

Neighborhood House**Request for Qualifications****Release Date: 9/21/2026****Number:****Service Needed: Data Architecture and Engineering****About Neighborhood House**

Neighborhood House is one of the oldest social service agencies in the Puget Sound region and has a long history of working with King County's immigrant, refugee and low-income communities. Our mission is to partner with diverse individuals and families to build community and achieve their goals for health, education and self-sufficiency. Last year, Neighborhood House served over 12,000 individuals of all ages through programming in the areas of early childhood education, youth development, family & social services, employment & adult education, housing stability, aging & disability services and community health.

Project Overview & Objectives

Neighborhood House is seeking proposals from qualified data architects, data engineers, and data engineering firms to design a modern, scalable data architecture and execute data integration across key organizational systems.

The primary goal is to establish a unified data platform on Microsoft Azure, ingest datasets from the SaaS systems holding our client and service delivery data (Apricot and ChildPlus), and build an integrated model capable of entity matching, deduplication, reporting, and data governance across both platforms. Microsoft Fabric is our anticipated direction for the central repository, but we are not committed to it and welcome alternatives, including a staged approach that adopts Fabric later.

An internal proof of concept for the Apricot pipeline is underway. The technical direction described below reflects our current thinking, not fixed decisions - we expect the selected partner to review our requirements and validate, revise, or replace that direction.

Fixed Constraints

- Apricot data is delivered over SFTP.
- ChildPlus data resides in a remote SaaS-hosted SQL Server
- The platform will store and process highly sensitive client data, including data subject to HIPAA. A signed Business Associate Agreement is required before any access to production data.
- Reporting is delivered through Power BI today and will continue to be, against the new platform.

Scope of Work

Phase 1: Architecture Consultation & Strategy Alignment

- Consult with Neighborhood House leadership and stakeholders to review current data sources and long-term organizational data strategy.
- Design a cloud-native, scalable analytics architecture on Azure.
- Recommend the appropriate Azure services for our requirements, data volume, and staffing - including whether Microsoft Fabric is the right platform for the central repository now, later, or not at all, and where else our current direction should change.
- Establish a layered repository design in which data from each source system lands first in raw, untransformed form that faithfully reflects the source, with all transformation applied in successive layers built on top of it.
- Define data governance, security, access control, and pipeline standards for future scalability.
- Design for long-term maintainability by a single in-house data engineer: infrastructure defined as code, version controlled and repeatable; documented, reviewable changes; no undocumented manual configuration.
- Account for anticipated future layers, notably automated program-specific transformations for grant and funder reporting, which are not in scope for this engagement.

Phase 2: Apricot System Data Ingestion

- Design a recurring, automated ingestion process for the nightly Apricot SFTP dump, including landing, validation, transformation, and load into the target environment.
- Stand up and configure the landing database required to execute the vendor MySQL database dump files.
- Land the extracted data in raw form, then build a transformation layer that uses Apricot's own metadata tables to resolve its generic table and field identifiers into meaningful names, exclude inactivated fields, and express record relationships in a structure analysts can query directly.

Phase 3: ChildPlus System Data Ingestion

- Extract program and participant records from the ChildPlus SQL database on a recurring basis.
- Land the extracted data in raw form, then build a transformation layer that removes unneeded tables and columns and aligns naming and structure with the equivalent Apricot layer where the two systems describe the same concepts.

Phase 4: Data Integration & Deduplication Framework

- Design a consolidated data model that combines Apricot and ChildPlus datasets into a single source-agnostic view of clients and the services they receive.
- Implement entity matching, resolution, and deduplication logic that records matched identities and a confidence rating for each match, and persists those match decisions so they are not recomputed from scratch on each refresh.
- Provide the data structures needed to support future staff review of uncertain matches, including storage for staff decisions and the ability to audit who changed a match and when. Building a user-facing review application is not in scope for this engagement.

