

Effective Health Care Research Reports

Number 14

Distributed Ambulatory Research in Therapeutics Network (DARTNet): Summary Report

Wilson D. Pace, M.D.

David R. West, Ph.D.

Robert J. Valuck, Ph.D.

Maribel Cifuentes, R.N., B.S.N.

Elizabeth W. Staton, M.S.T.C.

Research from the Developing Evidence to Inform Decisions about Effectiveness (DEcIDE) Network



Agency for Healthcare Research and Quality
Advancing Excellence in Health Care • www.ahrq.gov

July 2009

The DEcIDE (Developing Evidence to Inform Decisions about Effectiveness) network is part of AHRQ's Effective Health Care program. It is a collaborative network of research centers that support the rapid development of new scientific information and analytic tools. The DEcIDE network assists health care providers, patients, and policymakers seeking unbiased information about the outcomes, clinical effectiveness, safety, and appropriateness of health care items and services, particularly prescription medications and medical devices.

This report is based on research conducted by the University of Colorado DEcIDE (Developing Evidence to Inform Decisions about Effectiveness) Center under contract to the Agency for Healthcare Research and Quality, Rockville, MD (Contract No. HHS29020050037I TO2). The AHRQ Task Order Officer for this project was Gurvaneet Randhawa, M.D., M.P.H.

The findings and conclusions in this document are those of the authors, who are responsible for its contents; the findings and conclusions do not necessarily represent the views of AHRQ. Therefore, no statement in this report should be construed as an official position of AHRQ or the U.S. Department of Health and Human Services.

None of the authors has a financial interest in any of the products discussed in this report.

Suggested citation:

Pace W, West D, Valuck R, et al. Distributed Ambulatory Research in Therapeutics Network (DARTNet): summary report. (Prepared by University of Colorado DEcIDE Center under Contract No. HHS29020050037I TO2.) Rockville, MD: Agency for Healthcare Research and Quality. July 2009. Available at: effectivehealthcare.ahrq.gov/reports/final.cfm.

Contents

Executive Summary	1
Chapter 1. Introduction	5
Background	5
Chapter 2. DARTNet Description.....	7
Overview of Federated Database Concept.....	8
DARTNet Processes for Aggregating and De-identifying Patient Information	10
Types of Data Linked Through DARTNet	11
Clinical Decision Support and Data Extraction	12
Data Mapping With Each EHR.....	15
Data Standardization.....	15
Data Export and Presentation to the Grid	16
Natural Language Processing	17
Directed Point-of-Care Data Collection.....	17
Summary of DARTNet Structure	19
Validation of DARTnet.....	20
Governance	21
Chapter 3. How a Question Is Asked and Answered.....	21
Chapter 4. Lessons Learned and Future Directions	24
Future Direction and Challenges	30
Future Research Questions	30
Immediate Next Steps.....	31
Strategic Issues.....	31
Blueprint for the Future	32
Supporting DARTNet as a Learning Community	33
Conclusion	33
Acknowledgments.....	34
References.....	35
Glossary	36
Appendixes:	
A. DARTNet Design Specifications	
B. Pilot Research Project: DARTNet Phase 1	

Author affiliations:

Wilson D. Pace, M.D.^{a,b}

David R. West, Ph.D.^a

Robert J. Valuck, Ph.D.^c

Maribel Cifuentes, R.N., B.S.N.^a

Elizabeth W. Staton, M.S.T.C.^a

^a University of Colorado Denver School of Medicine, Department of Family Medicine

^b University of Colorado Denver School of Pharmacy

^c American Academy of Family Physicians National Research Network

Abstract

The Distributed Ambulatory Research in Therapeutics Network (DARTNet) is a federated network of electronic health databases created in 2008. Its purpose is to facilitate quality improvement of primary healthcare and to efficiently compile clinically-enriched data for comparative effectiveness research. A federated network such as DARTNet links geographically and organizationally separate databases to allow a single query to pull information from multiple databases while maintaining the privacy and confidentiality of each database.

The DARTNet federated network begins within each member organization, where a database assembles patient-level information (such as vital signs, social history, family history, and physical examination findings) from electronic health records, pharmacy utilization databases, and billing systems. This aggregated clinical information is then standardized, de-identified, and linked securely through the Web to similar databases in other DARTNet member organizations. Once linked, a database query can be sent to all federated databases at once.

In addition to facilitating queries among the federated databases, the DARTNet system can prompt clinicians to obtain specific information during a patient encounter. This capability allows research teams to collect additional data beyond what would normally be recorded during clinical care. Thus, DARTNet research projects can include observational research and practical clinical trials. DARTNet is also designed to help support a learning community, where health care providers from DARTNet member organizations can learn from the best practices of high-performing member practices, which are identified by comparing quality indicators of clinical care provided across the network.

DARTNet will also advance observational comparative effectiveness research (OCER) methods, which have traditionally depended on data sets created for other purposes, such as administrative data sets created to process insurance claims. DARTNet enhances these methods by providing a way to account for important clinical information that is missing from claims databases.

DARTNet's capabilities were demonstrated by a retrospective cohort study that evaluated patterns of use, comparative effectiveness, and safety of oral diabetes medications for adults with type 2 diabetes. The study was conducted in two phases. Phase 1 used a commercially available, integrated medical claims database to examine the comparative effectiveness and safety of oral diabetes medications. Phase 2 used DARTNet data for the same purpose, and showed that DARTNet can identify comparable panels of diabetic patients receiving various combinations of oral diabetes medications similar to the traditional OCER study completed in phase 1. It also showed that DARTNet can collect clinically-relevant data such as body weight, height, self-reported alcohol intake, and self-reported hypoglycemic events, which were absent in the claims database. The initial results from DARTNet data analyses indicate that the comparative effectiveness and safety of oral hypoglycemics is similar to that observed in the Phase 1 study.

The next steps for DARTNet include expanding its technical capabilities to complete the analyses of phase 2 as well as scale up the size and diversity of DARTNet clinical entities and population. Further, DARTNet will refine the final organizational structure of the network, define the selection process for research and quality improvement projects, and develop DARTNet's capacity as a learning community. By successfully combining the concepts of point-of-care data collection with secondary analysis of electronic medical record data from large populations of patients, DARTNet holds great potential for becoming a valuable tool for comparative effectiveness research and improving the quality of care provided by its members.

Executive Summary

The Distributed Ambulatory Research in Therapeutics Network (DARTNet) is a federated network of electronic health data from eight organizations representing over 500 clinicians and over 400,000 patients. A federated network is a collection of geographically and organizationally separate databases that are treated as one entity and viewed through a single user interface. Federated networks link separate databases in such a way that a single query can run on the separate databases and return results while conforming to each organization's privacy and confidentiality standards. DARTNet was formed in response to a request for proposals from the Agency for Healthcare Research and Quality and is funded through the Agency's Developing Evidence to Inform Decisions about Effectiveness Network.

The core of DARTNet is a network of medical practices that use electronic health records (EHR). Data from each practice's EHR are standardized and stored in a local database. This local database, called the Clinical Data Repository (CDR), connects to many other databases, such as practice management databases, hospital databases, and pharmacy fill/refill databases, thus centralizing and standardizing data across disparate systems. The CDR prepares data for presentation to the Internet for distributed queries. The distributed query process is handled through an application of the Globus[®] Toolkit, an open source software toolkit used for building computing Grids. Grid computing allows functions such as complex queries to be passed to any number of local nodes without crossing an organization's firewall to physically access the data to be acted upon.

The purpose of this report is to summarize the purposes for which DARTNet was created, including its primary aim of conducting observational comparative effectiveness research (OCER) and secondary aim of conducting a full spectrum of practice-based research. This report comes at the end of the initial development phase of DARTNet. The report provides a broad overview of the technical components of DARTNet, the evaluation and certification checks conducted, the results of an initial pilot OCER project that studied oral hypoglycemic medications, the long-term business model options for DARTNet, and some lessons learned and conclusions from this initial phase of work. This report can be considered a summary of the first phase of DARTNet. Funding for an expansion of DARTNet, further development, and additional research is already secured and underway.

To understand why a network such as DARTNet is needed, it is helpful to understand that classic OCER uses primarily insurance claims data from large populations to compare outcomes for large cohorts of patients. While very powerful, this type of research is often hampered by the lack of basic clinical data. DARTNet provides this missing clinical information and helps to quickly improve knowledge about health outcomes. DARTNet was also designed to detect important sentinel events, identify rare adverse effects of treatment, and assess long-term outcomes of treatments. The network can explore the impact of new innovations that affect small populations (for which no single care delivery system would have a sufficient sample size).

The aims of the initial DARTNet project were to: (1) develop a federated network of 200+ primary care clinicians who use EHRs; (2) analytically demonstrate how existing large-scale data sets can be enhanced by patient-level data from the federated primary care network to inform and expand knowledge of effective and safe medical therapeutics;

and (3) demonstrate the ability to collect specific data from clinicians or their staff on a clinically defined set of patients to enrich the EHR data set and answer effectiveness and safety questions concerning medical therapeutics.

How DARTNet Works: To link the disparate EHRs and related health information databases, DARTNet uses advanced clinical decision support tools available from Clinical Integration Networks of America, Inc. (CINA), a small corporation that provides patient specific clinical decision support at the point of care. At each DARTNet member organization, the data in the EHR are captured in a relational dataset referred to as the Clinical Data Repository (CDR). The CDR is a standardized database of relevant clinical information, not an image of the entire EHR database. In the CDR, data elements are standardized across EHR products. Currently 5 EHR products are supported within DARTNet, but CINA has established interfaces with many of the major EHR products in the country. This standardized dataset includes patient identifiers. Data from the CDR are transferred to another database also located within each organization: the electronic Primary Care Network Gateway database or “Gateway” for short. The Gateway presents de-identified data for query access through a secure Grid enabled web-portal. The movement of data from the CDR to the Gateway database is based on the Continuity of Care Record, which is a core dataset of the most relevant administrative, demographic, and clinical information about a person’s health care.

A full set of patient data never leaves the clinical sites where it is stored; however, using a secure Grid enabled web-portal, the DARTNet research team can query the de-identified federated databases. DARTNet builds on several technologies to create a true distributed clinical data repository. Data linked through DARTNet include data from EHRs, laboratory, imaging, medication fills/refills, demographic and billing systems. While DARTNet is currently based in primary care practices, practice specialty is not an inherent limitation to the design model. DARTNet is also not dependent on any particular EHR brand for data access.

How DARTNet Was Evaluated: DARTNet was subjected to ongoing testing and validation processes to warrant the integrity of the system. Validation was focused on three primary areas: data integrity, software functionality, and system security. The data integrity validation included the process of capturing data from each organization’s EHR, standardizing the data, presenting the data for secure Internet based queries and transferring query results back to the research staff. This set of activities was validated through 31 different steps. Given that DARTNet uses clinically derived data, it was important to perform detailed testing of all data processing steps to be sure that any dataset used for research is a valid representation of the original clinical data. Software functionality testing included evaluating overall system performance, characteristics of four separate components of the system as well as 16 functions and processes. As DARTNet expands it will be important that all software components work as expected with a minimal amount of support, to create a sustainable system. Testing of the system security is designed to ensure that DARTNet does not represent a security threat to patient data or to patients. Seven steps were used to verify system security.

The DARTNet Pilot Research Study: A two-phase pilot study was conducted to demonstrate the capabilities of DARTNet. In Phase 1, a retrospective, claims-based study

was conducted to: (1) determine the extent to which commercially available, integrated medical claims databases could be used to answer key research questions related to comparative effectiveness and safety of oral diabetes medications for adults with Type 2 diabetes; and (2) identify areas where such databases are limited, and for which DARTNet (through access to EHR and/or point of care data) may be useful for augmenting and improving upon the results that can be obtained from claims-database studies. This phase examined a limited set of data elements for a large number of individuals. We used the Ingenix/ICHIS Impact database for Phase 1.

The findings of Phase 1 suggest that there are no substantial, clinically significant differences in the adjusted effectiveness of any of the common oral diabetic medication (ODM) monotherapies. Furthermore, the adjusted differences in initial monotherapy versus combination therapy are minimal and suggest there is little reason to consider starting patients on combination therapy regimens. But while all monotherapies appear to be equally effective, sulfonylureas appear to have several characteristics that would indicate they may not be ideal initial therapy. The findings of our safety analyses suggest that sulfonylureas may put patients at increased risk of hypoglycemia and liver injury when compared to other ODM agents used as initial therapy for type 2 diabetes.

