

PopMedNet: A scalable and extensible open-source informatics platform designed to facilitate the implementation and operation of distributed health data networks

Daniel Dee¹, Dean Corriveau¹, Jeffrey R. Ladd¹, Lakshmi Nannapaneni¹, John Arkins¹, and Jeffrey Brown²

¹ General Dynamics Information Technology ² Harvard Pilgrim Health Care Institute

DOI: [10.21105/joss.04062](https://doi.org/10.21105/joss.04062)

Software

- [Review](#) ↗
- [Repository](#) ↗
- [Archive](#) ↗

Editor: [Daniel S. Katz](#) ↗

Reviewers:

- [@jhancock4d](#)
- [@Irasmus](#)

Submitted: 22 October 2021

Published: 04 March 2022

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC BY 4.0](#)).

Summary

PopMedNet™ is an open-source application used to facilitate multi-site health data networks ([Brown et al., 2010](#)). It uses a distributed network design that enables data partners to retain full control of their data. Investigators send questions to data partners for review and response. PopMedNet eliminates the need for assembling patient records in a centralized repository, thus preserving patient privacy and confidentiality.

Its development was pioneered by the Therapeutics Research and Infectious Disease Epidemiology (TIDE) group of the Department of Population Medicine (DPM) of the Harvard Pilgrim Health Care Institute (HPHCI), an appointing department of the Harvard Medical School.

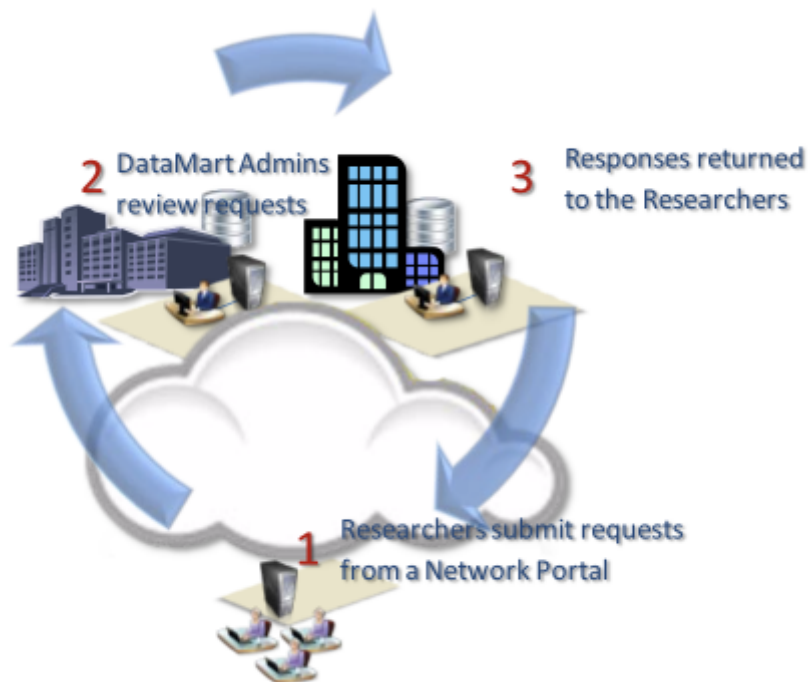


Figure 1: PopMedNet.

Statement of need

Distributed health network plays an important role in supporting a number of population health activities ([Brown et al., 2010](#)).

The goal of PopMedNet™ is to facilitate distributed analysis of electronic health data in order to support medical product safety surveillance, disease surveillance, comparative effectiveness, quality, medical resource use, cost-effectiveness, and related studies. It does this by enabling the creation of health data networks that allows the data owners to keep and control uses of their data, while at the same time allows customized access and governance for each network ([Davies et al., 2016](#); [Malenfant et al., 2015](#)).

The concerns of the network stakeholders are addressed in the design and governance models of the PopMedNet™ software platform.

PopMedNet™ features include distributed querying, customizable workflows, and auditing and search capabilities. Its flexible role-based access control system enables the enforcement of varying governance policies.

Architecture

PopMedNet™ has two components. It has a .NET web application component (Portal) and a Windows desktop application component (DataMart Client). The .NET web component uses a modern service oriented architecture (SOA) with a robust API serving a dynamic UI using the MVVM pattern. The DataMart Client (DMC), uses a plugin architecture to support different types of adapters for querying data. A variety of adapters are available, and able to support file transfer, ad-hoc SQL queries, and dynamically composed queries against local data sources (including SQL Server, Postgres and Oracle).

PopMedNet™ is a .NET Framework application, utilizing ASP.NET MVC5 for the implementation of the website application and WinForms for the desktop application. Entity Framework 6 is used for data access, and Typescript for client-side scripting.

How does PopMedNet™ work?

In PopMedNet™, the Portal handles all the governance and acts as a hub for all queries. That ensures that the coordinating center can expect:

- Consistency of information across data partners
- Quality of data
- Flexibility to handle different network structures with overlapping data partners
- Ability to track and report operational activities
- Managing multiple projects

The DMC allows data partners to control access to their data before sending it to the Portal. For the data partners, it will:

- Provide protection of Protected Health Information (PHI is defined as any personal health information that can potentially identify an individual)
- Prevent loss of control over proprietary information

The researcher will have no direct access to data but is shielded from the differences among networks' data models.

