticipating to convert into the new data warehouse. Where does the data for conversion currently reside in the State's HIE architecture (e.g., PostgreSQL, Redshift, the FHIR based Clinical Data Repository)?

A: *See current platform overview.*

Q: Would the State please clarify who the current third party vendor(s) are for the Clinical Data Repository and APCD, and whether those contracts will remain in place, be integrated with, or be terminated as part of this initiative?

A: MedicaSoft is DHIN's current vendor for the CDR and APCD. All three options are a possibility depending on the proposals received.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Would the State please clarify the intended data flow between the existing HIE and the new Data Warehouse, for example, will the HIE continue to perform real time ingestion and routing (via Mirth Connect) and simply forward or replicate data into the new Warehouse, or will the Warehouse itself assume ingestion responsibilities?

A: DHIN is the current existing HIE for Delaware. Mirth will continue to be the ingestion engine with the expectation that the selected solution will transform HL7 messages and CDA's into the FHIR data model.

Q: Would the State please provide a current system and data flow architecture diagram of the HIE to help bidders scope the integration boundary accurately?

A: *See current platform overview. Existing diagrams expose sensitive information*.

Q: Would the State please clarify the volume, size, format, and date range of historical clinical data (since 2007) and claims data (since 2013) that must be migrated or made accessible through the new platform?

A: *See current platform overview.*

Q: Would the State please clarify the current daily and peak transaction volumes by message type (HL7 v2.x, CCDA, FHIR API, X12 claims, ELR, syndromic surveillance) rather than the aggregate ""up to a million transactions daily"" figure, to support sizing?

A: *See current platform overview.*

Q: Would the State please clarify which HL7 v2.x message types beyond ADT, ORU, ORM, MDM, and VXU (such as SIU or DFT) should be anticipated, and whether a full list of currently supported message types and trading partners can be shared?

A: *See current platform overview.*

Q: Would the State please clarify what ""the entire HL7 FHIR data spectrum"" means in practical terms, that is, which specific FHIR resource types and implementation guides (such as US Core or CARIN BB for claims) must be supported at go live versus over time?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Our current environment is FHIR R4 with HL7 v2.51 and USCDI v2 and v3.1 transformations into FHIR. Live would need to be able to support these versions with a roadmap for continuously adapting to new versions and standards.

Q: Would the State please clarify the specific users/systems that would retrieve information directly from the data warehouse?

A: Direct connections to the clinical data repository are through an API and include the Community Health Record, Personal Health Record, Sensitive Data Viewer, Advanced Directives, eHealthExchange. Tableau will connect for reporting and visualization. Postman and DB IDEs for internal use.

Q: Would the State please clarify the expected scope and use cases for bulk FHIR retrieval (for example, research extracts or payer bulk data requests) to help bidders size infrastructure appropriately?

A: We do not have a current active use case for bulk FHIR but are looking at if for SMART on FHIR applications, FHIR to OMOP, and research.

Q: Would the State please clarify the specific integration requirements and interface specifications currently in place with Verato (EMPI), including transaction volume and expected SLAs for near real time synchronization?

A: Our integration engine is the primary tool for maintaining our EMPI. The data warehouse/clinical data repository should be able to remain in sync within 5 minutes with our EMPI using the notification API. There are currently 33M API calls to our EMPI monthly.

Q: Would the State please clarify whether DHIN currently has, or intends to separately procure, a third party Provider Directory product, or whether bidders should propose the Data Warehouse platform as the source of truth for this function?

A: The provider directory will be sourced from the data warehouse. Integrations and maintenance will be DHIN's responsibility. Any proposed functionality by respondents will be reviewed and weighed accordingly.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Would the State please clarify the specific AWS account structure, region(s), and any existing infrastructure as code or DevOps tooling bidders would be expected to work within if hosting in DHIN's AWS environment"

Q: Section 5.1 states that DHIN reserves the right to accept a portion or portions of a proposal, while Section 6 states that any proposal failing to respond to all requirements may be eliminated from consideration. Will a proposal that addresses every requirement in Section 2.1, designating certain capabilities as out of scope for the respondent and identifying a recommended path for DHIN to address them, be considered responsive? Relatedly, does DHIN anticipate a single award, or may it make multiple awards covering different portions of the scope?

A: Your description for addressing requirements in section 2.1 would be considered responsive. DHIN prefers a single aware but remains open to multiple awards if appropriate.

Q: Section 8.4 (No Subcontracting) and Section 4.2 (Project Staffing): Will DHIN consider proposals submitted by a team of organizations under a single prime respondent? If so, does the written consent contemplated by Section 8.4 need to be secured before proposal submission, or is it sufficient to disclose the proposed teaming arrangement in the proposal and obtain written consent prior to the subcontracted work beginning under an awarded contract?

A: Yes, you may disclose proposed structure and obtain written consent prior to subcontracted work beginning.

Q: Section 4.1 (References): Two questions on references. First, will DHIN accept references for engagements performed by a named subcontractor on the proposing team, in addition to references for the prime? Second, will DHIN accept reference organization names, services performed, and period of engagement at submission, with individual contact person, title, and telephone number supplied if the proposal advances to the reference-check stage described in Section 5.1?

A: If the reference for the subcontractor is relevant to the scope of the project then it would be accepted.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Section 7 references a FHIR-based Clinical Data Repository and All Payer Claims Database supported by a third-party vendor. Will DHIN identify that vendor and indicate whether its data is in scope for the conversion and migration referenced in Section 2? Specifically, should respondents plan for clinical ingestion via bulk FHIR export from the existing CDR, via direct consumption of source HL7 v2 and C-CDA messages, or both, and what is the intended end state for the existing repository?

A: DHIN's preference for the FHIR CDR is to re-ingest all HL7 messages and C-CDA documents. The claims data may be imported from a back-up or ingested from source files depending on the proposal.

Q: Section 2.1.1 (Data Ingestion) and Section 7: For claims, should respondents plan pipelines sourcing raw X12 transactions, tabular data already processed into the existing APCD, or both? If both, please describe the intended use case for each path. Additionally, to support a flat fee proposal, will DHIN provide approximate annual and historical volumes for claims lines, enrollment records, and clinical documents currently held?

A: Claims are currently loaded from tabular data. X12 ingestion should be a current capability or included in a road map for any submission. Overall volume can be found in the current platform overview.

Q: Section 2.1.3 states DHIN would prefer to host the data warehouse in its own AWS platform. Under that hosting arrangement, do the availability, disaster recovery, and geographic redundancy requirements in Section 2.2, and the HITRUST or equivalent certification preference in Section 3.1.2, attach to DHIN's environment rather than to the respondent's organization?

A: DHIN's environment. Proposals should include guidance on supporting availability, disaster recovery, and geographic redundancy.

Q: Section 7 identifies Mirth Connect as DHIN's primary integration engine and Verato as the existing EMPI. Does DHIN intend for the proposed platform to replace Mirth Connect for HL7 v2 ingestion, or to receive data from it? For Verato, is the near real-time synchronization in Section 2.1.5 intended as a read integration in which the warehouse consumes resolved identity, or as a bidirectional exchange in which the warehouse contributes identity attributes back?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: We will not replace Mirth Connect for ingestion. The warehouse will consume resolved identities and update identities through the notification APIs.