Phase 2 was a replication of the same study using DARTNet data. We focused on a smaller number of individuals but examined a broader range of data elements. We looked at EHR data using both coded data and natural language processing, and also tested the point-of-care data collection prompts.

Phase 2 findings show that DARTNet was able to identify similarly sized panels of diabetic subjects and subjects receiving various ODM to enable analyses of similar power to the claims based studies performed in Phase 1. The time to first regimen change for monotherapy and combination therapy groups in the DARTNet Diabetes Replication Cohort was generally similar to those seen in phase 1. The crude effectiveness outcomes showed that baseline and crude changes in hemoglobin A1C values were very similar in the DARTNet Diabetes Replication Cohort as in the claims cohort. Rates of hypoglycemia, liver injury, and liver failure were rare, and did not differ substantially across major ODM drugs or classes. No definitive conclusions can be drawn regarding the comparative safety outcomes in the Phase 2 replication cohort.

Lessons Learned to Date: Several lessons learned in this initial phase highlight the importance of collaborating with trusted colleagues. It is extremely important to negotiate the division of labor and monetary incentives for contractors early on. The DARTNet leadership is likely to further expand the detail in our performance-based contracts to clearly define key deliverables, delivery dates, and corresponding reimbursement in any future collaboration.

Among the most important lessons is a greater realization of some of the problems that can potentially arise with the intentional or unintentional misuse of this powerful new tool. Anticipated new challenges revolve primarily around the ability to potentially break or overload the DARTNet system, unintentionally compromise privacy or security resulting in potentially serious consequences, or degrade system functionality because of a lack of understanding of the underlying system.

In terms of the actual data being studied, experience has shown that the practice-based EHRs include incomplete, un-coded, highly variable data. It can be quite

challenging to aggregate data from separate practices, even those using the same type of EHR. There were wide variations in the way data from a given encounter were correlated, and it was nearly impossible to view an “episode of care” spanning multiple encounters. A further data quality issue was the capability of providers to enter selected data while leaving important data un-addressed within an encounter. A final area of data quality involves the lack of editing within systems for reasonableness of data.

This initial phase of work has shown the DARTNet can expect to work with primary care practices that have varied (and in some cases, quite minimal) IT capacity.

Next Steps: Critical to the process of realizing DARTNet’s potential is defining the appropriate next steps for the network. These next steps will take into account the need to expand technical capabilities of DARTNet, expand the patient base, define the final corporate/organizational structure of DARTNet, and define the next research questions. DARTNet may also need to expand its technical capabilities and data access. Important next steps include continued efforts to build the DARTNet research portfolio, secure additional funding, and ensure independence of the network. AHRQ has provided DARTNet with additional funding under a second task order to further expand the network to include a larger number of clinics, including specialty and general pediatric practices, to conduct an observational comparative effectiveness and safety study of different therapies for major depression, and to further test and develop DARTNet’s ability to collect point-of-care data. In addition, DARTNet will also work to create a more robust governance and management structure and an active learning community. We will search for “best practices” within the network and work with our members for how to disseminate these findings. In future grants we are including facilitators to help with this process. We will be exploring the use of social networking concepts and understand that we will likely need to identify sub-groups of like practices that we help organize into smaller learning groups. As we work with practices to improve care and add standardized data collection techniques to assist with these improvements these data should also improve our ability to conduct meaningful OCER.

Conclusions: DARTNet is a proven prototype that is capable of bi-directional electronic communication with EHR enabled physician practices, and holds great potential for becoming a valued tool for OCER, informing policymakers faced with making pragmatic health care decisions, and providing the information technology backbone for bringing populations of primary care practices together to design, track, and evaluate quality improvement initiatives. To evolve as a fully functional system, DARTNet requires further technical refinement and a significant expansion of the patient base. DARTNet has contributed to the understanding of patterns of use and effectiveness and safety of oral diabetes medications, and DARTNet investigators have identified important clinical data elements that will be useful in enhancing the state of the art in OCER, thus improving the understanding of the effectiveness of various clinical interventions – and improving the information available to policymakers in healthcare programs such as Medicare, Medicaid, and SCHIP.

Chapter 1. Introduction

The Distributed Ambulatory Research in Therapeutics Network (DARTNet) is a federated network of electronic health data from eight organizations representing over 500 clinicians and over 400,000 patients. A federated network links geographically and organizationally separate databases in such a way that a single query can run on the separate databases and return results while conforming to each organization's privacy and confidentiality standards.

The core of DARTNet is a network of medical practices that use electronic health records (EHR) and that have agreed to collaborate in research and quality improvement activities. Using DARTNet infrastructure, data from each practice's EHR is de-identified, standardized, and stored, then linked with other pertinent health data, such as practice management information, inpatient data, and drug fulfillment (pharmacy fills/refills) information. These data are then presented for distributed querying to answer questions concerning the safety and effectiveness of medications and medical devices. DARTNet was formed in response to a request for proposals from the Agency for Healthcare Research and Quality.

The purpose of this report is to summarize the prototype phase of DARTNet including the various reasons for its creation, its primary aim of conducting observational comparative effectiveness research (OCER) and secondary aim of conducting a full spectrum of practice-based research involving quality improvement, safety, effectiveness and descriptive work. This document will also provide an overview of the technical components of the DARTNet system, the required evaluation and certification checks conducted prior to using the system for comparative effectiveness research, the results of an initial pilot OCER project that studied oral hypoglycemic medications, the proposed aims and business model options to be explored in the 2nd phase of DARTNet, and some lessons and conclusions from this initial phase of work.

Background

Although a large number of papers are published each year describing individual studies on specific conditions, clinical situations, medications, and devices, many important questions on how to provide the safest, most effective health care still remain. Despite evidence from so many studies, there remains extensive and unexplained variability in what happens in hospitals and medical offices.¹ Furthermore, the clinical research enterprise in the United States is not producing adequate effectiveness knowledge to meet the needs of patients, physicians, payers, purchasers, health care administrators, and public health policymakers.¹ Two approaches to answering clinical questions may provide some of the missing information: practical clinical trials conducted through practice-based research networks and observational comparative effectiveness research.

Practice-based research networks provide a mechanism for examining everyday challenges of clinical care and for examining the effectiveness of particular features of clinical practice. A practice-based research network (PBRN) is a group of ambulatory clinicians or practices devoted principally to the care of patients and who work together to investigate questions related to community-based practice and to improve the quality

Effective Health Care Research Report Number 14

of care.² PBRN efforts link relevant clinical questions with rigorous research methods in community settings to produce scientific information that is externally valid.² Increasingly, PBRNs are supporting quality improvement activities, and are thus evolving from clinical laboratories into collaborative learning communities that identify, disseminate, and integrate new knowledge to improve care processes and patient outcomes.³ A hallmark of PBRN research is the ability to collect data on patients, conditions, and care as it occurs—at the point of care.

Practical clinical trials, which address questions about the risks, benefits, and costs of an intervention as they would occur in routine clinical practice, may provide useful and adequate information regarding clinical effectiveness.¹ Practical clinical trials compare clinically relevant alternatives, enroll a diverse study population, recruit from a variety of practice settings, and measure a broad range of relevant health outcomes. Practical clinical trials require substantial resources and a durable operational infrastructure that can be difficult to maintain independently. Practice-based research networks provide the central resources and operational infrastructure to conduct increasingly complex practical clinical trials and have continued to grow in number over the last three decades. Linking OCER research to an electronic research network provides alternative opportunities for infrastructure support for both activities.

Through its DEcIDE (Developing Evidence to Inform Decisions about Effectiveness) Network, AHRQ is providing funding to conduct accelerated practical studies about the outcomes, comparative clinical effectiveness, safety, and appropriateness of health care items and services. Traditional OCER uses available data, primarily insurance claims data (which may contain lab results) from large populations to create various cohorts of patients and then compare outcomes of each cohort. Generally, sophisticated statistical matching algorithms (such as propensity matching) are applied to members of the various cohorts to attempt to account for underlying clinical differences among the groups. While very powerful, this type of study is often hampered by the lack of basic clinical data to supplement the claims data. For instance, it is well known that depression is markedly under coded in claims data even when actively treated, thus confounding all the outcomes of any claims data analysis. This is not unique to depression. For instance, the validity of a diabetes diagnosis in claims data from one practice to another varies from over 98% valid to under 75% valid when compared to chart audit validation. The ability to improve the quality of data in claims data is limited as there is no interaction between those who “collect” the data – providers of care – and those who analyze the data. DARTNet provides an ongoing interaction between the researchers and the providers. Thus, we believe DARTNet offers the ability to develop enhanced data sets that far exceed the quality and content of claims data. A claims database lacks important clinical data such as alcohol use that can also account for liver toxicity. Medication side effects or effectiveness may be associated with body mass index (BMI) and this data is absent in claims databases. Table 1 presents some of the clinical variables present in DARTNet data but not in claims data.

Table 1. Variables available via DARTNet but not through claims

- Medication allergies
- Reason for appointment
- Family history
- Findings (blood pressure, weight, height, etc.)
- Social history (alcohol and tobacco use, etc.)
- Laboratory orders and results
- Prescribed medications (name, dose, frequency, number of refills authorized, etc.)
- Past medical history
- Date of onset of disease
- Referrals
- Provider-level data
- Practice-level data
- Data collected/prompted for collection at point of care

As part of the AHRQ DecIDE activities, the Colorado DEcIDE network (CO-DEcIDE) was funded to develop a prototype and pilot test a cooperative research network for conducting studies on therapeutic safety and effectiveness using electronic health information. The primary goals of the project are to improve public knowledge about health outcomes more quickly than traditional research approaches can, and to take advantage of the power of networks to increase the ability to detect important sentinel events such as rare adverse effects of treatment and assess long-term outcomes of treatments. The network spans multiple health care organizations to explore the impact of new innovations that affect small populations (for which no single care delivery system would have a sufficient sample size). This approach enhances and sustains both the OCER model and the PBRN model of practice improvement and practical clinical research. This approach may result in greater research effectiveness, quicker identification of safety problems and improved quality of care for patients. The network developed through this funding is the Distributed Ambulatory Research in Therapeutics Network (DARTNet), a network that can investigate questions using electronic health record data across multiple organizations and that, like a PBRN, can collect new data on individual patients at the point of care to answer research questions. DARTNet is a public/private partnership between the University of Colorado Denver, Department of Family Medicine and School of Pharmacy, The American Academy of Family Physicians National Research Network, The Robert Graham Policy Center, The University of Minnesota Center for Excellence in Primary Care, and Clinical Integration Networks of America.

The aims of the initial DARTNet project were to: (1) develop a federated network of 200+ primary care clinicians, all using electronic health records (EHR); (2) analytically demonstrate how existing large-scale data sets can be enhanced by patient-level EHR data to inform and expand knowledge of effective and safe medical therapeutics; and (3) demonstrate the ability to collect specific data from clinicians or their staff on a clinically defined set of patients to enrich the EHR data set and answer effectiveness and safety questions concerning medical therapeutics.

Chapter 2. DARTNet Description

The DARTNet prototype currently captures, codifies and standardizes over 150 unique data elements per patient for 48 months of time or more. In this first phase DARTNet was used to explore how currently available EHR data can be used to

supplement data from a large administrative dataset in order to answer questions concerning the safety and effectiveness of diabetes medications. DARTNet also tested its point-of-care data collection techniques to obtain clinical data at the patient level to supplement EHR data. The following sections provide a broad overview of the DARTNet concept and technological structure. The full report describing DARTNet’s design and specifications can be found in Appendix A.

Overview of Federated Database Concept

A federated (unified) database is a collection of databases that are treated as one entity and viewed through a single user interface. A federated database structure facilitates sharing and interchanges of data among autonomous databases, such as EHRs located within different organizations. A federated database architecture unites independent database systems in order to share and exchange information. The single interface provided by a federated database system allows a user to retrieve data from multiple geographically decentralized and heterogeneous databases with a single query.

In the case of DARTNet, the federation consists of EHR systems from individual organizations that are members of DARTNet. Each organization in the network controls its interactions with the federation (Table 2). The federated architecture provides mechanisms for sharing data, for sharing transactions, for combining information from several components of the system, and for coordinating activities among autonomous components.