A typical PopMedNet™ request-response cycle is as follows:

1. Request: Investigators distribute data requests constructed in the point and click web interface

2. Review: Data Partners receive and review requests in their DataMart Client
3. Respond: Data partners execute the requests against their local data and upload results
4. Download: Investigators access results in the query tool

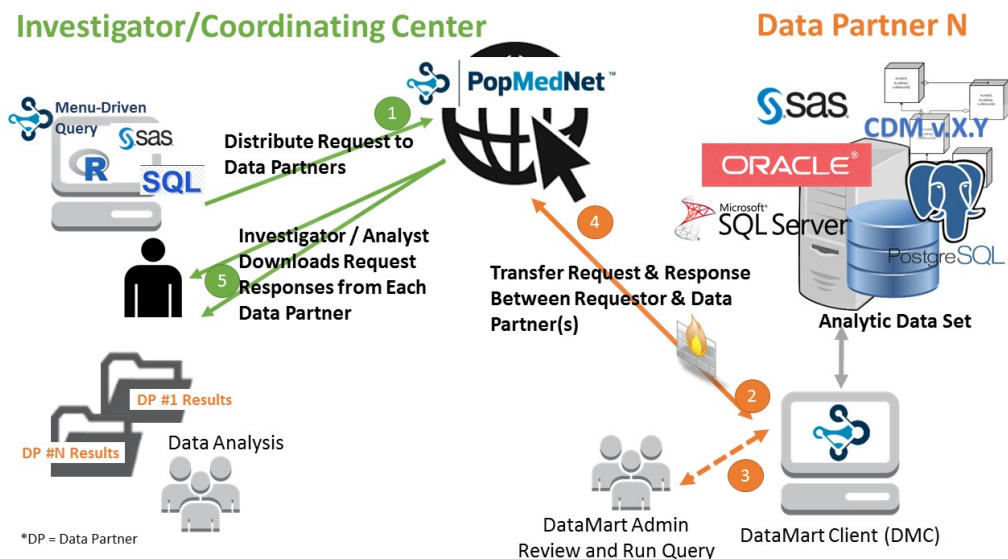


Figure 2: PopMedNet Workflow.

Comparison to i2b2 and SHRINE

i2b2 (Informatics for Integrating Biology & the Bedside) is a flexible query tool in use in many US hospitals. Each i2b2 instance comprises a collection of cells called a “hive”. This architecture supports extensibility via add-ons through additional cells.

SHRINE (Shared Health Research Informatics Network) is an i2b2 extension for supporting distributed queries. Each SHRINE instance consists of two hives, the client and the server. SHRINE instances communicate through the SPIN federated query language (McMurry et al., 2013).

Architecturally, PopMedNet™ uses a more tightly coupled hub-and-spoke distribution strategy while SHRINE uses a more loosely coupled, mesh strategy. While SHRINE requires immediate query response from the data sources in a single session, PopMedNet™ allows request-responses to be conducted over multiple sessions. From a user-interface perspective, i2b2/SHRINE uses a drag-and-drop model, while PopMedNet™ uses a menu-driven model.

A SHRINE-PopMedNet™ was prototyped to give SHRINE queries the ability to span sessions and conversely to give PopMedNet™ the ability to participate in a SHRINE network.

Acknowledgements

PopMedNet™ (“PopMedNet,” n.d.) was developed and has been extended as part of several contracts awarded by a range of federal, state, and industry stakeholders:

- FDA’s [Sentinel Initiative](#)
- PCORI’s National Patient-Centered Clinical Research Network ([PCORnet](#))

In addition, Harvard Pilgrim Health Care Institute serves as the coordinating center for several other initiatives that use PMN and its associated software ecosystem including:

- Biologics and Biosimilars Collective Intelligence Consortium (BBCIC)
- Office of the Assistant Secretary for Planning and Evaluation (ASPE) & Food and Drug Administration
- Privacy-protecting distributed analysis of biomedical big data project funded by the National Institute of Biomedical Imaging and Bioengineering of the National Institutes of Health

References

- Brown, J., Holmes, H., J., Shah, K., Hall, K., Lazarus, R., & Platt, R. (2010). Distributed health data networks: A practical and preferred approach to multi-institutional evaluations of comparative effectiveness, safety, and quality of care. *Medical Care*, 48. <https://doi.org/10.1097/MLR.0b013e3181d9919f>
- Davies, M., Erickson, K., Wyner, Z., Malenfant, J. M., Rosen, R., & Brown, J. (2016). Software-enabled distributed network governance: The PopMedNet experience. *The Journal of Electronic Health Data and Methods*. <https://doi.org/10.13063/2327-9214.1213>
- Malenfant, J., Davies, M., Swan, B., & Brown, J. S. (2015). PopMedNet: Collecting and using metadata in distributed research networks. *Journal of Patient-Centered Research and Reviews*. <https://doi.org/10.17294/2330-0698.1149>
- McMurry, A. J., Murphy, S. N., MacFadden, D., Griffin, W., Simons, W. W., Orcechia, J., Bickel, J., Wattanasin, N., Gilbert, C., Trevvett, P., Churchill, S., & Kohane, I. S. (2013). SHRINE: Enabling nationally scalable multi-site disease studies. *PLOS ONE*. <https://doi.org/10.1371/journal.pone.0055811>
- PopMedNet. (n.d.). In *PopMedNet*. Harvard Pilgrim Health Care Institute. <https://www.popmednet.org/>