Q: Will DHIN disclose the anticipated funding range or the Rural Health Transformation Program allocation supporting this procurement? Given the requirement that work be completed and invoiced by September 30, 2027, does DHIN anticipate separate funding for post-implementation support and enhancement beyond that date?"

A: DHIN does anticipate continued funding for post-implementation support and enhancements

---

Q: Data & Data Process Conversion: Can DHIN clarify whether "conversion of current data" refers to transforming unstructured data (e.g., HL7, CCD) into structured datasets or something else; what data is in scope for conversion; and the expected volume of data to be converted?

A: See current platform overview for volume. Expectation is that the system will transform HL7 and CCD to FHIR.

Q: Data Integration: The scope does not mention integration with data sources. Are vendors expected to integrate the platform with source systems/providers? How will data be made available for ingestion?

A: Mirth is our primary integration platform and is managed internally by DHIN's integration team.

Q: Data Integration: Please provide a list of data sources/providers that the data platform must integrate with.

A: Mirth is our primary integration engine and will handle ingestion of incoming data and integrations with data senders. Outgoing data is accessed from FHIR APIs, Internal custom APIs, Tableau, and standard SQL queries.

Q: Technical Architecture: Does DHIN have a preferred data warehouse platform (e.g., Databricks, Snowflake, AWS Redshift)?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Our current platform is Redshift, but we do not have a preferred platform.

Q: Technical Architecture: Is DHIN seeking a commercial off-the-shelf platform or is a custom solution on hyperscalers/cloud-native technologies acceptable? Has DHIN identified all required tools/components for the target architecture? If so, please provide details.

A: DHIN is looking for off-the-shelf platforms for FHIR interoperability but will consider custom solutions if based on existing platforms like HAPI or Couchbase.

Q: RHTP Flow-Down Provisions: Please provide the anticipated RHTP flow-down provisions or indicate when they will be available and what obligations respondents should incorporate into pricing and proposals.

A: TBC

Q: Award Contingency: If the State-DHIN RHTP contract is delayed beyond the anticipated award date (10/09/2026) or start date (11/01/2026), what notification process and timeline should respondents expect?

A: TBC

Q: In-Perpetuity Enhancement Obligation: Please clarify the scope of ongoing enhancement requirements “in perpetuity,” distinguishing regulatory/standards updates from separately commissioned services, and indicate the expected operating model (fixed capacity, fixed fee, T&M, etc.).

A: TBC

Q: Delivery Deadline of 09/30/2027: Does the September 30, 2027 deadline apply to the entire scope or only to specifically identified RHTP-funded deliverables?

A: RHTP funded deliverables.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Delivery Deadline of 09/30/2027: Would DHIN consider a phased approach that delivers non-RHTP scope after September 30, 2027?

A: Yes, as long as core functionality meets the deadline.

Q: Migration Trigger & Cutover: Please clarify when migration is considered applicable, which systems/services are in scope, the preferred cutover approach (big-bang, phased, parallel), any coexistence requirements, reconciliation/data-quality acceptance criteria, and how migration aligns with the September 30, 2027 deadline.

A: TBC

Q: FHIR Version Target: Is R5 expected at go-live or on a future roadmap? Are specific implementation guides (US Core, CARIN BB, Da Vinci, etc.) required? Does support mean ingestion, serving data, or both? Please define required FHIR versions, resources, profiles, operations, Bulk FHIR use cases, responsibilities, and implementation guides for initial acceptance.

A: TBC

Q: Ingestion End-State / MVP: For each ingested format, what level of processing is required at go-live versus roadmap (raw landing, discrete parsing, terminology standardization, cross-format mastering, MPI resolution, OMOP materialization)?

A: TBC

Q: Hosting Model Preference: Would DHIN consider platforms such as Databricks deployed within DHIN's AWS environment, or is a fully AWS-native approach preferred? Would either receive preferential evaluation?

A: DHIN would consider platforms deployed within DHIN's AWS environment

---

Q: OMOP Scope: Is OMOP required as a primary storage model at go-live, as a materialized export from the primary model, or simply enabled for future implementation?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: OMOP is not required at go-live. If it is not part of your existing platform, it can be excluded.

Q: Definition of Near Real-Time: Please quantify the target latency requirement and clarify whether it applies to ingestion-to-availability, query response, or both.

A: It applies to both and should be measured in minutes as opposed to hours.

Q: Accessibility Patterns / MVP Consumers: Which consumer groups must be supported at go-live, which retrieval mechanisms map to each group, what downstream systems require integration, and is an integration-ready posture required for external consumers?

A: TBC

Q: Verato EMPI Integration Details: Please describe the current Verato integration pattern and expected direction(s) of synchronization.

A: Verato is the master source of data and the data warehouse should only be pulling data. Additions and updates should be managed through the integration engine.

Q: USCDI Compliance Cadence: Is USCDI compliance a point-in-time requirement, ongoing SLA, or rolling standard? How will compliance be measured, and does it apply to CCD/C-CDA, FHIR US Core, or both?

A: USCDI is a rolling standard. We want to be within two versions of approved standards and have road maps in place for updating and testing version releases. It is relevant to CCD/C-CDA and FHIR US Core.

Q: Zero-Trust Architecture: Does DHIN expect alignment with a framework such as NIST 800-207 or specific control families (identity, device, network, workload, data)?

A: TBC

Q: Missing Attachments B and C: Are Attachments B and/or C available to bidders, or is the bidder expected to author the SOW/template?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Bidder is expected to author SOW and MSA

Q: MSA Baseline: For bidders without an existing MSA, does DHIN have a preferred MSA template or required terms it will provide?

A: DHIN can provide a MSA template.

Q: Page Limits and Format: Are there page limits, formatting standards, or appendix rules beyond the specified PDF/Word submission format?

A: No

Q: Baseline Environment Artifacts: Can DHIN provide architecture diagrams, data volumes, growth rates, and retention policies to support solution sizing and pricing?

A: see current platform overview

Q: Business Objectives & Success Measures: What business outcomes and measurable use cases must be operational by September 30, 2027 versus delivered in later phases?

A: TBC

Q: Timeline & Scope: What specifically must be completed and invoiced by September 30, 2027 (deployment, integrations, migration, UAT, training, support transition, etc.)?

A: TBC

Q: Scope & Requirements: Which capabilities are mandatory for initial acceptance, preferred, or roadmap/future enhancements?

A: Meeting current functionality with roadmap for adding capabilities.

Q: TEFCA / QHIN / National Exchange: Is DHIN expecting production connectivity by September 2027 or only an architecture and roadmap that supports future connectivity?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: DHIN is actively working with eHealthExchange on connectivity. It will be an expectation to continue with the new platform.

Q: Reporting & Analytics: Does scope include delivery of BI/reporting tools, dashboards, and semantic layers, or only data access and platform capabilities?

A: Data access and platform capabilities

Q: Users, Roles & Volumes: Please provide projected user populations, concurrent user estimates, and which stakeholder groups will directly access the platform.

A: TBC

Q: Data Ingestion Inventory: Please provide an inventory of expected interfaces and sources, including organization, feed count, format, status, and onboarding responsibilities.

A: TBC

Q: Data & Transaction Volumes: Please provide current and forecast storage volumes, record counts, daily transactions, growth projections, throughput requirements, and retention periods.