Table 2. DARTNet organizations and corresponding EHRs

DARTNet organizations	EHR product
Medical Clinic of North Texas	NextGen [®]
WellMed Medical Group	SmartClinic [®]
Tiena Health	Allscripts Professional [®]
Wilmington Health Associates	Allscripts Professional [®]
University of Colorado	Allscripts Enterprise [®]
University of Minnesota	Allscripts Enterprise [®]
Cranford Family Medicine	e-MDs [®]
Family Health Center of Joplin	e-MDs [®]

EHR= electronic health record

The federated database architecture of DARTNet was created using several technologies, which are briefly explained here.

The DARTNet architecture is based on Grid computing, a method of creating very large-scale distributed networks. Distributed networks allow a central site to simultaneously query data stored locally for each member of the Grid. Grid computing facilitates flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources. Such dynamic situations present unique

challenges related to authentication, authorization, resource access, and resource discovery. These types of challenges are addressed by Grid technologies.

Grid computing is distinct from Internet access (which is primarily a communication tool), wide area networks, or virtual private networks, which allow access rights behind a firewall. Grid computing allows functions such as complex queries to be passed to any number of local nodes without crossing an organization's firewall to physically access the data to be acted upon. The query can be executed locally and the output can either be stored on a local computer or, if allowed, returned to the central location. Since the Grid utilizes parallel processing, queries that may take hours to run against a central database can often be completed in minutes across a Grid system.

The specific challenge that underlies the Grid concept is coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations. This sharing is not primarily file exchange but rather direct access to computers, software, data, and other resources. This sharing is highly controlled, with clear definitions of just what is shared, who is allowed to share, and the conditions under which sharing occurs.

Grid architecture uses protocols to define the basic mechanisms by which organizational users and resources negotiate, establish, manage, and utilize sharing relationships. A standards-based open architecture facilitates future growth and code sharing and is interoperable and portable. Application programming interfaces and software development kits are available for the Grid. An application programming interface is a set of procedures, methods, and protocols that support requests made by computer programs, and a software development kit is a set of development tools that allows software engineers to create applications. Together, these technologies and architecture constitute "middleware" (that is, the services needed to support a common set of applications in a distributed network environment). The Globus[®] Toolkit has emerged as the dominant middleware for Grid deployments worldwide.

The open source Globus[®] Toolkit lets people share computing power, databases, and other tools securely online across corporate, institutional, and geographic boundaries without sacrificing local autonomy. The Toolkit includes software services and libraries for resource monitoring, discovery, and management, plus security and file management, and it is a substrate on which leading IT companies are building significant commercial Grid products.

The Toolkit includes software for security, information infrastructure, resource management, data management, communication, fault detection, and portability. It is packaged as a set of components that can be used either independently or together to develop applications. The Globus[®] Toolkit has grown through an international open-source strategy similar to the Linux operating system.

The DARTNet Grid application was built with the Globus[®] Toolkit. The primary interface for accessing the Grid system is called the electronic Primary Care Research Network (ePCRNet) Portal. The ePCRNet portal was initially developed by the University of Minnesota with funding from the National Institutes of Health in collaboration with the University of Birmingham, UK. The ePCRNet portal provides the architecture that accesses the Gateway databases for all users, as well as the query capabilities and the security systems. The Gateway database presents de-identified patient data for query access. DARTNet uses the connectivity and distributed query capabilities of the ePCRNet

Portal, which are handled through a specifically designed application of the Globus Toolkit.

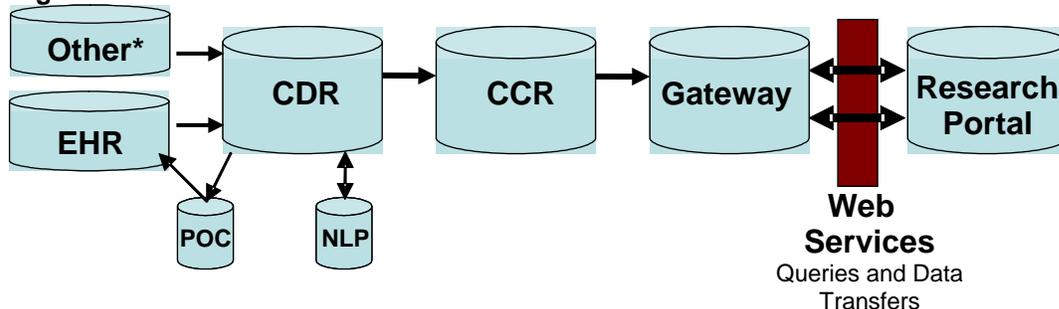
DARTNet Processes for Aggregating and De-identifying Patient Information

DARTNet makes use of advanced clinical decision support tools available from Clinical Integration Networks of America, Inc. (CINA). CINA is a small corporation dedicated to providing patient specific clinical decision support at the point of care independent of a particular EHR. First, at each member organization, the data in the EHR are captured in a relational data set referred to as the Clinical Data Repository (CDR). The CDR is a proprietary product of CINA. In the CDR, data elements are standardized across EHR products. DARTNet has proven interfaces with twelve of the major ambulatory EHRs in the United States. We have standardized over 150 selected EHR data elements. This standardized data set, which remains on the local servers of each organization, includes patient identifiers. Second, data from the CDR are transferred to another database also located within each organization: the ePCRN Gateway database or “Gateway” for short. The Gateway presents de-identified data for query access through a secure Grid enabled web-portal.

The movement of data from the CDR to the ePCRN Gateway database is based on the Continuity of Care Record (CCR), which was standardized by ASTM International—a voluntary standards development organization. The CCR is a core dataset of the most relevant administrative, demographic, and clinical information about a person’s health care. A CCR provides a means for aggregating and moving all of the pertinent data about a patient.

A full set of patient data never leaves the clinical sites where it is stored; however, using the secure Grid enabled web-portal, the DARTNet research team can query the de-identified federated databases. The DARTNet system is currently based in primary care practices, but practice specialty is not an inherent limitation to the design model. DARTNet is not dependent on any particular EHR brand for data access. Figure 1 below summarizes the relationships between data sources and data access points.

Figure 1. Relationship of data sources and the DARTNet architecture within a single organization



*Other data sources include billing data, hospital data, medication fill/refill data, and other third party databases.

ABBREVIATIONS:

EHR = electronic health record
CDR = Clinical Data Repository
CCR = Continuity of Care Record
POC = point-of-care data
NLP = natural language processing data

Types of Data Linked Through DARTNet

Data interfaces for DARTNet organizations are described below. All interfaces operate for clinical purposes. These interfaces are not maintained by the DARTNet infrastructure per se. They are described here for the reader to understand how clinical data are available for use by DARTNet.

Electronic Health Record (EHR) Data

DARTNet has elected to work only with EHRs that include coded problem lists, electronic prescribing, and laboratory interfaces. The EHR system must also allow read-only access to a data extraction/standardization system. Virtually all EHRs that meet these minimum requirements can be supported. EHRs that are known to be compatible with the current data extraction system include Allscripts Professional[®], Allscripts Enterprise[®], eMDs Chart[®], GE Centricity[®], Meditech[®], Misys EMR[®], NextGen[®], Practice Partner[®], Medent[®], NextGen[®], SmartClinic[®], and SOAPware[®]. DARTNet organizations must commit to using their EHRs in a way that will support an advanced clinical decision support system, meaning that the organization uses limited locations for data elements and consistent terminology throughout the EHR. Where use is highly variable by member clinicians, the organization must develop a plan to improve data integrity to support an advanced clinical decision support system. Note that the minimum standards are couched in clinical decision support terms to help solve clinical problems facing DARTNet clinicians with the expectation that this will improve clinical care, while also providing higher quality data for research.

Laboratory, Imaging, and Medication Fulfillment Data

An electronic laboratory interface must be in place for the primary laboratory used by each organization, and preferably for all laboratories used by the organization. Electronic interfaces for an organization's secondary labs can be created to store data directly in the CDR if an EHR interface is not available. These labs would then support the clinical decision support/data extraction process.

Electronic interfaces with the imaging centers used by the organization are preferred but not required. Even electronically transmitted imaging data arrives as a text file and therefore requires processing for data extraction.

SureScripts-RxHub is a partner in this project and all DARTNet sites were encouraged to install SureScripts-RxHub capabilities as the organization's EHR and region allow. These activities are part of the EHR installation and not within the control of the DARTNet infrastructure development. Once SureScripts-RxHub capabilities are installed, medication fulfillment data can flow either into the organization's EHR (preferred method) or into the CDR until the EHR can support these data. In either case, the medication fulfillment data are available for clinical decision support, quality improvement, and research purposes.

Hospital Inpatient Data

Hospital inpatient data interfaces are not included in most ambulatory EHRs. In EHRs where inpatient data are included, the enormous volume of inpatient data has created difficulties in finding useful ways to sort inpatient and outpatient data for efficient review. Selected hospital inpatient data elements can be useful, but the extent of the data required for outpatient clinical care is markedly less than the needs during many inpatient admissions. Therefore, we continue to explore the need and usefulness of specific data interfaces for hospital data that would feed into the CDR and be accessed for quality improvement and research. In this first phase we have focused our data collection to inform us of severe adverse events that originated in the ambulatory arena rather than developing a robust inpatient related distributed database. Such interfaces require development with each institution in terms of data elements, frequency of data transfers and local data storage.

Demographic and Billing Data

The CDR in use by DARTNet is able to link to most billing and accounts receivable systems. In general, the data contained in these systems are also included in the EHR. Therefore, in this first phase, interfaces to these databases were not required of DARTNet participants, though some already had these interfaces in place. DARTNet will be moving rapidly to put these interfaces in place in all organizations to assist with visit-level interpretation of clinical data in its future work.

Clinical Decision Support and Data Extraction

Practice-level implementation of the CINA clinical decision support system provides an avenue to communicate with clinicians at the point of care. DARTNet capitalizes on the clinical decision support system to facilitate bi-directional communication with clinicians to collect EHR data from their practices, and to provide local clinical decision support services for clinicians to use for their own quality improvement initiatives. Communication with clinicians at the point of care can be tailored for each patient visit based on analysis of many different data elements from various sources, including the EHR, the billing system, direct patient data entry, drug fulfillment data, and laboratory data that is not linked to the EHR. Thus, the CINA software is a critical middle layer for DARTNet, collecting, and standardizing data and populating the federated database, as well as providing a method through which DARTNet can manage point-of-care data collection.

Using an Open Database Connectivity connection, CINA connects at the data level to ambulatory EHRs. CINA already has proven interfaces with many of the major ambulatory EHR products in the country as noted above. For this project the primary data sources were each organization's EHR, augmented where necessary by billing data and medication fulfillment data over time.

CINA completes data extraction and connectivity through a software package called the CINA Protocol Engine (PE). This software system runs locally on each practice's CINA server. Through connectivity to each organization's EHR and ancillary databases the PE fills the CINA CDR. After the CDR is filled or each time new data are added, the PE performs data standardization activities and runs a series of clinical

decision support functions (including research specific algorithms) against each patient for whom new information has been added to the CDR since the previous update. All outputs from the clinical decision support function (such as the need for a particular lab test or the need for additional data collection for a research project) are stored in the CDR and reported to clinicians at the time of the next patient visit or can be exported for batch usage (such as recall letters) and reporting activities. If there is a change in one of the protocols in the PE then the system will run the new protocol against all eligible patients updating their clinical decision support results. On a timed, automatically triggered basis—typically early each morning—the CINA system within each organization checks with the central CINA server to see if any protocol changes have been made to the organization’s PE. If so, the local PE is updated and the new or changed protocols run against the CDR. Thus, allowing DARTNet to rapidly implement new clinical decision support and research functions.

The CINA CDR is a normalized and standardized database of relevant clinical information (not an image of the entire EHR database). The CINA CDR can be deployed in a number of different compliant databases; DARTNet uses a Microsoft Sequel deployment. The CDR stores the raw data format for each data element (in some cases data elements may be represented in the EHR in many different ways) along with all outputs related to clinical decision support protocols. This allows us to recode any data element in the CDR if we discover a need to present the information in a more granular fashion. For instance, we are coding all Hemoglobin A1c values with one SNOMED CT code, though some tests are run in a central laboratory and some are performed using an analyzer at the point of care. If in the future it becomes important to tell the difference between these two circumstances we can look back into the CDR data and recode for each.

