Pulling Back the Covers:
Technical Lessons of a Real-World Health Information Exchange

Atif Zafar, M.D.[a,b], Brian E. Dixon, M.P.A.[a]

[a]Regenstrief Institute, Inc., Indiana
[b]School of Medicine, Indiana University, Indiana

Abstract

Several nations and local communities are striving to achieve widespread, secure exchange of clinical data between various health care providers and public health organizations. Most of the literature on health information exchange focuses on the financial, political, and privacy aspects of these initiatives. Perhaps just as important are the technical and organizational factors that have influenced development of data exchange methods and results.

One mature network in the Midwestern United States has had success in establishing consistent, secure exchange of clinical data for more than ten years. Presented here are the technical lessons learned and design decisions made from this initiative with the hope that they can be used by others striving to connect disparate clinical information systems for the improvement of health care quality and safety.

Keywords:

computerized medical records systems,
computer communication networks, equipment reuse

Introduction

Everything in health care ultimately revolves around the accessibility and effective use of clinical data. When physicians, nurses, and other health care professionals have the information they need when they need it, they serve patients better, in terms of both the quality and the safety of the care they provide. Making these data available electronically, then, appears to make good sense.

Electronic clinical data abound. The problem is that they are often inaccessible to providers, because health care organizations tend to house their clinical data in distinct, isolated repositories. Many providers and policy makers now recognize that the sharing of data among hospitals, doctors, and other health care organizations in a given city, state, or region often referred to as health information exchange (HIE) can make health care safer, more efficient, and more effective [1].

The Indiana Network for Patient Care (INPC)

Indianapolis has pioneered an extremely successful HIE initiative, the Indiana Network for Patient Care (INPC), launched in 1993 under the leadership of the Regenstrief Institute. Initially, the INPC provided data from one hospital to providers in emergency departments at three other hospitals. By 2005, a more mature network, with a membership comprising 95 percent of all hospital and emergency care in Indianapolis, expanded to include providers in other parts of the state. By the end of 2006, the INPC contained more than six million distinct patient registration records, 850 million discrete observations, 17 million text reports, 50 million radiology images, and 40 million orders.

A number of factors have contributed to the INPCs success, including political and legal dynamics, which have been addressed in other publications [2][3][4]. Here we outline the reasons for the technical design decisions and functionality of the INPC, highlighting the technological and organizational factors that have contributed to the networks growth, ease of use, and sustainability.

Technological factors

Regenstrief has examined, deployed, refined, and evaluated a variety of operating systems, programming languages, software applications, and database management systems over its thirty-year history. The philosophical approach has been to *select a technology, stick to it, and make it work*. This means we rarely make radical changes (e.g., redesign a program written in one language using another language) unless there is a clear need (e.g., the new language is far superior to the old one). For example, early use of the Web (prior to 1995) for results aggregation brought many challenges. To enable asynchronous communication between clients and servers, we developed customized tools that possessed modern asynchronous JavaScript and XML (AJAX) functionality. Only recently have we begun to redesign our tools to utilize current AJAX frameworks.

Although sometimes slow to change, *our organization is not afraid to experiment*. We have worked with state-of-the-art image-compression technologies, such as JPEG

2000 and Wavelets, and motion video (MPEG) for display of radiology images and cardiac echo movies, respectively, in Web browsers. We have also experimented with voice-recognition and voice-annotation for clinical notes and display of laboratory results on mobile devices in wireless settings. Currently, we are experimenting with nomadic computing technologies so that clinicians can have an access-anywhere system for clinical data.

The INPC is composed of many moving parts, some of which are legacies while others are more modern. It is not the formula of a certain operating system with a specific database management system that has produced success for the INPC. Rather, the technical success of the INPC may be attributed to its adherence to these philosophies and the principle that, *when possible, one should build upon existing infrastructures rather than inventing or implementing new ones*. Below we discuss this principle in the context of the INPCs security, speed, flexibility, and reusability.

Security

Secure exchange of information between the INPC and participants operates using point-to-point connections. In the past, the INPC has employed T1 lines (data pipes as weve referred to them in other publications). We are now phasing out T1 lines in favor of virtual private networks (VPNs).

INPC security policy dictates the use of up-to-date encryption methods and good password hygiene and RSA keys. Currently the INPC relies on 128-bit SSL encryption to protect data on the rare occasion we use the public Internet to exchange information. Passwords must be changed at specified intervals, require a certain combination of letters and numbers, have a minimum length, and cannot be reused by the same individual.

The INPC also requires users to sign a confidentiality agreement and devices to be equipped with time-limit controls to prohibit unauthorized access. However, users do not log in directly to the INPC. User authentication is done via providers. First, users login to a provider portal or local area network (LAN), then they access INPC applications through a gateway created between the provider and the INPC network. The INPC relies on providers to implement access and time-limit controls on devices and ensure that users have signed a confidentiality agreement, which is required of them anyway to access the providers electronic resources.

Federated data sharing model

At its core, the INPC is a series of federated vaults, sometimes referred to as edge proxies or silos, storing data from the various participating institutions. Each institution has its own privileged silo where only data from that institution resides. The architecture of each silo closely resembles that of the Regenstrief Medical Record System (RMRS), a well known electronic medical record system [5]. A simplified data model of the RMRS is presented in Figure 1.

Figure 1 - Simplified data model for the Regenstrief Medical Record System (RMRS)

Each silo represents mirrored data from one participating institution. Patient registry data, such as name, medical record number, and date of birth, and clinical data, like laboratory results, immunizations, and free-text notes recorded by the doctor during an encounter (e.g., clinical visit), are stored in the silos.