A: see current platform overview

Q: Claims Ingestion: Which X12 transaction types, implementation guide versions, claim categories, and historical periods are required for the initial implementation?

A: We currently ingest flat files for the APCD. Please see links for the current platform overview for details on the file templates.

Q: Integration Platform: What is DHIN's intended future role for Mirth Connect: retain, augment, or replace?

A: Retain

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Current-State Platform: Which existing components (FHIR CDR, APCD, PostgreSQL, Redshift, etc.) should be replaced, migrated, retained, integrated, or operated in parallel?

A: FHIR CDR and APCD will be replaced, other systems can be replaced or retained and integrated.

Q: Provider Directory: Should the proposed platform become the authoritative provider directory, or integrate with an existing solution? Please define hierarchy, stewardship, and update requirements.

A: TBC

Q: External Integrations: Please identify all organizations/networks requiring production integration and clarify responsibility for access, specifications, testing, certification, and coordination.

A: TBC

Q: Consent & Privacy: Please describe required consent, privacy, sensitivity-classification, segmentation, and enforcement requirements across APIs, queries, and extracts.

A: TBC

Q: Post-Go-Live Support: Please define expectations for hypercare, support hours, SLAs, monitoring, managed services, and division of responsibilities between DHIN and the selected vendor.

A: TBC

Q: Schedule Dependency: If award or project start is delayed due to RHTP contract timing, will the September 30, 2027 deadline remain fixed, and how should bidders account for that risk in their work plans?"

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Section 4.2 refers to an MSA attached as “Attachment C” but there is no Attachment C provided with the RFP. Please clarify.

A: TBC

Q: Section 4.5 directs respondents to submit proposals through the proposal submission form on the DHIN Website (https://dhin.org/who-we-are/rfps/). Can DHIN please provide a direct link to the proposal submission form as we are unable to locate it on the website?

A: Form will be posted by 9/23/26 in the same location as the LOI form.

Q: Is DHIN interested in a managed services arrangement under which the successful vendor provides monitoring and maintenance while DHIN hosts? If yes, what level of support might DHIN require?

A: DHIN is expecting some level of managed services, primarily around patches, upgrades, and improvements to the FHIR clinical data repository

---

Q: What are your top priorities for this modernization project, and what technology gaps are most prominent for your team today? Can you share your current vendor stack?"

A: see current platform overview. Data quality and HL7/C-CCDA mapping to FHIR are areas we are focusing on.

Q: Is there an incumbent or desired Enterprise Data Platform that is in place today or a preferred platform to be utilized in the development of the DHIN Healthcare Data Warehouse?

A: DHIN will evaluate all submitted platforms equally.

Q: Is the intent to have the platform implemented in alignment with the Core Capabilities described in section 2.1 by 9/30/27 with a subset of overall data sources flowing into it and then continue to add additional data sources over time (during the ongoing Maintenance and Support Phase)?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Yes, this is an accurate understanding of the RFP

Q: If yes, what is the initial subset of data sources that are expected to be live in the warehouse by 9/30/27?

A: FHIR model clinical data repository with API endpoint, All Payer Claims DB, and all data migrated.

Q: Is the intent for the platform to serve as a ""data warehouse"" that joins multiple datasets together for analytical querying, extraction, reporting, and dashboarding, or is the intent to build a ""data hub/data exchange"" that is serving as a transaction processing service for all the various data sources that it contains?

A: Both

Q: In section 2.1.2, the RFP states that ""Provenance must be maintained with each data element"" - does this mean that the intent is for the data warehouse to be built utilizing a data vault model?

A: DHIN is not currently using a data vault model and relies on logs and traceability of original sources to final state. Standardizing the provenance for all sources would be a benefit.

Q: In section 2.1.4 the RFP states that ""Near real-time data access & retrieval for all transactional data"" is required.

Is the requirement for real-time access to the data in the warehouse or is this saying that all sources feeding into the warehouse need to have ""near real-time"" ingestion from the source into the warehouse?

A: The FHIR clinical data repository needs “near real-time ingestion, data access, & retrieval.” The APCD and data for analytics can have a delay.

Q: The current timeline may limit or eliminate or prevent us from responding. Given the project's depth and scope, will DHIN grant an extension to the proposal due date?"

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: DHIN cannot grant an extension to the proposal due date because of time constraints placed on funding from the Rural Health Transformation Grant.

Q: Can DHIN share the specific RHTP flow-down provisions expected to apply (e.g., FAR/2 CFR 200 clauses, reporting, audit, Buy American, record retention), or a draft of the State–DHIN contract terms, so bidders can price compliance accurately?

A: TBC

Q: Is DHIN's expectation a full replacement of the existing CDR/APCD platform, an augmentation alongside it, or a greenfield warehouse that coexists with current systems? Which specific systems, if any, are being retired?

A: A full replacement of the existing CDR/APCD.

Q: Please quantify the historical data to be migrated — total volume (TB), record counts by domain (HL7 messages, CCDAs, X12 claims), date ranges, and current physical formats/schemas. Is migration in scope for the 09/30/2027 deadline?

A: Migration is in scope for the 9/30/27 deadline. Please see current platform overview for more details.

Q: Please clarify what ""ongoing product development & enhancement in perpetuity"" means contractually. Is that a commitment sought within the initial term, or a separate multi-year support agreement? What is the anticipated total contract term and renewal structure?

A: Support for patches, version upgrades, and HL7/C-CDA mapping to FHIR data elements. Support agreements should be detailed separately from the initial core implementation of the data warehouse.

Q: Is there a budget range or not-to-exceed amount for this engagement, and how is the RHTP award amount allocated between implementation and ongoing support?

A: Responders are asked to provide their typical annual, all-in costs and should separate implementation from ongoing support.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Does DHIN want implementation and ongoing maintenance/support priced as one flat fee, or as separate line items with distinct payment milestones?

A: Implementation and ongoing maintenance/support should be separate line items.

Q: Up to a million transactions daily during peak"" — can DHIN break that down by transaction type (inbound HL7, CCDA, API query, bulk export) and provide average vs. peak, plus expected growth over the next 3–5 years, for sizing and pricing?

A: see current platform overview

Q: What is the current Verato integration pattern (real-time API, batch, event stream)? Can DHIN share the interface specification and expected MPI transaction volumes? Does ""near real-time synchronization"" carry a specific latency target?

A: Expected latency between Verato and the data warehouse should be less than 5 minutes. Volumes can be found in the current platform overview.

Q: Is Entra ID integration required via SAML, OIDC, or SCIM (or all three)? Is DHIN's Entra ID tenant to be the sole identity provider for all warehouse users, including external stakeholder access?

A: TBC

Q: Is the OMOP CDM requirement a full parallel OMOP instance kept in sync, or an on-demand ETL/export to OMOP? Which OMOP version is targeted, and what is the expected refresh cadence and research user population?

A: TBC

Q: ""Near real-time data access & retrieval for all transactional data"" — what is DHIN's defined latency target (seconds, minutes) from message receipt to queryability? Does the same target apply to claims data as to clinical?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: ""Ability to build ad-hoc stored procedures, functions, and triggers"" — is this a DHIN-staff self-service requirement, and should bidders assume DHIN analysts will have direct SQL access to production? What SQL dialect(s) are acceptable?