The point-of-care output of the clinical decision support system is a report that can be displayed on a web page, embedded within an EHR, or printed for use. Figure 2 shows a sample report. All DARTNet practices have found that printing the output is the best approach since it allows the information to be acted upon by the front desk staff, the patient, a nurse and the clinician. Paper is an easy and proven way to move the information between this diverse set of individuals during an office encounter. Also, several practices use the system to highlight online patient education and support resources and services based on diagnoses or other clinical data. In these practices the paper form is given to patients at the end of a visit for educational and personal clinical information tracking purposes.

Figure 2. Sample point of care report from the CINA protocol engine
 University of Colorado Health Science Center

Patient Recommendation Report

DOB: _____ Age: 78 Sex: F Seen By: _____
 Report Date: 10/21/2008 PCP: _____

Appointment Date: 10/22/2008 11:00:00 AM

<p>Active Diagnoses</p> <p>diabetic macular edema of the right eye (250.00) Diabetic Peripheral Neuropathy (250.60) Diabetic Peripheral Neuropathy Type I (250.50) Diabetic Retinopathy (250.50) Type II Diabetes Mellitus (250.00) Atherosclerosis Of The Femoral Artery (440.9) Hypertension (401.9) Hyperlipidemia (272.4) Osteoporosis (733.00) Abdominal Pain (789.00) Adhesive Capsulitis Of Shoulder (726.0) Anal Fissure (565.0) Anemia (285.9) Atypical Mycobacterial Infection (031.9)</p> <p>MORE</p>	<p>Risk Factors</p> <p>Considered CHD 10Yr Risk > 20% Breast CA (Age 40-80) Colon CA (age 50-80) Influenza (age > 50 or Risk Dx) Pneumonia (Age > 64 OR Risk Dx)</p>																																							
<p>Active Meds</p> <table border="0"> <tr><td>NovoLog Flexpen 100 U</td><td>0</td><td>03/27/07</td></tr> <tr><td>Lantus 100 UNIT/ML</td><td>54 On</td><td>01/07/04</td></tr> <tr><td>Enalapril Maleate 10 MG</td><td>1 Twice</td><td>01/04/08</td></tr> <tr><td>Hydrochlorothiazide 25</td><td>0.5 Onc</td><td>09/22/06</td></tr> <tr><td>Aspirin 81 MG</td><td>1 Once</td><td>09/12/07</td></tr> <tr><td>Advair Diskus 500-50 MC</td><td>1 Twice</td><td>05/14/04</td></tr> <tr><td>Albuterol 90 MCG/ACT</td><td>0</td><td>01/07/04</td></tr> <tr><td>Azithromycin 250 MG</td><td>0</td><td>09/17/08</td></tr> <tr><td>Bisacodyl EC 5 MG</td><td>1</td><td>09/26/07</td></tr> <tr><td>Fluarix</td><td>0</td><td>10/10/08</td></tr> <tr><td>Fluticasone Propionate 5</td><td>0 Once</td><td>09/17/08</td></tr> <tr><td>Insulin Syringe 31G X 5/1</td><td>0</td><td>04/22/08</td></tr> <tr><td>Novofline 31 31G X 6 MM</td><td>0</td><td>07/31/08</td></tr> </table> <p>MORE</p>	NovoLog Flexpen 100 U	0	03/27/07	Lantus 100 UNIT/ML	54 On	01/07/04	Enalapril Maleate 10 MG	1 Twice	01/04/08	Hydrochlorothiazide 25	0.5 Onc	09/22/06	Aspirin 81 MG	1 Once	09/12/07	Advair Diskus 500-50 MC	1 Twice	05/14/04	Albuterol 90 MCG/ACT	0	01/07/04	Azithromycin 250 MG	0	09/17/08	Bisacodyl EC 5 MG	1	09/26/07	Fluarix	0	10/10/08	Fluticasone Propionate 5	0 Once	09/17/08	Insulin Syringe 31G X 5/1	0	04/22/08	Novofline 31 31G X 6 MM	0	07/31/08	<p>Goals</p> <p>Goal not met: BMI > 30 Goal not met: LDL >70 Goal met: BP <130/80 Goal met: A1c <= 7.0% Nonsmoker</p>
NovoLog Flexpen 100 U	0	03/27/07																																						
Lantus 100 UNIT/ML	54 On	01/07/04																																						
Enalapril Maleate 10 MG	1 Twice	01/04/08																																						
Hydrochlorothiazide 25	0.5 Onc	09/22/06																																						
Aspirin 81 MG	1 Once	09/12/07																																						
Advair Diskus 500-50 MC	1 Twice	05/14/04																																						
Albuterol 90 MCG/ACT	0	01/07/04																																						
Azithromycin 250 MG	0	09/17/08																																						
Bisacodyl EC 5 MG	1	09/26/07																																						
Fluarix	0	10/10/08																																						
Fluticasone Propionate 5	0 Once	09/17/08																																						
Insulin Syringe 31G X 5/1	0	04/22/08																																						
Novofline 31 31G X 6 MM	0	07/31/08																																						
<p>Labs</p> <table border="0"> <tr><td>Trig</td><td>102 mg/dl</td><td>3/13/08</td></tr> <tr><td>Chol</td><td>173 mg/dl</td><td>3/13/08</td></tr> <tr><td>LDL</td><td>108 mg/dl</td><td>3/13/08</td></tr> <tr><td>HDL</td><td>47 mg/dl</td><td>3/13/08</td></tr> <tr><td>Gluc, Fasting</td><td>87 mg/dl</td><td>9/12/08</td></tr> <tr><td>Gluc, Random</td><td>126 mg/dl</td><td>7/25/08</td></tr> <tr><td>HbA1c</td><td>6.6 %</td><td>9/12/08</td></tr> <tr><td>MicroAlb/Cr</td><td>6 mg/g</td><td>7/16/07</td></tr> </table>	Trig	102 mg/dl	3/13/08	Chol	173 mg/dl	3/13/08	LDL	108 mg/dl	3/13/08	HDL	47 mg/dl	3/13/08	Gluc, Fasting	87 mg/dl	9/12/08	Gluc, Random	126 mg/dl	7/25/08	HbA1c	6.6 %	9/12/08	MicroAlb/Cr	6 mg/g	7/16/07	<p>Action Items</p> <table border="0"> <tr><td>___ Document Advanced Directive status</td><td>PREV</td></tr> <tr><td>DOC: Add Obesity Dx / Plan to Problem List</td><td>DM</td></tr> <tr><td>DOC: Document or perform Diabetic Foot Exam</td><td>DM</td></tr> <tr><td>MED: Start Lipid lowering med* due to LDL goal not met</td><td>CAD</td></tr> <tr><td>LAB: Order Lipid Panel (q 6 mos for CHD risk > 20%).</td><td>CAD</td></tr> <tr><td>LAB: Order Urine Microalbumin (yearly) for DM</td><td>DM</td></tr> </table>	___ Document Advanced Directive status	PREV	DOC: Add Obesity Dx / Plan to Problem List	DM	DOC: Document or perform Diabetic Foot Exam	DM	MED: Start Lipid lowering med* due to LDL goal not met	CAD	LAB: Order Lipid Panel (q 6 mos for CHD risk > 20%).	CAD	LAB: Order Urine Microalbumin (yearly) for DM	DM			
Trig	102 mg/dl	3/13/08																																						
Chol	173 mg/dl	3/13/08																																						
LDL	108 mg/dl	3/13/08																																						
HDL	47 mg/dl	3/13/08																																						
Gluc, Fasting	87 mg/dl	9/12/08																																						
Gluc, Random	126 mg/dl	7/25/08																																						
HbA1c	6.6 %	9/12/08																																						
MicroAlb/Cr	6 mg/g	7/16/07																																						
___ Document Advanced Directive status	PREV																																							
DOC: Add Obesity Dx / Plan to Problem List	DM																																							
DOC: Document or perform Diabetic Foot Exam	DM																																							
MED: Start Lipid lowering med* due to LDL goal not met	CAD																																							
LAB: Order Lipid Panel (q 6 mos for CHD risk > 20%).	CAD																																							
LAB: Order Urine Microalbumin (yearly) for DM	DM																																							
<p>Measures / Calculations</p> <table border="0"> <tr><td>BP</td><td>124/70</td><td>9/29/08</td></tr> <tr><td></td><td>135/70</td><td>9/17/08</td></tr> <tr><td>CHD Risk</td><td>14%</td><td></td></tr> <tr><td>BMI (Wt)</td><td>34.1 (183.7lb)</td><td>9/29/08</td></tr> <tr><td> Ideal Wt.</td><td>135-156</td><td></td></tr> <tr><td>Est. CrCl</td><td>67.91</td><td>7/25/08</td></tr> </table>	BP	124/70	9/29/08		135/70	9/17/08	CHD Risk	14%		BMI (Wt)	34.1 (183.7lb)	9/29/08	Ideal Wt.	135-156		Est. CrCl	67.91	7/25/08	<p>Insurance: MEDICARE A & B</p>																					
BP	124/70	9/29/08																																						
	135/70	9/17/08																																						
CHD Risk	14%																																							
BMI (Wt)	34.1 (183.7lb)	9/29/08																																						
Ideal Wt.	135-156																																							
Est. CrCl	67.91	7/25/08																																						
<p>Diagnostic Testing</p> <table border="0"> <tr><td>Bone Density</td><td>11/30/06</td></tr> <tr><td>Colonoscopy</td><td>3/04/08</td></tr> <tr><td>Mammogram</td><td>1/24/07</td></tr> <tr><td>PAP</td><td></td></tr> </table>	Bone Density	11/30/06	Colonoscopy	3/04/08	Mammogram	1/24/07	PAP		<p>Routine Visits: Next Visit: Last Visit: 10/17/2008</p>																															
Bone Density	11/30/06																																							
Colonoscopy	3/04/08																																							
Mammogram	1/24/07																																							
PAP																																								
<p>Vaccine</p> <table border="0"> <tr><td>Tetanus</td><td>10/03/07</td></tr> <tr><td>Tdap</td><td></td></tr> <tr><td>Pneumoccal</td><td>10/17/06</td></tr> <tr><td>Flu</td><td>10/17/08</td></tr> </table>	Tetanus	10/03/07	Tdap		Pneumoccal	10/17/06	Flu	10/17/08	<p>Comp. Exam Visits: Next Visit: Last Visit: 05/03/2006</p>																															
Tetanus	10/03/07																																							
Tdap																																								
Pneumoccal	10/17/06																																							
Flu	10/17/08																																							
	<p>Next Appt. Date:</p> <table border="0"> <tr><td>3 mos</td><td>CAD</td></tr> <tr><td>3-6 mos</td><td>DM</td></tr> <tr><td>6 mos</td><td>HTN</td></tr> </table>	3 mos	CAD	3-6 mos	DM	6 mos	HTN																																	
3 mos	CAD																																							
3-6 mos	DM																																							
6 mos	HTN																																							

*--Unless contraindicated

Data Mapping With Each EHR

Data elements identified as applicable to the clinical decision support process or for current or future studies are mapped within each EHR and standardized upon being imported into the onsite CDR. Such mapping is needed because not every EHR, or even every installation of a given EHR, stores specific data elements in the same location. For example, some EHRs create a new table in the EHR database each time a new template for data capture is created. Thus, in one practice that allows each clinician to develop his/her own templates, the data required for documenting a diabetes foot exam were spread across several tables that only exist for that practice.

The labor-intensive process of mapping has included identifying the variations on data content and storage locations occurring within each EHR and deciding how to apply standardized nomenclature to each data variant. The EHR mapping process (which uses the *CINA Mapper*[®]) utilizes pattern matching to locate, verify and translate both codes and text into the standardized nomenclature maintained within the CDR. Whether or not data are stored in constrained fields, the *CINA Mapper*[®] allows likely matches to be viewed by a clinician who then determines whether the “match” is congruent with the concept desired. Once congruence is established, the protocol engine extracts these data elements on a regular (typically nightly) basis from the EHR to the CDR.