- Ensure that staff overrides recorded in the future take precedence over automated matching results and survive recurring pipeline refreshes without being overwritten.
- Validate data accuracy, completeness, and query performance in the final target repository.

Required Technical Qualifications & Experience

We are seeking breadth across the Azure data ecosystem and the judgment to recommend the right tools for an organization of our size. Experience in the following areas is expected so that the consultant can evaluate and, where appropriate, challenge our current direction:

- **Azure Data Engineering:** Building scalable, secure data pipelines on Azure, with enough breadth across the available services to compare options and recommend a fit for our requirements rather than a default.
- **Microsoft Fabric:** Hands-on development of end-to-end pipelines and analytical stores in Fabric, with Power BI as the consumption layer, and a candid view of where Fabric is and is not the right choice today.
- **Layered Data Architecture:** Designing repositories that separate raw source-faithful data from conformed and curated layers, and maintaining that separation through automated, repeatable refreshes.
- **Cloud-Native Architecture:** Scalable and secure cloud design principles, including containerized workloads where appropriate and scheduled job orchestration.
- **Secure Azure Networking:** Private connectivity between Azure services, controlled outbound egress, and site-to-site VPN connectivity to external systems.
- **Secrets & Identity:** Centralized secrets management with rotation and least-privilege, identity-based access - no credentials embedded in source code, images, or scripts.
- **Database & Connectivity Expertise:** Advanced SQL skills, experience restoring and executing MySQL dumps, and working with remote SQL Server connections over VPN.
- **Data Modeling & Entity Resolution:** Dimensional and integrated data modeling, ETL/ELT pipeline design, and multi-source record matching and deduplication.
- **Regulated Data:** Handling of sensitive data subject to HIPAA or comparable requirements.

Deliverables

1. **Architecture Blueprint:** System topology, data flow diagrams, and a platform recommendation with rationale, including cost and long-term maintainability tradeoffs.
2. **Ingestion Pipelines:** Fully documented, repeatable ETL/ELT pipelines for both Apricot and ChildPlus, defined as code and version controlled.
3. **Unified Data Repository:** A layered analytics repository holding raw source data, system-specific transformed layers, and the combined cross-system model, consumable by Power BI.
4. **Deduplication Engine:** The matching framework identifying and flagging duplicate profiles across sources, with match confidence and match decisions retained and structured for future staff review and override. Development of the staff-facing review interface is excluded.

5. **Operational Readiness:** Monitoring and alerting, runbooks for routine operation and failure recovery, and least-privilege access design. Recurring refreshes must be validated before replacing current data, with the ability to roll back.
6. **Technical Handoff & Training:** Hands-on training sufficient for our in-house data engineer to independently operate, troubleshoot, modify, and extend the platform; operational cross-training for a systems engineer covering routine recovery and escalation; and orientation for agency data analysts connecting Power BI to the new platform.

Evaluation Criteria

- Meets Qualifications.
- Level of Experience (record of past performance).
- Availability.
- Business Integrity.
- Cost/Rate for Services.

Questions

Alejandro Chavarria <AlejandroC@NHWA.ORG>

What to Include In Submissions

- Rates: Price per Hour.
- Description of Experience.
- Others materials: Bidders discretion.

Submission Close Date: 11/15/2026

Submission Requirements

Proposals should include:

- * Company / Individual Profile: Resume or firm overview showcasing relevant past projects.
- * Technical Approach: Brief summary of how you plan to approach the four project phases.
- * Cost Estimate: Breakdown by phase or milestone, along with hourly rates for additional scope.
- * Timeline: Estimated start date and schedule to completion across all four phases.
- * References: Contact details for 2–3 recent clients with similar scopes of work.

Additional Requirements (if selected): Neighborhood House may require any or all of the following: a written contract; proof of professional and/or general liability insurance; pass Background Checks, clear a check of the federal EPLS (Excluded Parties List System), proof of Washington State business license or comparable license from another state; other checks and reviews appropriate to the nature and scope of the engagement.