A: TBC

Q: For DS4P — which sensitivity categories are in scope (42 CFR Part 2 SUD, behavioral health, HIV, reproductive health, minors), and is Part 2 compliance a contractual requirement?

A: TBC

Q: Regarding TEFCA/QHIN: is DHIN currently connected to a QHIN, and if so which one? Is the bidder expected to deliver QHIN-readiness within this engagement, or only a documented roadmap?

A: TBC

Q: Does DHIN require the platform to be FedRAMP authorized or StateRAMP listed? Is HITRUST certification of the bidder, of the platform, or of DHIN's hosting environment expected — and is it scored as a differentiator or a pass/fail?

A: TBC

Q: Will DHIN accept a bidder's standard MSA and SOW templates as the starting point, or does DHIN have its own preferred forms it would rather bidders mark up? Can DHIN provide those forms before 09/30/2026?

A: TBC

Q: ""All staff will be US-based"" — does this permit US-based staff who are not US citizens, and are there any background check, citizenship, or security clearance requirements tied to the RHTP funding?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: How much on-site presence at DHIN's Dover offices does DHIN expect, for which roles and at what cadence? Should travel costs be included in the flat fee or billed separately?

A: TBC

Q: Will DHIN consider a mutual indemnification and an aggregate liability cap (e.g., fees paid under the agreement), with customary carve-outs for confidentiality breach, IP infringement, and willful misconduct? As written the indemnity is one-way and uncapped.

A: TBC

Q: Regarding the requirement to ingest unstructured data and convert incoming care summaries/CCDs into discrete data elements: Does DHIN expect the proposed warehouse platform to natively include an automated clinical extraction / NLP pipeline capable of parsing narrative clinical text and static PDF documents into discrete, structured FHIR R4/R5 resources, or will DHIN's upstream integration engine (e.g., Mirth Connect) perform this extraction prior to warehouse ingestion?

A: TBC

Q: The RFP specifies that ""Provenance must be maintained with each data element."" Could DHIN clarify its expectations for data element provenance—specifically, is DHIN seeking cryptographic audit trails (e.g., content-addressed hash tracking/immutable versioning) linking transformed warehouse elements back to raw source payloads, or standard metadata timestamps and source facility identifiers?

A: TBC

Q: Regarding data segmentation (HL7 DS4P) and sensitivity tagging: Does DHIN require automated, attribute-based redaction/filtering at the query level to support multi-jurisdictional privacy mandates (such as sensitive reproductive or behavioral health records for cross-border patients), and should the warehouse support privacy-preserving/zero-knowledge query mechanisms for secondary research access?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Regarding support for the OMOP Common Data Model (CDM) for research: Is DHIN expecting an automated, continuous transformation pipeline (ETL) that standardizes inbound claims (X12) and clinical feeds (HL7/FHIR) into OMOP tables on a scheduled basis, and what is the anticipated scale/cadence for research cohort extraction (e.g., clinical trial feasibility queries)?

A: TBC

Q: Regarding near real-time synchronization with DHIN's existing 3rd-party EMPI (Verato): Can DHIN specify the current integration architecture with Verato (e.g., RESTful API webhooks vs. HL7 v2 PIX/PDQ feeds), and does DHIN expect the warehouse platform to manage and log identity mismatch resolution workflows locally?

A: TBC

Q: To accurately size the data migration effort, could DHIN provide the approximate historical data volume (in terabytes or total record/encounter counts) currently residing in the transactional PostgreSQL and AWS Redshift repositories (including claims from 2013 forward and clinical data from 2007 forward) that will be in scope for initial migration into the new warehouse?

A: TBC

Q: The RFP states a preference for hosting in DHIN's AWS environment. Can DHIN confirm whether the selected vendor will be granted access to deploy automated infrastructure-as-code (e.g., Terraform/CloudFormation) into a dedicated DHIN-managed AWS VPC, and does Section 8.6 apply solely to DHIN healthcare data and project-specific deliverables, leaving vendor core pre-existing IP and proprietary platform software protected?

A: TBC

Q: Given that funding is tied to the Rural Health Transformation Program (RHTP) requiring all work to be completed and invoiced by September 30, 2027: Will DHIN consider a milestone-based flat-fee structure for the implementation phase (through Sept 30, 2027)

paired with a separate recurring annual SaaS/maintenance fee for post-implementation support in subsequent years?"

A: Yes

Q: For some RFP capabilities, there may be a one-time implementation cost and an ongoing service cost that is dependent on data volume or transaction volume. Would DHIN prefer to include only implementation costs (to be completed by September 30, 2027) in the RFP?

A: TBC

Q: In RFP Section 2.1.1, there is a Data Ingest requirement for “Public Health Reporting.” Is it correct to assume that this is just the data ingest capability for ELR and syndromic surveillance data? In other words, development of the actual reports is not part of this RFP’s scope."

A: TBC

Q: (Sections 2, 7) Is the Data Warehouse intended to replace the existing third-party FHIR-based Clinical Data Repository and All Payer Claims Database, to coexist with them as an aggregation and analytic layer, or is that determination expected to come out of the solution architecture phase? If replacement is contemplated, what is the contract end date and renewal posture of the incumbent, and is transition expected before September 30, 2027?

A: TBC

Q: (Section 7) What is the intended future state for Mirth Connect, PostgreSQL, and Redshift? Should respondents assume Mirth remains the primary integration engine, that it is in scope for replacement, or that its disposition is open?

A: TBC

Q: (Section 2) On ""Conversion/migration of current data (if applicable)"": is migration of historical clinical data (2007 forward) and claims data (2013 forward) in scope, and must

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

full historical migration complete by September 30, 2027, or may historical backfill continue after go-live for priority use cases?

A: TBC

Q: (Sections 2, 4.5) The RFP uses product language throughout (""product enhancement in perpetuity,"" ""routine product upgrades/releases""). Will DHIN accept a response in which the core platform is a commercially licensed third-party technology (licensed by DHIN and running in DHIN's own AWS account), with the respondent providing architecture, implementation, integration, enablement, and ongoing managed services? Will such a response be scored equivalently to a single-vendor proprietary product response?

A: TBC

Q: (Sections 2.1, 2.1.4) DHIN acknowledges that not all platforms are built for all use cases. Will a composite architecture (a warehouse/lakehouse core combined with purpose-built components for FHIR API serving, IHE XCA exchange, and on-the-fly C-CDA generation) be evaluated favorably where the respondent is accountable for the integrated whole? Specifically, must the XCA and C-CDA capabilities be native to the warehouse platform, or may they be served by an adjacent interoperability component sourcing from it?

A: TBC

Q: (Sections 8.4, 4.2) Does the use of commercially licensed software or cloud services constitute subcontracting under Section 8.4? Relatedly, does the requirement that all staff be employees of the bidding organization permit platform-vendor professional services or support personnel in a defined advisory or escalation capacity, subject to DHIN's written consent?

A: TBC

Q: (Sections 2.1.3, 4.3) Who is expected to own and operate the AWS infrastructure and platform after implementation: DHIN staff, the selected bidder under a managed service, or a defined transition from bidder-operated to DHIN-operated? Is independent DHIN operation of the platform a stated goal of this engagement?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: TBC

Q: (Section 2.1.4) On ""Near real-time data access & retrieval for all transactional data"": what latency target does DHIN have in mind, and for which use cases? Is DHIN distinguishing between (a) sub-second point-of-care record retrieval, (b) minutes-latency operational reporting and event-based sharing, and (c) analytic and research workloads? Explicit tiering would materially improve architectural fit across all responses.