Data Standardization

The CDR is primarily populated with data elements used for clinical decision support. All data elements in the CDR are standardized (cross-walked) to one of several coding systems; ICD-9 CM for diagnoses, RxNORM and GCN codes for drugs and SNOMED CT codes for all other data elements. SNOMED CT codes are used for laboratory tests, selected imaging studies and procedures, history, allergies, and family history. SNOMED CT codes have been mapped and reviewed by SNOMED STS, the training and certifying organization for SNOMED CT operated by the College of American Pathologists. In the future, we can also cross-walk the SNOMED-CT codes to LOINC for lab results if this is deemed appropriate. Medication cross-walking poses the greatest challenge since information is needed both at the drug classification and the individual medication level and medications come as single entities and combination drugs. We have coded all single agent medications of interest for general clinical decision support and for our first research project. We are continuing to explore the best use of SNOMED CT versus commercially available drug codes for group classifications and combination drugs, neither of which is currently included in RxNORM. As there is no single coding system that is comprehensive and without disadvantages, we are currently planning to cross-walk medications to both RxNorm and GCN codes as well as capturing NDC codes, when available. We also must incorporate the ability to link both prescription generation data from the EHR to dispensing information from SureScripts-RxHub. As new data elements are added, or if greater detail is required from currently cross-walked codes, new entries will be added to the data dictionary and the CDR will be populated with these codes. All codes are reviewed by in-house physicians, our in-house coding expert, and intermittently by external coding consultants.

Data Export and Presentation to the Grid

The movement of data from the CDR to the ePCRN Gateway database is based on the ASTM standardized Continuity of Care Record. The CINA tools access data in the EHR, standardize the data, and then create Continuity of Care Records locally for each eligible patient at each of the clinical organizations. Continuity of Care Record files consist of an XML string, which is passed to the ePCRN Gateway database (created in MySQL), and the file is parsed into fields that are selectively available to outside Grid enabled queries, effectively de-identifying the dataset.

Data Query

Grid services are used for registering, discovering, and querying databases in DARTNet member organizations. The OGSA-DAI (Open Grid Service Architecture – Data Access & Integration) middleware is used to allow the databases to be exposed as Grid services in a highly secure manner. Queries are dynamically created by the ePCRN Research Portal application. This is a web-based interface that allows complex queries to be developed through an intuitive interface. Once a query is developed the ePCRN Research Portal application submits the query to the OGSA-DAI server, which passes them to each node within DARTNet to be run against the Gateway database. All queries run locally and simultaneously. If the query is designed to return only aggregate data the results are returned in the same session to the OGSA-DAI server. If the results return de-identified patient-level data the query can only be passed into the local server and must be activated by someone behind the local firewall. Results of these queries are returned locally and then, after local permission is given, transferred to the research team. These two additional steps guarantee local control over any patient-level data used for research purposes.

Data Transfer

Aggregate practice-level data are currently being transferred using two different mechanisms, with the expectation that eventually the Globus system will handle all data transfers. While the full Globus installation at each organization is being finalized and fine-tuned, DARTNet can also use a secure FTP transfer from each location to CINA offices and from CINA to a University of Colorado Department of Family Medicine secure server. We use standard secure FTP software for this transfer and have set up the receiving server to transfer received files every 5 minutes to a secure data server, which is not visible to the Internet.

As the system matures, we will use the OGSA-DAI queries to return aggregate data directly from the organizational-level Gateway databases to the University of Minnesota (UMN). We can then aggregate further and move the data between the UMN and other research sites using either the Globus transfer capabilities or secure FTP.

Patient-level data are always extracted to a local directory within each organization. If the select statement is passed over the Globus system it must be activated by clinic personnel or a business partner. Once clinic personnel approve the data extraction, de-identified data are transferred to the research portal using Globus secure transport functions, a secure File Transfer Protocol (FTP) embedded within the Globus security architecture. If the data extraction occurs directly against the CDR (as it does

currently) it must be physically run by CINA personnel or a practice representative (these queries are written by CINA using a stored procedure to create the data set). This data set is then transferred through secure FTP. This redundant approach is more labor intensive but provides a backup data extraction system as we fine-tune the Globus environment.

Natural Language Processing

The use of Natural Language Processing (NLP) is a well established means of extracting data from textual data embedded in medical records. Studies have demonstrated that NLP is a valid and perhaps more valid than chart abstraction by trained abstractors. It is our belief that NLP derived data is no less valid than diagnostic data from claims databases. The free text found in EHRs has significant benefits for clinical research and care, as has been shown for heart failure and chronic angina. Even relatively simple natural language processing methods can be used to “unlock” the valuable patient- and disease-specific information from an EHR. The primary issue to address for ambulatory records is not the efficacy of NLP but the potential paucity of data in ambulatory records.

The DARTNet natural language processing system has been designed initially for processing the text of physical examination, procedural, or history of present illness data. The system has been constructed by adapting publicly available software platform – Unstructured Information Management Architecture (UIMA). One of the unique advantages of UIMA is that it supports distributed Grid-enabled applications, which is critical for the interface with the ePCRN system. Furthermore, our group has prior experience with implementation of natural language processing technology using UIMA.

Our work with text data to date indicates that extracting EHR textual data is highly challenging. Not only is the use of text fields highly variable from system to system, but some systems incorporate patient identifiers into each section of text while other systems’ text fields are completely free of identifying information. DARTNet is in the process of adding a local de-identification system to those systems that embed HIPAA identifiers prior to moving to the next step of submitting the data for natural language processing.

Directed Point-of-Care Data Collection

Even in a highly coded EHR there will be data elements that are essential to understanding the effectiveness or safety related to therapeutic agents that are not likely to be included in a clinical note. Examples include a PHQ-9 score for all depressed patients or a standardized assessment of bipolar symptoms in patients starting to take an antidepressant. In these situations the ability to prompt data collection during an office visit will enable the collection of additional critical data to supplement routine clinical data for selected patients. This ability will allow DARTNet to essentially create a controlled trial environment within the routine care process.

The CINA protocol engine (PE) creates a clinical decision support report for every patient visiting DARTNet practices (Figure 2). This report can also prompt providers to collect specific data at the patient or data element level, based on existing data within the CDR. This ability to fill-in missing data and supplement clinical data is one of the reasons the DARTNet leadership required that any clinical decision support system be able to support point-of-care recommendations. These are two common

possible scenarios: (1) the data are frequently collected during routine care but the results are missing on a large portion of the population of interest or (2) the data are not typically collected to the degree of standardization that is needed for the study in question.

For the first scenario, where the data are often collected during routine care, the CINA PE would be programmed to look for these data within the specified period of time when a patient eligible for the study is being seen. If the data element is present and timely then no prompt would be included on the point-of-care report. If the data element is lacking, a request to collect the information would be generated. The Action Items from Figure 2 are highlighted in Figure 3.

Figure 3. Clinically indicated point of care data collection prompt

Action Items	
___ Document Advanced Directive status	PREV
<hr/>	
DOC: Add Obesity Dx / Plan to Problem List	DM
DOC: Document or perform Diabetic Foot Exam	DM
MED: Start Lipid lowering med* due to LDL goal not met	CAD
LAB: Order Lipid Panel (q 6 mos for CHD risk > 20%).	CAD
LAB: Order Urine Microalbumin (yearly) for DM	DM

In the second scenario, where data standardization needs to be improved (such as a standardized assessment of depression severity) the CINA PE would be programmed to print the full data collection form for patients meeting study criteria. See Figure 4. Depending on the questions to be asked and other factors this could be done only with patients who have provided their consent to participate, or could be done as an extension of clinical care. The results of this standardized assessment would then be entered into the EHR for extraction to the CDR and eventually to the study team.

Figure 4. Study specific point of care data collection instrument

Diabetes Point-of-Care Questionnaire

Please take a moment to answer the following questions about your diabetes. Return to your provider when finished. Thank you!

1. In the last month how many times have you needed to eat or drink something because you felt your blood sugar was too low?

- 0
- 1-2
- 3 or more

2. In the last month has your blood sugar fallen so low that you required assistance from someone else?

- Yes
- No

If you answered yes, how many times in the past month did you have to have someone help you, dial 911 for an ambulance, visit the emergency room, or stay in the hospital? ____

3. Have you taken any of the following herbs or over-the-counter medications in the last month? Please check all that apply.

- | | |
|---|---|
| <input type="checkbox"/> Ginkgo biloba | <input type="checkbox"/> Fenugreek |
| <input type="checkbox"/> Chromium | <input type="checkbox"/> Ginseng |
| <input type="checkbox"/> Cinnamon or cinnamon extract | <input type="checkbox"/> Vanadium/ Vanadyl/ Panchromium |
| <input type="checkbox"/> Garlic extract | <input type="checkbox"/> Vitamin E (not multivitamin) |
| <input type="checkbox"/> Gymnema sylvestre | <input type="checkbox"/> Magnesium |

Are you taking any other herbs or over-the-counter drugs for your diabetes?

- Yes
- No

If you answered yes, please list what else you are taking: _____

Summary of DARTNet Structure

The DARTNet system is readily expandable, utilizing local parallel processing and a two stage data extraction and de-identification process. We estimate that a single central data support site (currently the University of Minnesota) can handle up to 1,000 Gateway databases. While the computer technology allows for management of many more Gateway databases through this central “supernode” it may become logistically difficult to support and track a network of this size through a single central support site. Nonetheless, adding more Gateway databases does not materially slow the distributed query process, which is controlled by the size and complexity of each local Gateway database and the complexity of the query being run.

By adding additional central support sites (or supernodes) the network is essentially infinitely expandable. Furthermore, the data interfaces are not specific to primary care and can be expanded to sub-specialty data when they are available electronically. In fact, the current prototype version includes data from rheumatologists, obstetrician/gynecologists, infectious disease, endocrinology, general surgery and other specialties in limited numbers. When taken to scale, DARTNet will be able to explore both rare safety events in low usage medications and the safety and efficacy of commonly used ambulatory therapies.

Validation of DARTNet

In this initial phase, DARTNet underwent and continues to undergo an evaluation process to ensure provision of data that are valid. Processes were put in place such that DARTNet is subject to ongoing testing and validation processes sufficient to warrant the integrity of the system. Through an approved validation process, DARTNet was “certified” to conduct comparative effectiveness research.

The validation of DARTNet focused on three primary areas: data integrity, software functionality, and system security. The data integrity step included the process of capturing data from each organization’s EHR, standardizing the data, presenting the data for secure internet based queries and transferring query results back to the research staff. This set of activities was validated through 31 different steps. Given that DARTNet uses clinically derived data, it was important to perform detailed testing of all data processing steps to be sure that any dataset used for research was a valid representation of the original clinical data. Software functionality testing for DARTNet included ensuring the performance characteristics of four separate components of the system, and included evaluating 16 functions and processes. As DARTNet expands it will be important that all software components work as expected with a minimal amount of support, to create a sustainable system. Finally, it was critical that when working with a large number of clinicians and health care organizations that DARTNet does not represent a security threat to their patient data or to their patients. Seven steps were used to verify system security.

The Grid security was previously evaluated as part of an earlier NIH contract, therefore retesting the system for unauthorized entry was not done. As for the creation of possible “rogue programs” users are not able to actually develop “programs” on the system. They are limited to the development of queries from a controlled interface. As part of this initial validation work, we have developed and tested a number of queries to attempt to extract data elements that are not “visible” to the Grid query system. None of these attempts were able to extract data from hidden fields. Through this testing we have verified that should a researcher ask for a very large set of data from a single organization it could overwhelm the receiving portal server. Governors on activity of this type will be developed in the second phase of DARTNet work.

The general approach to validating these areas involved User Acceptance Testing. In general, acceptance testing involves running a series of tests on a completed system. Each individual test, known as a Functional Test Case, exercised a particular operating condition or system feature, resulting in a pass or fail. For each of the specifically engineered Functional Test Cases, results were recorded for further analysis and tracking. At the conclusion of the case we analyzed the results and determine if a specific functional requirement was satisfactorily met. User Acceptance Testing may occur at the system level (e.g. tested in one or two locations only) for relatively static processes that are identical from installation to installation. It may occur at the source database level (e.g. once for each EHR, or once for all CDRs) for activities that varied only at that level. Finally, it may occur at the organization and practice level where use of even the same source database may be highly dynamic. This was the case for all clinical data, which can typically be stored in multiple locations within a single EHR and in multiple versions (e.g. a Hemoglobin A1c result that is sent to a central lab for testing and one that is performed at the point of care by the office are typically coded differently and stored in

different locations within an EHR.) The report describing DARTNet's evaluation and validation process can be found in Appendix A.

Governance

In this initial phase, DARTNet was governed by a Board of Directors made up of eleven individuals. Six members were clinicians within member organizations. Four members of the board were members of the various research organizations which support DARTNet. One member was from the Agency for Healthcare Research and Quality. The Board met monthly via conference call. The Board set over-arching policy for the network as a whole. Decisions required a supra-majority of at least nine members to be enacted.