A: TBC

Q: (Sections 2.1.4, 3.1.1) For FHIR API serving including Bulk FHIR: does DHIN expect a fully conformant FHIR server (search, subscriptions, versioning, capability statement), or a FHIR-conformant API layer over warehouse data? Is ONC Health IT Certification of the FHIR endpoint a requirement or a preference?

A: TBC

Q: (Sections 2.1.1, 7) Can DHIN provide current and projected volumes by data type: annual HL7 v2 message counts, C-CDA counts, claims lines, total stored volume in TB, and expected growth? The RFP cites up to one million transactions daily at peak; a breakdown by type is needed to size and price accurately.

A: TBC

Q: (Section 2.1.3) For the OMOP requirement: does DHIN intend both clinical and claims domains to be represented in OMOP, and does DHIN hold (or intend to obtain) UMLS/SNOMED-CT licensing sufficient to distribute vocabularies to downstream researchers? Is OHDSI tooling (ATLAS, Achilles, Data Quality Dashboard) expected as a deliverable?

A: Yes clinical and claims within the OMOP model. It could be phased to incorporate clinical first, then claims. OHDSI tooling is not expected as a deliverable, but can be included. DHIN maintains vocabulary licensing and will utilize existing licenses and address any additional vocabularies that may be necessary.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: (Section 2.1.5) For the Verato EMPI integration: which Verato product and integration pattern is in use today (API, streaming, batch), and does DHIN's current Verato agreement accommodate the transaction volume and near-real-time synchronization contemplated here without a license expansion?

A: TBC

Q: (Section 2.1.6) How are consent and data sensitivity captured and enforced at DHIN today, including Delaware-specific consent policy and opt-out handling? Does DHIN currently hold, or expect to hold, data subject to 42 CFR Part 2? Is DS4P tagging applied today, or is establishing it part of this scope?

A: TBC

Q: (Section 2.2) For the SaaS-conditional requirements (≥99.9% uptime, RTO ≤ 4 hours, RPO ≤ 15 minutes, multi-region): if DHIN hosts in its own AWS account, do these apply as contractual SLAs for components under respondent management, or only to a vendor-hosted SaaS configuration? Where infrastructure is DHIN-operated, how does DHIN envision SLA accountability being allocated?

A: SaaS configuration

Q: (Section 3.1.2) Is HITRUST or equivalent certification expected at the level of the respondent organization, the underlying platform technology, or the deployed DHIN environment? Does DHIN hold a current certification or audit attestation for its existing AWS environment that the new platform would extend or inherit?

A: DHIN is currently HITRUST certified, but our current data warehouse is externally hosted and also HITRUST certified. Any SaaS would be expected to meet that standard and any internal system would need to be configured to maintain our current certification.

Q: (Section 3.1.2) The RFP references a data governance committee that will review CMS quality measure changes and approve implementation roadmaps. Does this body exist today? Is the respondent expected to stand it up, facilitate it on an ongoing basis, or participate in it?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: the DHIN data governance committee is currently in planning. Respondent is not expected to stand it up, but any functionality related to version control, data integrity, and validation should be included in the proposal.

Q: (Sections 3.1.3, 4.3) What is the expected scope of enablement beyond training materials? Specifically: approximately how many DHIN staff require training, by role (data engineering, operations, analytics, support, business)? And are external stakeholders (participating providers, payers, public health, researchers) in scope for training and adoption support?

A: TBC

Q: (Sections 4.3, 9) What DHIN internal resources will be available to the project, by role and approximate FTE allocation? DHIN-side capacity is a primary determinant of a credible fixed-fee work plan against a September 30, 2027 deadline.

A: TBC

Q: (Sections 2.1.1, 4.3) Will stakeholder organizations (hospitals, labs, payers, public health) be required to change their interfaces or data submission processes as part of this initiative? If so, is the respondent expected to manage that external engagement, communication, and cutover coordination?

A: TBC

Q: (Section 1.1) Is TEFCA/QHIN readiness work in scope for this engagement, or is the requirement that the architecture not preclude it? Has DHIN determined whether it will pursue QHIN designation directly versus participating through an existing QHIN?

A: TBC

Q: (Sections 1, 3.1.3) Are there specific, prioritized use cases or service lines DHIN intends to deliver first on the new platform? A stated priority sequence would allow all respondents to propose a more credible phased plan within the funding window.

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: (Section 2, cover note) RHTP funds require work to be completed and invoiced by September 30, 2027. Are software license, subscription, and cloud consumption costs allowable and invoiceable under RHTP, or is funding limited to implementation services? May multi-year licenses be prepaid before that date? Additionally, can DHIN share the anticipated RHTP flow-down provisions, or the expected timing of the State–DHIN contract, before proposals are due?

A: TBC

Q: (Section 4.4) Is there a budget range or not-to-exceed amount for this engagement? If DHIN prefers not to publish a figure, would DHIN indicate whether the anticipated cost is materially above or below a stated threshold?

A: TBC

Q: (Section 4.4) DHIN prefers a flat fee, but cloud infrastructure and consumption-based platform costs vary with data volume and workload. Would DHIN prefer respondents (a) present implementation services as flat fee with infrastructure and licensing separately estimated as pass-through, (b) include an estimated consumption allowance inside the flat fee, or (c) propose a not-to-exceed structure with reconciliation?

A: TBC

Q: (Sections 9, 8.10, 8.11) Is the 10% schedule-overrun penalty calculated against total contract value, the value of the delayed milestone, or remaining unpaid amounts? Will the contract include a change-control process confirming that DHIN-caused, stakeholder-caused, or third-party-caused delays adjust the baseline schedule? We also note an apparent tension between Section 8.10 (termination on 15 days' notice for any untimely performance) and Section 8.11 (15 business days to cure); can DHIN confirm the cure period applies before termination under Section 8.10?

A: TBC

Q: (Section 5.2) Will DHIN publish relative weightings for the seven scoring criteria?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: DHIN is not planning on publishing the relative weightings.

Q: (Section 5.1) ""DHIN reserves the right to accept a portion or portions of a proposal."" Is DHIN contemplating multiple awards across scope components (platform, implementation, ongoing support)? Should respondents structure proposals to be severable by component?

A: TBC

Q: (Section 4.5) The RFP notes respondents should call out preferred configurations, citing an example where a platform ""runs on AWS but runs more optimally on Snowflake."" Should respondents read this as openness to any major cloud data platform, or does DHIN have an existing preference, prior evaluation, or commercial relationship with a particular data platform vendor that respondents should be aware of?"

A: TBC

Q: Does DHIN expect all capabilities identified in Section 2.1 to be fully implemented by the September 30, 2027 project completion date, or may certain capabilities be delivered in subsequent phases as part of the product roadmap?

A: TBC

Q: Can DHIN provide estimated current and projected data volumes for claims, clinical data, HL7 transactions, FHIR transactions, and other major data sources to assist respondents in sizing and pricing the proposed solution?

A: See current platform overview.

Q: Can DHIN provide additional information regarding the volume, age, and structure of data expected to be migrated from the current clinical repository and All-Payer Claims Database environment?