The original plans called for a Scientific Committee made up entirely of clinicians and staff from member organizations. This committee was to review all research projects for general clinician interest and for practice burden. However, in this initial phase the Board elected to review all research projects. It is expected that this committee will be created once DARTNet matures further.

The Board of Directors also provided guidance to the Executive Committee which implemented board policy and made day to day operating decisions. The Executive Committee was made up of five individuals representing each of the research and infrastructure partners. Members of the Executive Committee could not serve on the Board of Directors. The Executive Committee met weekly.

Chapter 3. How a Question Is Asked and Answered

The following section provides a broad overview of a diabetes pilot study conducted during this initial phase of work as a proof of concept of how the DARTNet system can be used to ask and answer a research question. The full report describing the pilot study's methods, analysis, and results can be found in Appendix B.

The first study completed by DARTNet was a retrospective cohort study on patterns of use, comparative effectiveness, and safety of oral diabetes medications for adults with type 2 diabetes. Diabetes is a priority condition under Section 1013 of the Medicare Modernization Act, and an AHRQ-funded Comparative Effectiveness Review (CER) on oral diabetes medications was recently published.⁴ The CER served as the framework for identification of specific aims and hypotheses for the DARTNet pilot research project.

The pilot study's specific aims were to:

1. Examine the comparative effectiveness of single drug and two or more drug combinations of oral hypoglycemic medications as measured by glycemic control.
2. Examine comparative safety of medication related adverse events as measured by liver failure, liver injury and/or elevated hepatic enzymes and rates of hypoglycemic events.
3. Examine patterns of drug utilization over time and by age and gender.

The pilot diabetes research project had two phases. In Phase 1, a retrospective, claims-based study was conducted with the dual goals of: (1) determining the extent to which commercially available, integrated medical claims databases could be used to answer key research questions related to comparative effectiveness and safety of oral diabetes medications for adults with Type 2 diabetes; and (2) identifying areas where such databases are limited, and for which DARTNet (through access to EHR and/or point of care data) may be useful for augmenting and improving upon the results that can be obtained from claims-database studies. This phase examined data on a large number of individuals, using claims data, and examined a limited set of data elements. We used the Ingenix/ICHIS Impact database for Phase 1.

Phase 2 was a replication of the same study using DARTNet data. DARTNet provides more extensive clinical information than can be found through claims databases (see Table 1). We focused on a smaller number of individuals but examined a broader range of data elements. We looked at EHR data using both coded data and natural language processing, and also tested the point-of-care data collection prompts.

Summary findings from the Phase 1 study of oral diabetes medications (ODM) patterns of use, comparative effectiveness, and safety include the following:

GENERAL

- The diabetic cohort from the Ingenix Impact Database is sufficient in size and scope to enable the study of several important aspects of patterns of use, comparative effectiveness, and safety of ODM. Approximately 100,000 diabetic subjects comprised the utilization and safety aims of the study, and a subset of approximately 14,000 subjects comprised the effectiveness aim.
- Among diabetics prescribed ODM, nearly 80% were initiated on monotherapy while 20% were initiated on combination therapy regimens.
- Persistence with initial ODM therapies differed across specific monotherapy and combination therapy groups. Subjects initiated on biguanides or thiazolidinediones (TZDs) had greater persistence than those initiated on other monotherapies; subjects initiated on sulfonylurea (SU)+Biguanides or Biguanides+TZDs had greater persistence than those initiated on other combinations.

COMPARATIVE EFFECTIVENESS

- Unadjusted reductions in Hemoglobin A1C (H-A1C) from baseline to lowest value were similar to previous findings reported in the literature for both monotherapy and combination therapy subjects.⁵⁻⁷ Use of any single ODM resulted in unadjusted reductions in H-A1C of approximately 1%; use of two-drug ODM combinations resulted in unadjusted reductions in H-A1C of about 2%; and use of three-drug combinations resulted in unadjusted reductions in H-A1C of about 2.6%. Adjusted reductions in H-A1C (either baseline to lowest or baseline to last) attenuated some of the crude differences observed by number of ODM received and resulted in the various drug groups becoming more similar in terms of observed, real-world effectiveness. Changes from baseline to last H-A1C were somewhat lower for all agents either alone or in combination.

- Multivariate modeling results on the primary effectiveness outcome showed slight differences across individual ODM drugs or combinations in comparison to metformin monotherapy (statistically significant findings were numerous but many were of questionable clinical significance). Other factors associated with achievement of greater H-A1C reduction included: propensity score, persistence and compliance with therapy, baseline H-A1C, and number of diabetes-related MD and diabetes education visits.

COMPARATIVE SAFETY

- Crude rates of hypoglycemia, liver injury, and liver failure were relatively low (ranging from 0.007 to 0.015 events per person-year of therapy or follow up in the entire study cohort). Unadjusted rates of all three safety outcomes were similar among diabetic subjects whether treated with ODM or not.
- Multivariate modeling results showed that as compared to those receiving metformin monotherapy, users of sulfonylureas (either alone or in combination with other ODM) were at greater risk of hypoglycemic events and liver injury, but not liver failure. No such increases in risk (relative to metformin monotherapy) were observed for patients receiving TZDs or for those receiving statins concurrently. Other factors associated with adverse safety outcomes included renal dysfunction and certain specific other diagnoses and medications associated with these outcomes.
- Propensity adjustment was used in all comparative effectiveness and safety models; neither stratification of effectiveness and safety model results by propensity score quintiles, nor sensitivity analysis using a multiple propensity score approach resulted in significant changes to the principal findings or conclusions of the study.

Phase 2 results are summarized as follows. We identified a large panel of subjects (N=35,215) that met the study criteria. They had similar age and gender distributions, Chronic Disease Indicator (CDI) scores,⁸ and patterns of initial ODM prescribing as the subjects from the claims database. The Chronic Disease Indicator score is a chronic disease index that approximates the number of chronic diseases a subject/patient has, using prescription medication fill/refill data (as contrasted with indices relying on ICD-9-CM diagnosis codes). It is interesting to note that even in its “proof of concept” stages, the DARTNet system was able to identify similarly sized panels of diabetic subjects and subjects receiving various ODM to enable analyses of similar power to the claims based studies performed in Phase 1.

The time to first regimen change for monotherapy and combination therapy groups in the DARTNet Diabetes Replication Cohort was generally similar to those seen in phase 1 related to drugs most frequently used and mean time to regimen change. However, median time to regimen change was shorter in the DARTNet Diabetes Replication Cohort.

The replication of crude effectiveness outcomes from Phase 1 using the DARTNet Diabetes Replication Cohort shows that baseline HA1C values were very similar in the DARTNet Diabetes Replication Cohort as in the claims cohort, as were crude changes in HA1C from baseline to lowest and baseline to last. Slightly smaller reductions were observed in the crude DARTNet results but the patterns were similar

(e.g., greater reduction from baseline to lowest than from baseline to last, substantial variation in reduction, and slightly greater reduction for combination regimens versus monotherapy regimens).

Rates of hypoglycemia, liver injury, and liver failure were determined from EHR data. As was the case in Phase 1, rates of hypoglycemia, liver injury, and liver failure were rare and did not differ substantially across major ODM drugs or classes. Some slight differences in crude rates of these outcomes was noted, but as these are crude (unadjusted) rates, no firm or definitive conclusions can or should be drawn at this time regarding the comparative safety outcomes in the Phase 2 replication cohort. Total numbers of safety events are high enough in the replication cohort to suggest that safety studies can be conducted in the DARTNet system. Liver injury and liver failure events could be further explored in greater detail in the DARTNet system to enable case/outcome validation for these severe outcomes. Also, hypoglycemic events appear to be very low, as was observed using the claims data in Phase 1, suggesting that point of care and/or patient-reported sources of clinical data are likely needed to better study hypoglycemia outcomes in future studies.

These initial pilot study results lead to us to conclude that the DARTNet prototype tested and confirmed its ability to conduct observational comparative effectiveness studies that yield data that are valid and reliable using our secondary data analysis as a benchmark. This analysis in fact provided important confirmatory evidence regarding the effectiveness and safety of various ODM therapies. We expect that the significant benefits of using DARTNet data over claims data will become more evident when we examine research questions in which clinical data is an essential part of the analyses.

Chapter 4. Lessons Learned and Future Directions

We anticipate the important lessons learned to date will help us in establishing a highly adaptable, collaborative, and result-driven organizational structure for DARTNet.

Several lessons from this initial phase point to the importance of working with trusted colleagues capable of working collaboratively for the good of the whole. By engaging in frequent and frank communication, usually via weekly teleconference, the partners have been able to work through sensitive issues related to the sharing of intellectual property, appropriate divisions of labor, and agreeing on a joint vision for DARTNet. Sharing of intellectual property presented itself as an important issue related to DARTNet's contractual relationships with CINA and University of Minnesota. The DARTNet system is currently dependent on the capabilities of CINA's clinical decision support system, which is the most important asset of CINA as a corporation. Therefore, the contractual agreement established with CINA had to recognize and protect its proprietary interests while at the same time ensuring that DARTNet would have access and rights to the components necessary for the replication and propagation of DARTNet should CINA's involvement with DARTNet come to an end. It's important to note that although we have searched for other alternatives for the extraction, transformation and CCR generation functions of the CINA software used in DARTNet, at this time we can find no other software that can perform this full set of activities. Thus, the only

alternative to using the CINA software would be to accept CCRs directly from an EHR (CCRs are coded and created from software embedded in some EHRs). There is only limited experience with this as most EHRs still cannot produce CCRs, but the limited experience that we could find indicates that the embedded systems do not take into the account the various locations and data coding options within an EHR and thus are very poor representations of the true record.

Similar issues are true for University of Minnesota. While the Globus Toolkit is an open source software, the University of Minnesota has made substantial investments and has proprietary rights to the ePCRN Gateway, which is a critical component of DARTNet's ability to transfer data from CINA's CDR to the Grid to make it available for distributed query. University of Minnesota's contract also had to be carefully negotiated to protect University of Minnesota and to allow for DARTNet to continue should their involvement end. As for the use of the Grid, we can see two alternative options: first, we could use secure FTP of files directly from the CDR and second, we could develop a web services interface. We have used the FTP option while awaiting stabilization of the Grid system and found that it works for small numbers of sites but quickly becomes unwieldy. As the number of sites grows it requires hands-on work from CINA at each site. The second option requires developing and coding a new interface, security and data model. At this point, we believe the choice to use the Grid solution is the logical choice to stay synchronized with other major networks using this technology, such as those sponsored by the CDC.

The division of labor between CINA and University of Minnesota has also presented interesting lessons for DARTNet. Although these two entities had limited previous experience with each other, they have successfully worked collaboratively to facilitate each other's contractual obligations while advancing DARTNet as whole. This collaborative approach has evolved rapidly between these two very capable contractors, each with proprietary interests, confident in their own abilities, but particularly at the beginning still uncertain of one another. It is extremely important that an appropriate division of labor and monetary incentives be negotiated early on with contractors, which takes into account their unique contributions while still prioritizing the immediate needs of DARTNet. Based on this experience, the DARTNet leadership is likely to further expand the detail in our performance-based contracts to clearly define key deliverables, delivery dates, and corresponding reimbursement in any future collaborations.

This initial phase of work has also shown the challenges in working with primary care practices with varied IT capacity. Practices with fewer than five providers tended not to have dedicated IT staff. In some cases, these practices had contracted with competent outside resources and in other cases IT support was provided by a friend, family member, or a small unsophisticated consultant. In the latter cases, it was sometimes difficult to get access or appropriate permissions needed on the local computer system. In addition, small practices often used a local cable or phone company that served as the Internet service provider and provided the communications hardware used in the clinic. This hardware often had poor or no documentation and nobody at the phone or cable company could assist or answer questions. In all cases, these barriers were overcome, but sometimes with great difficulty and a lot of time and effort. Another example of practices' technological limitations presented when it came time to install the Gateway database and software, a Globus service that generally resides locally at the practice or at

the practice's designated technology center. In order to install the Gateway, practices needed to obtain a static IP address in keeping with Internet 2 criteria for secure information exchange. Many small primary care practices were unable to secure a static IP address for their Internet connection. However, through modification of the Globus server software, these limited practices could still participate in DARTNet in a secure way through a centralized technical support center provided by CINA. The capability of allowing practices to participate independently through a single technical support center provides enhanced flexibility to the DARTNet software. Across the country, many small practices are loosely affiliated to share the expenses of technical support but still maintain independent administration. This additional capability enhances the potential for participation in DARTNet by practices with this organizational model.