A: See current platform overview.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: The RFP indicates a preference for hosting the data warehouse within DHIN's AWS environment. For solutions deployed within DHIN's AWS environment, what infrastructure, cloud services, security services, and operational responsibilities will be provided by DHIN versus expected of the selected vendor?

A: TBC

Q: For the required Development, UAT, and Production environments, does DHIN expect the selected vendor to provision and manage these environments, or will DHIN be responsible for providing and maintaining the underlying infrastructure?

A: Vendor management for SaaS, DHIN management for proposals intended for DHIN's AWS environment,

---

Q: Of the interoperability objectives identified in the RFP, including TEFCA alignment, QHIN readiness, and participation in national exchange frameworks, which capabilities are considered highest priority during the initial implementation phase?

A: eHealthExchange is our highest priority.

Q: The RFP references support for the OMOP Common Data Model for research purposes. Is OMOP support required at initial go-live, or may it be delivered as a later implementation phase?

A: Can be delivered later.

Q: Can DHIN provide additional details regarding the current integration patterns, APIs, and synchronization requirements with the existing Verato EMPI platform?

A: TBC

Q: Does DHIN envision the data warehouse becoming the system of record for provider directory management, or would integration with an external provider directory solution also satisfy this requirement?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Section 8.4 requires prior written approval for subcontractors. Does DHIN permit the use of subcontractors for specialized services such as implementation, integration, data migration, cloud engineering, or support services, and should proposed subcontractors be identified within the initial proposal response?"

A: TBC

Q: "Is there a required format or page limit for the technical proposal required elements?

A: No limit, PDF preferred

Q: Is there a preferred template / format for the statement of work?

A: PDF preferred

Q: Should the flat pricing fee include the ongoing maintenance in 2.2, or only the period of engagement ending on 09/30/27?"

A: Ongoing maintenance or support fees should be included separately.

Q: Given DHIN's stated preference to host the data warehouse within its own AWS environment, will DHIN consider or prefer solutions built on open, interoperable data formats that allow DHIN to retain direct ownership and access to underlying data files within DHIN-controlled S3 buckets?

A: Yes, DHIN will consider a non-AWS solution with full access to data and user roles and permissions.

Q: Additionally, should respondents address the ability to support governed access to DHIN data through multiple approved analytics, integration, and operational tools without requiring proprietary data extraction processes, unnecessary data duplication, or a single-vendor access model?

A: Yes

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: To support advanced population-health analytics, predictive modeling, clinical NLP, and research-grade OMOP Common Data Model transformations, will DHIN evaluate platforms based on their ability to provide integrated support for machine learning and scalable Python, R, SQL, and distributed data-processing workflows directly against governed platform data?

A: It will not be a major component of the evaluation, but will be factored into the final decision.

Q: Please clarify whether DHIN expects advanced analytics and OMOP transformation workloads to operate within the selected platform, or whether integration with separate analytics, data-science, or model-development tools is acceptable.

A: The selected platform does not need to address advanced analytics or OMOP transformations, but should be robust enough to handle the computational resources necessary to support those tools.

Q: As Delaware's designated Health Data Utility, DHIN shares data with diverse stakeholders, including state agencies, academic researchers, healthcare providers, payers, public-health entities, and other authorized parties. Will DHIN evaluate platforms based on their support for secure, standards-based, vendor-neutral data-sharing capabilities that allow authorized external stakeholders to access governed datasets in real time or near real time?

A: Yes

Q: If so, should respondents address the ability for recipients to access authorized data without requiring the recipient to adopt a particular proprietary software platform, while maintaining DHIN's control over privacy, consent, auditing, authorization, and revocation?

A: Yes, specifically related to the FHIR CDR.

Q: Considering that DHIN manages structured claims and HL7 feeds, semi-structured FHIR resources, unstructured C-CDA clinical documents, and analytical models, does DHIN

expect a centralized governance capability that supports unified data cataloging, row- and column-level security, attribute-based access control, audit logging, data lineage, and data-quality monitoring across these assets?

Please also clarify whether DHIN has existing governance, identity, security, metadata, or logging tools with which the selected platform must integrate.

A: TBC

Q: Given the requirement to ingest high-volume HL7 v2.x, FHIR R4/R5, and Bulk FHIR datasets, and DHIN's current peak volume of up to one million daily transactions, will DHIN evaluate solutions based on their ability to provide streaming or event-driven ingestion, automatic compute scaling during peak periods, and efficient operation during lower-volume periods?

If so, are there target throughput, latency, availability, recovery, or processing-performance measures respondents should address in their proposals?

A: TBC

Q: DHIN's organizational background describes a current ecosystem including AWS, Mirth Connect, PostgreSQL, Redshift, a FHIR-based Clinical Data Repository, and an All-Payer Claims Database supported by a third-party vendor. Is the intent of this RFP to replace all or some of these components, including the third-party CDR/APCD, or to augment the current environment with a new enterprise data warehouse and analytics platform?

A: Mirth is outside of the scope of this RFP, but a FHIR based clinical data repository and integrated solution with the APCD are the focus of solutions.

Q: If replacement, consolidation, or migration is in scope, please identify the components DHIN expects to address during the initial implementation versus components that will remain in place and integrate with the selected solution.

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: What service-level targets, latency ranges, or performance expectations should respondents use for "near real-time" processing and retrieval? If possible, please distinguish expectations for:

- HL7 v2.x ingestion and availability; **near real time**
- FHIR API ingestion and retrieval; **near real time**
- EMPI synchronization and identity updates; **near real time**
- Claims ingestion and availability; **Future state for "near real time"**
- Longitudinal patient-record retrieval: **near real time**
- Downstream reporting, analytics, and exchange workflows. **24 Hours**

A: See bold responses above. Near real time should be on the order seconds to minutes

---

Q: The RFP states that the solution must support all FHIR APIs, including Bulk FHIR retrieval. Could DHIN clarify the intended functional scope? Specifically:

A: TBC

Q: Does DHIN expect a FHIR server and persistence layer, a FHIR API layer over the warehouse, integration with an existing FHIR service, or a combination of these capabilities?

A: Yes DHIN expects a FHIR server

Q: Which FHIR interactions and operations are expected at implementation, such as read, search, create, update, delete, history, subscriptions, transactions, terminology services, and Bulk Data $export?

A: TBC

Q: Which FHIR versions, implementation guides, and profiles are required at go-live, including R4, R4B, R5, US Core, USCDI, Da Vinci, or public-health implementation guides?

A: TBC

Q: Are Bulk FHIR import and export both required?

A: Import no, export yes

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Are there expected security, identity, or authorization requirements, such as SMART on FHIR, OAuth 2.0, or UDAP?

A: TBC

Q: What volumes, frequency, and performance targets should respondents assume for Bulk FHIR operations?

A: TBC

Q: DHIN notes that it maintains an MPI of more than five million people, approximately three million active records, and handles up to one million daily transactions during peak usage. Does DHIN expect transaction volumes, data volume, source-system count, query concurrency, or Bulk FHIR activity to grow over the contract term?

A: Yes to all examples.

Q: If available, please provide planning estimates or ranges for expected annual growth in HL7, FHIR, claims, C-CDA/document, API, and other transaction volumes; new data sources; historical data-retention requirements; peak and typical user or application concurrency; and Bulk FHIR ingestion and retrieval activity."