Many lessons have been learned about using EHR data for research in this first phase of work. The major historical problems with using data from practice-based EHRs have included incomplete, un-coded, highly variable data, and the inability to aggregate data from separate practices. By looking at actual EHR system use by the eight study practices (even those using the same EHR system), it was re-confirmed that wide variations in system use exist. This is also true for individual providers within the same practice. A single event or concept (e.g., lab value, procedure, family history, behavioral factor) was likely to be recorded in multiple locations and in multiple ways even within a single EHR. Further, there were almost no consistent underlying coding systems used within any of the EHRs. Consequently, extracting information directly from the systems would not likely yield a useful or complete representation of the recorded patient data. However, the *CINA Mapper*[®] software was able to locate and classify data with a high degree of accuracy and completeness, resulting in the ability to extract comprehensive and accurate data and to store it in a standardized and consistent database. This allowed data to be aggregated across practices and further analyzed.

In addition to the completeness and variability problems discussed above, there were wide variations in the way data from a given encounter were correlated. It was nearly impossible to view an "episode of care" spanning multiple encounters. Much of the problem with correlating data may be addressed by the CINA software, but correlating data from multiple encounters into an "episode" would require some degree of inference. Additionally, we found that the concept of a set of data related to a "visit" is more difficult to ascertain from an EHR than expected. To improve this ability DARTNet is moving rapidly to include billing data from each organization, which better aggregates information into a visit paradigm.

A further issue with data quality was the configuration of several systems that limited providers to entering only selected data, while leaving important data un-addressed within an encounter. It appears that many of the EHR systems currently being used permit users to record data they consider important, while not recording other data. There doesn't appear to be adherence to the concept of a "minimum data set" of information that should be required for each encounter. It should not be assumed that "not recorded means not done," yet for analysis purposes in this initial work one may have to make this assumption. We have observed that production of a point-of care reminder system, either designed into the EHR or as a separate function, may encourage providers to enter more complete information for each encounter, thereby improving the quality of data for analytic purposes.

An additional area of potential weakness in using EHRs as a source for research data is in the area of free-text narrative: typed or dictated notes. One of the study questions dealt with extracting clinical data from these sources. One solution would be to apply natural language processing technology to free-form text. However, we found that EHRs typically had one or more specific areas that clinicians would use to record these data, and that clinicians were, in most cases, willing to enter the required data in those defined areas even if they had already addressed these in their text notes. Further, the amount of redundant data entry of this type was fairly limited. While natural language processing appears to be an important component of DARTNet in the long run, we were not able to find significant value in our initial work, primarily because the free text data in the systems was very sparse.

Our pilot research indicates that prompted point-of-care data collection will greatly enhance the ability of DARTNet to provide new comparative effectiveness data beyond those available from claims data. Using EHR data we found that the rate of hypoglycemic events appears to be very low. While the ability to capture events such as mild hypoglycemia in an EHR is difficult, point-of-care data collection presents an opportunity to prompt for documentation of these events within an EHR, including prompting users to ask patients about such events. Traditional OCER work using claims data finds only major events that result in an emergency department or hospital visit or are actually coded for within an ambulatory visit. Our preliminary point-of-care work shows that many more hypoglycemic events occur than are captured using claims data. Thus, it is clear that DARTNet will extend the ability to capture various events that may not result in a charge and that occur as part of an illness to a much greater extent than our current claims based approaches.

A final area of data quality involves the lack of editing within systems for reasonableness of data. We encountered data that was clearly inaccurately entered, such as improbable heights or weights or illogical lab results, usually in areas where data were entered manually by clinicians (versus being imported from other data systems). Based on this, it is safe to assume that some data that is within normal limits may nevertheless be inaccurate. Analysis of these data should consider the possibility of data entry error. Only limited numerical data elements are hand entered—vital signs and point of care lab tests primarily—but these data could be further improved if logical range checks were included against these fields within an EHR. Nonetheless, keeping in mind the issues mentioned previously, we think that, with the proper tools to correct for variations within EHRs, valuable research data may be obtained from EHRs.

Additionally, improved analytic capability of EHR products is needed to allow the use of data contained in the record by practices, by networks of practices, by researchers, and by others (ignoring for the moment any objections that may exist to such data sharing). The DARTNet prototype gives us a glimpse into the possibilities of using commercial middleware to solve many data aggregation, sharing, and analysis problems that current EHR products cannot address. It is our hope that further linkages and interoperable functions will be released into the marketplace that can be used to further leverage the power of DARTNet by allowing it to be more easily connected to EHR products, and to facilitate the growth of the network. Until that point, DARTNet will seek public and private partnerships to link and aggregate EHR data across practices to

activate both quality improvement at the practice level, and the aggregation of clinical data for OCER work.

Although this project was primarily focused on data obtainable from EHRs, we also sought to determine the extent to which other data available to the practice could be used as a source for comparative effectiveness research. With the exception of a few values not represented in administrative data (e.g., height, weight, blood pressure, smoking status, alcohol use), there are proxies for many clinical measures in administrative data. For example, we can determine that a procedure such as a colonoscopy had been performed if there were a charge for the procedure. However, any value that is not individually billed cannot be inferred from billing data. On the other hand, administrative data from third parties often contains more or different information than can be found from an EHR. For example, if administrative data from payers were available to practices, providers may be able to determine that certain procedures had been performed by other providers, even if the patient forgot or neglected to mention it to the primary care provider or if the provider failed to ask the patient. Likewise, data from a payer or pharmacy benefits management company reflects prescriptions actually filled if the payer paid for the medication. Providers typically have to depend on patient's statements and do not have access to this type of corroborating data. With the potential for medication fulfillment and other payer data being available to practices under the emerging "patient centered medical home" model, the opportunity of obtaining clinically relevant and valuable data from administrative data may become possible and should be a useful addition for OCER research.

Among the most important lessons is a greater realization of some of the problems that can potentially arise with the intentional or unintentional misuse of this powerful new tool. Anticipated new challenges revolve primarily around the ability to potentially break or overload the DARTNet system, unintentionally compromise privacy or security resulting in potentially serious consequences, or degrade system functionality because of a lack of understanding of the underlying system. One example is a researcher who asks complex questions about many different or unrelated data elements. Although a new Globus patch that allows for timing out and stopping a distributed query has been integrated into the system, a mischievous or poorly trained researcher could potentially degrade the entire DARTNet system including potentially overloading and shutting down central system servers with a few poorly conceived queries. Another scenario is that of the "greedy researcher"—who out of ignorance or ambition simply says "just give me all your data." The current software version does not address this satisfactorily. When a researcher requests data to be transferred, multiple threads of information will come streaming in simultaneously from practices all across the country. The two step process for obtaining patient level data makes this scenario unlikely, though still possible. It would not take a large number of large clinical systems all moving their data at the same time to overload the Internet bandwidth of any University. Creating mechanisms to protect from these possibilities were not part of the initial pilot and feasibility testing, but they will be essential to the delivery of high quality and consistent service in a research production environment where demands are high and hence a primary task in DARTNet's second phase of work.

We learned a number of lessons in the conduct of OCER studies within DARTNet. As OCER research is an evolving and multidisciplinary science, there are

many possible approaches to studying any given research question. Thus, there is also considerable debate as to the relative merits of each approach and the many decisions that must be made to implement an OCER study regardless of the chosen approach. It would be very useful to establish an “ongoing” peer review process for OCER studies within DARTNet, or the AHRQ comparative effectiveness group overall, to establish mechanisms for soliciting and providing expert advice on clinical and methodological matters before conducting complex mathematical analyses rather than after the fact (as is the case with traditional peer review). Further, it is important to note that even after OCER studies are completed considerable disagreement and debate may remain regarding the merits and value of individual studies; as such, no single OCER study can (or should) be viewed as definitive. Even with the promise of DARTNet and the addition of previously unavailable clinical/EHR data there will continue to be debate about the ideal approach for OCER studies.

Our first review of clinical data indicates that as the DARTNet system is refined there will be significant additions to claims databases for OCER work. For instance, alcohol intake, smoking status, disease severity and minor drug side effects, as well as self treated hypoglycemic events are types of data that should allow extensions of traditional OCER studies. It has also become clear that claims data from network providers will significantly add to the clinical data from EHRs. The addition of these types of data is a high priority for DARTNet and plans for data acquisition are already underway.

In its second phase of work, the DARTNet team has been tasked with developing a more robust management and governance structure that will support its potential rapid growth and be responsive to various stakeholders with interest in working with the network. We recognized that our current Board of Directors (BOD) does not exercise the authority nor does it function as a typical fiduciary or corporate board. We believe that as DARTNet expands and aspires to become a self-sustaining organization, a number of possible corporate models may be considered—including a 501 (c) (3), Limited Liability Corporation, or maintenance of University or other academic affiliation. Given the likely scenario of some form of new corporate structure, the formation of a governing BOD would be appropriate. For the near future, DARTNet continues to be a federally-funded research project that needs a governance structure that allows the project PI and AHRQ’s Task Order Officer to exercise the authority to make all final decisions. As such, the governance of DARTNet in the next phase of work will consist of an executive committee, an oversight and advisory committee (instead of a BOD) whose functions will be to guide and make recommendations to the executive committee about DARTNet’s short and long-term operating policies and goals, formulate and enforce policy to assess, address and monitor potential conflicts of interest, and evaluate the benefits of engaging in future partnerships and research projects. DARTNet’s revised management structure will be comprised of four core work groups that will execute the day to day operations of the network: Administrative Core, Technical Core, Research Core and Practice Network Core. The University of Colorado will be the administrative, research, and technical home for DARTNet. The clinical network home will be the AAFP National Research Network. The four work groups will report to the Executive Committee.

Future Direction and Challenges

DARTNet is now a proven prototype capable of bi-directional electronic communication with EHR enabled medical practices and holds great potential for becoming a valuable tool for Observational Comparative Effectiveness Research (OCER) to inform healthcare policy and provide the information technology backbone necessary for clinical practices to design, track, and evaluate quality improvement initiatives.

AHRQ is the most important funding source for the network and its continued support in DARTNet's second phase of work will be critical to the realization of DARTNet's potential. AHRQ has provided DARTNet with additional funding under a second task order to further expand the network to include a larger number of clinics, including specialty and general pediatric practices, to conduct an observational comparative effectiveness and safety study of different therapies for major depression, and to further test and develop DARTNet's ability to collect point-of-care data.

Also critical to advancing DARTNet's future is defining the appropriate next steps and laying down the necessary foundation for the network to accomplish its aims in the second phase of work:

1. Expanding technical capabilities: The ability of DARTNet to move into production mode rather than operate as a prototype will require technical additions and refinements.
2. Expanding the patient base: The primary care networks that comprise DARTNet must have adequate patient numbers and diversity to assure that inquiries are capable of informing policy for general and specific populations, including those served by Medicare, Medicaid, and SCHIP.
3. Defining future research questions: The specific research questions that DARTNet would address next will be based upon the needs of stakeholders, emergence of new clinical issues as well as the availability of EHR data. They will be critical factors for further development of the network and its research infrastructure.
4. Defining a future corporate and organizational structure: DARTNet will require a structure that is nimble, self-sustaining, allows for efficient use of resources with reasonable overhead requirements, and is capable of receiving funding from multiple sources, such as the federal government, not-for-profit organizations, foundations, etc.

With these critical factors in mind, DARTNet is poised to undertake a rapid expansion in the next 18 months in its second phase of work. A full proposal for the new task order can be referenced for a detailed description of this new scope of work.

Future Research Questions

In its second phase of work DARTNet will conduct an observational comparative effectiveness and safety study of different therapies for major depression. First we will use commercially available national claims data to conduct a population-based cohort study of new episodes of major depression. This study will describe the patterns of use of antidepressants and combinations of other psychotropic drugs, both alone and in combination with psychotherapy. We will examine patterns of treatment utilization for pediatric and adult populations separately, and by primary care versus specialty mental

health care diagnosing of major depression and prescribing psychotropics. Next we will use DARTNet data to evaluate the responsiveness and remission rates for major depression treatment, beginning with population-based data and enhanced by clinical data. We will measure rates of adverse effects to include the most common medical side effects causing treatment discontinuation, and suicidality, including ideation and attempts.

We will evaluate the ability of DARTNet to track episodes of depression care and relate patient specific outcomes to these episodes. Using DARTNet, we will attempt to delineate episodes of care for depression, distinguish advent of new from on-going episodes of care, and track process and clinical outcomes within and at the end of each episode of care. This information will inform depression care and the evaluation of comparative effectiveness, given that reaching end points of care and managing the chronicity of depression are central goals for clinical care.⁹

Immediate Next Steps

The next steps to address the necessary components of the depression study as well as future studies are to ensure:

1. **Access to prescription fulfillment data:** Plans and negotiations to gain access to this data using the services of SureScripts/RxHub are well underway. This work had been initially planned for the prototype phase but delayed due to a merger of these two organizations.
2. **Streamline and automate existing capabilities:** To move past prototype phase, DARTNet must streamline manual and complex technical processes, certify DARTNet for additional inquiries, and refine point-of-care and natural language processing capabilities.
3. **Secure ongoing participation of practices:** As DARTNet expands it must develop and implement mechanisms to incent and engage current and new practices to participate. This will include the creation of a robust learning community and the integration of other academic units into the network.
4. **Expand the population base:** Currently DARTNet has access to the health records of approximately 400,000 patients. Considering the relative rarity of suicide attempts, we propose that undertaking the depression study described above will require a substantial expansion of the DARTNet network. DARTNet is poised to accomplish up to a six-fold increase by partnering with a large and stable Regional Health Information Organization (RHIO). The initial groundwork for this expansion with a viable RHIO has already been laid.

Strategic Issues

DARTNet currently exists within the administrative structure of the University of Colorado Denver. We recognize that in the future the corporate and organizational structure of DARTNet must allow for:

1. **Building of a diverse research portfolio:** To date, DARTNet has secured three additional grants/contracts (two from AHRQ and one from the AAFP Foundation) to address specific policy questions. Collaborative efforts with HMO Research Network (HMORN, a consortium of 15 HMO organizations)

are currently underway to explore the feasibility of using DARTNet for gathering disease data for surveillance purposes. DARTNet will continue seeking out grants/contracts to conduct at least three studies per year to cover its minimum financial requirements related to personnel and resources for the DARTNet infrastructure.

2. **Assured independence and efficiency:** While currently the University of Colorado is a suitable home for DARTNet, there may be distinct advantages if DARTNet were to be organized within a private corporation model. These advantages include more efficient overhead rates, less bureaucracy-related expenses, and greater independence. For any such corporate “spin-off” to coincide with DARTNet’s aims, access to government and foundation funding streams would need to be assured. Possible models could include remaining with the University, affiliation with or becoming a 501 (c) (3), or becoming a private corporation capable of subcontracting to an academic institution or 501(c) (3) (e.g. a limited liability corporation). Clearly, the governance of any such “spin off” would evolve from the current Board of Directors structure to one best suited to the policy and fiduciary role of a private corporation.

We believe AHRQ’s renewed financial investment in the network for a second phase of work will go a long way to move DARTNet past the “tipping point” to become a viable enterprise that is capable of sustaining itself with grants and contracts in the future.

Blueprint for the Future

DARTNet has received a contract from AHRQ to expand its size, to conduct OCER related to depression care, and to further refine its governance and sustainability models. DARTNet leaders recognize the immediate need to rapidly expand the population base of DARTNet as well as the scope of practitioners included within the network. Expansion will focus on inviting large multi-specialty groups, regional health information organizations, other academic partners and Independent Practice Associations (IPA) to join the system.

Sustainability will be addressed by solidifying the governance structure and diversifying funding sources as well as the research team. The inclusion of additional academic partners will both add large group practices as well as address the need to expand the DARTNet research team.

The financial model of DARTNet recognizes three distinct phases of development. The first phase was creating the basic infrastructure/software of the federated system (described in this report). The second phase of development for DARTNet is the addition of clinical data, which will be an on-going process; for example, the latest task order supports minor expansions of the current data set and medication fulfillment data will expand over time. This phase of the network will continue to be built over time as other stakeholders seek to utilize the DARTNet system and as other researchers obtain funds for projects requiring new data elements. The data acquisition component of DARTNet also involves the ongoing development of linkages to various EHRs, and new members. These costs are covered in a variety of ways including funding for various research projects, infrastructure support grants and through member decisions to personally pay for connectivity. The third phase of the DARTNet

financial model is funding ongoing operational costs. Once the federated system is fully in place, routine maintenance of the technical component is relatively inexpensive, though not free. Ongoing funding of personnel to support the system may require the development of an organization outside any one academic home as discussed above. Given the broad interest in DARTNet among key stakeholders, DARTNet should be able to continue to obtain funding as long it continues to improve the data within the network.

Supporting DARTNet as a Learning Community

All DARTNet practice and clinician members will be invited to participate in quality improvement activities to form a learning community. Learning activities will include the use of benchmarking reports to identify best practices and top performers. Top performers will be invited to present their work models, tips and advice to other DARTNet members via webinars, blogs of personal conversations and informal knowledge exchanges, and other social networking activities. Member practices will be surveyed to identify perceived learning needs. The learning community will support consultations and visits from practice facilitators, online resources, webinars, and communication vehicles to promote learning within and among DARTNet practices.

In the near future, we envision DARTNet as an established learning laboratory that can provide ways to better understand variations in care and outcomes across practice settings, practice types and organizational models. We plan to use benchmarking reports to further explore practice variabilities through quantitative and qualitative means. Practice facilitators will also be used to assist practices in using the data and knowledge exchange derived from their participation in DARTNet to enhance their own quality improvement efforts and improve the care of their patients.

Conclusion

The DARTNet is a federated network of electronic health record (EHR) data from eight organizations representing over 500 clinicians and over 400,000 patients. DARTNet has demonstrated it can answer questions concerning the safety and effectiveness of medications and medical devices and can bring together a community of practices to improve clinical care.

DARTNet is a proven prototype that is capable of bi-directional electronic communication with EHR enabled medical practices. The DARTNet prototype holds great potential for becoming a valuable tool for Observational Comparative Effectiveness Research (OCER), to inform healthcare policy and provide the information technology backbone necessary for primary care practices to design, track, and evaluate quality improvement initiatives.

In developing and testing the prototype, DARTNet has contributed to the understanding of patterns of use, effectiveness, and safety of oral hypoglycemics in treating patients with type 2 diabetes. DARTNet investigators have also identified critical clinical data elements derived from EHR data that will be useful in enhancing the state of the art in observational comparative effectiveness research (OCER). Through this type of research, DARTNet's ultimate aim is to improve the understanding of the effectiveness of various drugs and devices—and by this means improve the information available to policymakers in healthcare programs such as Medicare, Medicaid, and SCHIP.

Critical to the ability of DARTNet to evolve as a fully functional system is the need for further technical refinement of the system itself, as well as a significant expansion of the patient base. Only through such expansion will DARTNet realize its potential to inform decision-makers regarding relatively rare events, and to provide representative information for subsets of the general population.

AHRQ has provided DARTNet with further financial support to move forward with its development and expansion during a second phase of work. Once fully established, DARTNet will have the capability to inform multiple decision-makers and help bring about better informed health care policy decisions, and will also become a valuable tool to transform and improve care in the emerging patient-centered medical home.

Acknowledgments

We extend our deepest appreciation to the Agency for Healthcare Research and Quality for funding this work and to our task order officer, Dr. Gurvaneet Randhawa, for providing steadfast vision and leadership from the inception of this network to its current form.

We're grateful to the eight medical practices that constitute the core of DARTNet: Medical Clinic of North Texas, WellMed Medical Group, Tiena Health, Wilmington Health Associates, University of Minnesota, Cranford Family Medicine, Family Health Center of Joplin, and University of Colorado. Without these practices and their champions this network would not exist.

We also acknowledge the many valuable contributions of our Oversight Committee: Cathy Bryan, Robert Eidus, Terrence Feehery, James Galliher, Diane Madlon-Kay, Richard Penaloza, Gary Piefer, Robert Phillips, Lister Robinson, and Brian Webster.

Lastly, we would like to thank a number of individuals who have worked shoulder to shoulder with us for the past year and half to make this network a reality. Our special thanks to Elias Brandt, Jessica Huff, Jim May, Cathy Bryan, Kevin Peterson, Mark Janowick, Heather Orton, Elizabeth Staton, Bill LeBlanc, Stephanie Mitchell, and Sherry Holcomb.

References

1. Tunis SR, Stryer DB, Clancy CM. Practical clinical trials: Increasing the value of clinical research for decision making in clinical and health policy. *JAMA*. 2003 Sep 24;290(12):1624-1632.
2. Agency for Healthcare Research and Quality (AHRQ). AHRQ Support for Primary Care Practice-Based Research Networks (PBRNs). Available at: <http://www.ahrq.gov/research/pbrn/pbrnfact.htm>. Accessed May 13, 2009.
3. Mold JW, Peterson KA. Primary care practice-based research networks: Working at the interface between research and quality improvement. *Annals of Family Medicine* 2005 May-Jun;3(Suppl 1):S12-20.
4. Bolen S, Wilson L, Feldman L, et al. *Comparative Effectiveness and Safety of Oral Diabetes medications for Adults with Type 2 Diabetes*. Rockville, MD: Agency for Healthcare Research and Quality; 2007 Jul. Available at: www.effectivehealthcare.ahrq.gov/reports/final.cfm. Accessed: 2007.
5. Lincoff AM, Wolski K, Nicholls SJ, Nissen SE. Pioglitazone and risk of cardiovascular events in patients with type 2 diabetes mellitus: a meta-analysis of randomized trials. *JAMA* 2007 Sep;298(10):1180-1188.
6. Karter AJ, Moffet HH, Liu J, et al. Glycemic response to newly initiated diabetes therapies. *Am J Manag Care* 2007 Nov;13(11):598-606.
7. Bolen S, Feldman L, Vassy J, et al. Systematic review: comparative effectiveness and safety of oral medications for type 2 diabetes mellitus. *Ann Intern Med* 2007 Sep 18;147(6):386-399.
8. Malone DC, Billups SJ, Valuck RJ, Carter BL. Development of a chronic disease indicator score using a Veterans Affairs Medical Center medication database. IMPROVE Investigators. *J Clin Epidemiol* 1999 Jun;52(6):551-557.
9. Agency for Health Care Research and Quality. Distributed Ambulatory Research in Therapeutics Network (DARTNet): Informing Practice and Improving Depression Care. Draft Abstract published 8 May 2009: Research in Progress; Available at: <http://effectivehealthcare.ahrq.gov/healthInfo.cfm?infotype=nr&ProcessID=94>.

Glossary

AAFP NRN	American Academy of Family Physicians National Research Network
ASTM-International	An international standards setting organization, formally called the American Society for Testing and Materials
CCR	Continuity of Care Record, a core dataset of the most relevant administrative, demographic, and clinical information about a person's health care
CDR	Clinical Data Repository, a database of de-identified, standardized data stored locally at each organization; the CDR prepares data for presentation to the Internet for distributed queries
CINA	Clinical Integration Networks of America, Inc.
Data cross walking	Labeling data elements with recognized standards such as SNOMED CT or RxNORM
Data mapping	Finding various locations and representations of like data within an electronic health record
EHR	Electronic Health Record
ePCRN	electronic Primary Care Research Network - a group of practice-base research network installing a Grid enabled data system
Federated distributed database	Centrally accessible database where the data is distributed across a wide area
FTP	file transfer protocol – a protocol for transferring data from one computer to another
Gateway database	Grid enabled local database created by ePCRN
Globus Consortium	international open access consortium developing Grid solutions for distributed database access and queries
Grid	Internet based software system that connects similar or disparate databases to allow secure central access to a distributed database system
ICD-9CM	International Classification of Diseases 9th edition clinically modified –codes used to classify diseases and a wide variety of signs, symptoms, etc.
HMORN	HMO Research Network - a consortium of 15 HMO organizations that have formal, recognized research capabilities
MySQL	Open access database
ODM	Oral diabetes medications
OGSA-DAI	Open Grid Architecture Data Access and Integration
PBRN	Practice-based Research Network
PE	CINA Protocol Engine
POC	Point of Care data collection - data collection for research purposes, which takes place during clinical care
RHIO	Regional Health Information Exchange
RxNORM	A National Library of Medicine coding system for medications
SNOMED CT	Systematized Nomenclature of Medicine Clinical Terminology
UIMA	Unstructured Information Management Architecture – a natural language processing software system that supports distributed Grid-enabled applications
UMN	University of Minnesota Center for Excellence in Primary Care
XML	Extensible Markup Language, a general purpose specification for creating custom markup languages to help share structured data