A: TBC

Q: The RFP states DHIN ""would prefer to host the data warehouse in its own AWS platform, but [is] open to alternative hosting arrangements based on the solution provided."" How much weight will this preference carry in scoring, will a non-AWS proposal be evaluated on equal footing if it meets all other requirements?

A: Yes, other options will receive equal consideration.

Q: Section 2.1.3 also lists ""support for storage of data as an OMOP common data model for research related purposes."" Is this an active, existing use case with a current research partner, or is it a future-proofing requirement with no current research use case attached?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: We have current use cases, but do not have an OMOP CDM. It does not need to be part of the submitted proposal, but the capability to implement OMOP in our data warehouse is necessary.

Q: Regarding the required integration with DHIN's existing EMPI platform (Verato): does DHIN have an existing API/interface specification or integration guide available for Verato, or would defining that integration be part of the development scope under this engagement?

A: There is an existing API endpoint with comprehensive documentation.

Q: What level of ongoing managed services and support should vendors assume as part of their base pricing versus what should be proposed as an optional add-on? For example, should base pricing assume standard platform maintenance and support only, with elevated SLAs, dedicated support staff, or enhanced managed services quoted separately?

A: TBC

Q: Section 4.1 requires ""references from at least three organizations for which the bidder has performed similar work"" and notes that ""organizations that resemble DHIN are preferred."" Can DHIN confirm whether references must reflect healthcare-specific data warehouse/interoperability work, or whether references demonstrating similar technical scope (e.g., data warehousing, systems integration, regulatory/compliance reporting) in other sectors would satisfy this requirement?

A: At least one reference should be able address experience with health care data.

Q: Is the minimum of three references a firm requirement, or would two strong, closely relevant references be acceptable in lieu of three if a bidder's most applicable experience is concentrated in fewer engagements?

A: DHIN will accept two references if they are from an HIE or HDU.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Section 7 describes DHIN's current environment as including a ""FHIR-based Clinical Data Repository and All Payer Claims Database supported by a third-party vendor."" Is the intent of this RFP for the new data warehouse to replace the existing CDR/APCD (or any other current data system), or would it operate alongside them?"

A: It is intended to replace them.

Q: Is DHIN looking to replace its current Clinical Data Repository and All Payer Claims Database (or any other data system), or would the new data warehouse work alongside them?

A: We are looking at replacements.

Q: Does DHIN expect all existing data to be migrated by September 30, 2027, or would a phased migration be acceptable?

A: All existing data will need to be migrated by September 30th. All existing FHIR functionality and transformations will also be required.

Q: Can DHIN share the approximate amount of data that will need to be stored, along with expected future growth?

A: TBC

Q: What level of ongoing managed services and support should vendors include in their base pricing?

A: TBC

Q: How much weight is put into the preference for hosting in their own AWS warehouse?

A: DHIN will evaluate all proposals regardless of platforms equally.

Q: Is the support for OMOP CDM an existing use case with a research partner or future-proofing?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: OMOP is not required as part of the submission.

Q: Does Verato have an API/interface guide together or would that be part of development?"

A: Our EMPI vendor has extensive API documentation for additions, updates, queries, maintenance, and notifications averaging 33 million API calls per month.

Q: Is DHIN seeking primarily a replacement for its existing clinical data repository and data warehouse, a new enterprise platform that will coexist with selected current systems, a modernization of the broader HIE architecture, or a phased combination of these objectives?

A: TBC

Q: Which existing platforms does DHIN expect to replace, retain, or integrate with, including Mirth Connect, PostgreSQL, Redshift, the current FHIR repository, APCD platform, Verato, analytics tools, and internally developed applications?

A: TBC

Q: Does DHIN expect a single platform to support both operational HIE workloads and analytical/research workloads, or may respondents propose separate but integrated operational and analytical platforms?

A: TBC

Q: Please provide current and projected volumes for HL7 messages, CDA/C-CDA documents, FHIR resources, API transactions, X12 claims, public health feeds, files, and historical data to be migrated.

A: TBC

Q: What transaction volumes, concurrent users, API request rates, and query workloads should the solution support at go-live and for future growth?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: TBC

Q: What measurable definition of ""near real-time"" should be used for data ingestion, longitudinal record updates, FHIR access, analytics, notifications, and Verato synchronization?

A: TBC

Q: What performance and response-time requirements exist for FHIR APIs, Bulk FHIR exports, SQL queries, materialized views, longitudinal record retrieval, and C-CDA generation?

A: TBC

Q: Which systems and repositories will require migration, and what are the associated data volumes, retention requirements, formats, and data quality considerations?

A: TBC

Q: Is all historical data expected to be migrated, or is a defined historical lookback period acceptable?

A: All historical data.

Q: Will incumbent vendors be required to provide complete exports, data dictionaries, mappings, and technical assistance to support migration?

A: TBC

Q: What reconciliation and acceptance criteria will be used to validate migrated data?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: How many production interfaces and data sources are currently in scope for implementation?

A: We have approximately 120 data sources, but the primary integration engine for ingesting data from these sources is Mirth.

Q: How many new interfaces and participating organizations are expected to be onboarded prior to project completion?

A: TBC

Q: Will DHIN provide existing interface specifications, routing rules, mappings, message samples, and test scripts?

A: TBC

Q: What specifically is meant by “the entire FHIR spectrum” and “all FHIR APIs”?

A: All core profiles

Q: Is support for FHIR R5 required at initial go-live, and must R4 and R5 coexist simultaneously?

A: R4 is required at initial go-life. Simultaneously running R4 and R5 for a window of time would be preferable, but a more traditional upgrade is also acceptable.

Q: Which FHIR implementation guides, profiles, and operations are required?

A: US Core – required, SMART on FHIR future state, QI-Core/DEQM future state. All Core foundational profiles, Read, VRead, Search (get & post)

Q: Is SMART on FHIR required?

A: SMART on FHIR capability is preferred, but not required. Future compatibility is required.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: What are the expected Bulk FHIR use cases, data volumes, frequencies, and performance requirements?

A: Bulk FHIR is not currently being used. We do not have a defined scope yet.

Q: Should claims data be incorporated directly into the operational longitudinal record, maintained separately for analytics, or both?

A: Both

Q: For consolidated C-CDA generation, what templates, claims content, implementation guides, and consumer requirements must be supported?

A: FHIR R4 to USCDI v3.1

Q: Is OMOP expected to be a primary repository, a research repository, an analytical mart, or a downstream export?

A: research repository with ATLAS as an interface and ability to export cohort data.

Q: Which OMOP version, vocabularies, refresh frequency, historical depth, and validation requirements are expected?

A: OMOP

Q: Will DHIN provide OMOP vocabularies and mappings, or should respondents include those capabilities?

A: DHIN will provide OMOP vocabularies and mappings.

Q: What specific Clinical Quality Language (CQL) capabilities are required?

A: Clinical Quality Measurement is the primary requirement.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Does DHIN require execution of eCQMs, clinical decision support, CDS Hooks, or other CQL-driven applications?

A: No, the clinical data repository should have CQL capability, but integration with CQL-driven applications is not expected at go live.

Q: Which initial clinical quality measures and clinical decision support use cases are anticipated?

A: Risk scores, integration of SDOH data, and closing care gaps.

Q: Will Verato remain the authoritative master patient index during the implementation and beyond?

A: Yes

Q: What integration methods, transaction volumes, latency requirements, and governance processes currently exist for Verato?

A: Our EMPI vendor has extensive API documentation for additions, updates, queries, maintenance, and notifications averaging 33 million API calls per month. Latency requirements should be 10 minutes or less.

Q: Which system should be considered authoritative for patient identifiers, merges, unmerges, and identity governance?

A: Our empi system. The data warehouse will need to be in sync with that.

Q: Does DHIN expect the selected solution to become the provider directory system of record or simply integrate with an existing solution?

A: It will ideally be the master source of data, but does not need to have an interface.

Q: What provider directory relationships, affiliations, hierarchies, endpoints, and organizational structures must be supported?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: Organization and Provider affiliations. NPI, Taxonomy, and State Licensing information.

Q: What discrepancy resolution and governance processes should be used when provider information differs between DHIN and external sources such as NPPES or CMS?

A: Notifications and change logs

Q: Please provide the consent policies that must be enforced, including any organization-specific, source-specific, data-class-specific, or purpose-of-use-specific requirements.

A: TBC

Q: At what level must privacy controls and data sensitivity restrictions be enforced: document, resource, section, field, or individual data element?

A: TBC

Q: What DS4P implementation guides, security labels, and privacy requirements are expected?

A: TBC

Q: Must consent and sensitivity controls apply consistently across FHIR APIs, SQL access, Bulk FHIR, analytics, exports, and document generation?

A: Yes expectations are to use FHIR Confidentiality Codes and Sensitivity Tags on the data with the ability to the ability to access control rules. Automated data tagging capabilities are preferred.

Q: If deployed into DHIN's AWS environment, what account structures, regions, connectivity requirements, security controls, and approved AWS services should respondents assume?

A: DHIN will manage infrastructure and security controls for systems deployed in DHIN's AWS environment.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Does DHIN expect the selected vendor to provide software only, implementation services, managed services, or full operational support?

A: DHIN is evaluating all options

Q: How should responsibilities be divided between DHIN, the selected vendor, AWS, incumbent vendors, and any approved implementation partners?

A: Responsibilities will be dependent on submitted proposal.

Q: Does the 99.9% uptime requirement apply only to SaaS deployments or also to customer-hosted deployments?

A: SaaS deployments

Q: Do the stated RTO and RPO targets apply to all platform capabilities and data types?

A: TBC

Q: Is active-active multi-region architecture required, or is a primary/secondary disaster recovery model sufficient?

A: We currently use a multi-availability zone architecture, but will considera more traditional redundancy model.

Q: What specifically constitutes project completion by September 30, 2027?

A: Data migration, existing functionality matched, testing and validation plan completed.

Q: Does completion require all interfaces migrated, all data converted, production go-live, final acceptance, training completion, and operational transition?

A: TBC

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Q: Will phased implementation approaches be considered if core capabilities are completed within the funding window?

A: Yes

Q: What DHIN resources will be assigned to architecture, testing, data mapping, validation, security, and acceptance activities?

A: Executive sponsor, project manager, dedicated data engineer, and varying effort from data quality analysts, integration engineers, a solutions architect, and applications analysts.

Q: What turnaround times can DHIN commit to for reviews, approvals, testing, issue resolution, and acceptance of deliverables?

A: DHIN is committed to the project and will provide timely responses and remain actively engaged. Detailed timelines are dependent on volumes and complexity.

Q: Will delays caused by DHIN, participating organizations, incumbent vendors, or state approvals be excluded from schedule-overrun penalties?

A: TBC

Q: What RHTP flow-down requirements are anticipated and when will those requirements be made available to bidders?

A: TBC

Q: If state contracting delays occur, will implementation deadlines and completion milestones be adjusted accordingly?

A: TBC

Q: Are implementation partners permitted if approved by DHIN, or must all project resources be direct employees of the bidder?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: TBC

Q: Does the ownership clause apply only to DHIN data and project deliverables, or would it also apply to pre-existing vendor intellectual property, software, templates, data models, and methodologies?

A: TBC

Q: What evaluation criteria will carry the greatest weight among technical capability, implementation approach, healthcare/HIE experience, pricing, and long-term support?

A: Healthcare experience, technical capability, implementation approach, pricing, long-term support.

Q: Does DHIN have a preferred pricing model, such as flat fee, subscription-based, managed services, consumption-based pricing, or a combination of these approaches?

A: DHIN hosted would be flat fee, managed services would be a combination of approaches.

Q: Should respondents provide separate pricing for software, cloud infrastructure, implementation services, migration services, managed services, and optional capabilities?

A: Respondents should include as much detail in pricing as possible and clearly detail one time vs. ongoing costs.

Q: What planning assumptions should respondents use for future growth in patient population, participating organizations, transaction volumes, and analytical workloads?

A: DHIN is looking to expand participating organizations to include more Ambulatory practices and growth expectations should be significant over the next three years.

Q: What future use cases beyond those listed in the RFP are driving DHIN's vision for a statewide Health Data Utility over the next three to five years?

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

A: DHIN is currently the state HIE and looking to continue and expand support for practices, providers, patients, and the state by adopting the HDU model.

Q: Does DHIN envision the selected platform serving as the foundation for future TEFCA participation, QHIN connectivity, public health modernization, value-based care initiatives, and research activities?

A: Yes, the platform will be the foundation of our data ecosystem and will support all DHIN initiatives as an HIE and APCD supporting the state of Delaware.

Q: What would DHIN consider the most important measure of success one year after implementation is completed?"

A: A robust platform that is scalable, secure, and reliable with a clear roadmap for maintaining data quality, completeness, reporting, and analytics.

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

Current Platform Overview

AWS Cloud Environment (US-East-1 VA and US-East-2 OH)

- Integration Engine: Mirth (internal)
- Transactional DBs: Postgresql (internal)
- Datawarehouse: Redshift (internal & third party)
- Clinical Data Repository: Couchbase FHIR (third party)
- Document Storage: S3 & SFTP (internal & third party)
- EMPI: Verato (third party)

Storage:

- Databases (Claims, Clinical RDS, logs, utilities): > 40TB
- FHIR Repository: ESTIMATE 15TB based on 3.5M unique clinical patient data and size of compressed archive.
- RAW HL7 & CCD files historical archive: 15TB compressed

Volume:

- HL7:
    - 393M messages from 2014 to present
    - Current 4.7M messages per month (~10-12% growth per year)
    - Currently ingest ADT & ORU only, capable to ingest ORM, VXU, MDM, and most others. Looking to expand what we receive and ingest.
- CCD:
    - 14M documents from 2019 to present
    - Current 220K documents per month (2027 expected growth ~20%)
    - USCDI v3 capable
- EMPI
    - 5.5M unique identities, 32.9M source and native ID records
    - 611 sources (includes some with multiple under a single parent organization)
    - ~6% growth per year
- APCD
    - 700M records across medical, pharmacy, eligibility, and provider
    - 1.8M unique patient identities

API:

- EMPI: 32M calls per month
- FHIR: ~4M calls per month

Submitted Question Responses for the DHIN Data Warehouse RFP.

Last Updated: 09/21/2026

More Information:

- https://dhin.org/dhin-data/
- https://dhin.org/resources/