Silos can be created technically using a variety of methods. Hierarchical databases using large flat files running on clustered VMS nodes have been employed in the past. More modern relational databases can also be used to develop silos. Each silo can function as a separate database, or indices can partition data stored in a table on separate physical disks.

The specific technologies employed to create silos are not as important as the concept. Creating federated vaults gives participants peace of mind that their data will be segregated and secure. Yet data in federated vaults can exist within a single network access storage (NAS) unit, reducing latency when retrieving data during a clinical encounter. To date, the centralized, federated model developed by Regenstrief has yielded better performance (speed) when compared to decentralized federated networks used by other HIE initiatives.

A centralized, federated model also simplifies the process of data standardization ensuring identical blood urea nitrogen (BUN) results from various laboratories are interpreted the same way. The burden of reconciling various tests is shifted from individual provider organizations to Regenstrief. We can employ a single data model and dictionary and resolve errors as they arrive at the central hub for processing. Troubleshooting and mapping data elements to the standardized model requires overhead approximately 1-3 FTEs for the volume of messages we process. However, the costs and occasional painstaking mapping efforts are justified, because *we value the quality of the data stored in the INPC*. Our methods ensure that the

data retrieved from the INPC is reliable and standardized, which builds trust amongst network users and makes retrieval and delivery methods efficient.

Standards

Since its inception, the INPC has strived to provide optimal access to relevant clinical data at the point of care. To that end, the INPC has invested significant time and resources into the development and use of health information technology (health IT) standards (e.g., HL7, LOINC, CPT, etc.). These standards permit disparate systems to share data among one another, making them *interoperable*. They also permit the INPC to quickly add new types of data by reducing the time required to create customized interfaces for information delivery to providers.

Standards also enable data reusability, the ability to store a single concept and use it multiple times in a variety of applications. For example, a physician may order an HIV 1 AB (LOINC #7917-8) to indicate the presence of HIV in a patient. Once the result is reported to the INPC from the lab, three separate actions can be taken using the same HL7 message and LOINC code. First, the result would be stored in the institutional silo corresponding to the provider identified in the message. Second, the result could be delivered electronically to the physician using a clinical messaging application. Finally, the result could be reported to one or more public health agencies.

The example demonstrates that a single element, a standardized clinical message, can be used by three very different components of the INPC to store and exchange clinical information. Standards are employed so that the provider, physician, and health department interpret the result in the same way (e.g., all three receive a message indicating a positive value for LOINC #7917-8). This reuse of the same data is efficient, flexible, and cost-effective.

At the time of inception for the INPC, standards were *immature* and *limited*. We had to *invent* standardized methods for transmitting and mapping data between networked provider organizations. One such invention, the LOINC standard, was created because SNOMED and other existing terminologies lacked breadth for laboratory and some clinical concepts. We hope that other organizations can benefit from early experiments by us and other organizations. We believe that *field tested standards*, such as HL7 2.x and LOINC, can help others create interconnected systems in less than half the time it has taken us to develop the INPC.

Although we are experimenting with HL7 Version 3, current INPC members continue to transmit data using HL7 2.x. We encourage continued development and refinement of standards, and we will support them as they mature and become adopted by INPC participants.

Applications

An important lesson learned from building an **aggregated, standardized** data repository is that *data can then be re-used for many applications*. For example, the same dataset that is reported out to clinicians using a clinical messaging application, we call ours DOCS4DOCS, can be sent to the State Health Department for communicable disease reporting. Similarly, data received from private practices could be aggregated and presented to an ER physician for delivering emergency care. These compelling applications allow INPC *stakeholders* to get some value-added by joining the collaborative, with an understanding that *their data* will only be used according to the agreed upon terms within the data-sharing contract.

Component-based architecture

The INPC employs a number of component technologies to process much of the data that travel across the network. Technologies like *interface engines, message processors*, and a *global patient and provider index* perform specific tasks that are generic enough to be re-used from application to application. Thus each component can be optimized for its task and easily modified to include a newly interfaced system. This creates a network in which components are not only interoperable but also reusable.

The idea of developing and reusing components is not unique to Regenstrief or the INPC. The object-oriented paradigm has influenced software development practices, with conventional modular techniques abandoned in favor of component-based approaches [6]. This is especially true in the open source software movement [7].

Regenstrief has embodied the philosophy of component-based development into the INPC, which has enabled the network to remain flexible. Expansion over the last thirteen years has involved the addition of new participants (e.g., hospitals, laboratories), new applications that use the data for a variety of tasks, and new forms of data (e.g., we added pathology reports in 2003 and dictated notes in 2005). With each new addition, the network has required slight modification. Development time is shorter, because components can be reconfigured and redeployed faster than monolithic programs. New components can be developed more quickly, because insertion into the network does not require recompilation or reconfiguration of other components.

Organizational factors

Designing, constructing, and operation of a working technical infrastructure for interoperable exchange do not guarantee success. In addition to its technical infrastructure successes, the INPC has also benefited from a number of organizational factors that have shaped its development over the last 13 years.

Incremental evolution

Incremental change has played a significant role in the INPCs long-term success. What began as an experiment to connect emergency rooms together slowly evolved into a large network that provides clinical information to emergency rooms, hospital staff in other departments, and ambulatory providers. This growth was guided by steady leadership that focused consistently on the INPC vision rather than on trends in the budding HIE industry.

Some HIE projects can and will evolve more quickly than the INPC. However, leaders of such projects should temper expansion with a clear vision for their network and agreement from all their partners.

Human resources

Technology is not the single most critical factor for successful HIE. To succeed in the development of a broad clinical data exchange, the INPC employed many capable people to manage and support the technology of the network. A knowledgeable staff is necessary on both ends of the network, at each participating organization as well as at the data exchange entity. Technical difficulties and bugs are inevitable, so capable humans are needed to troubleshoot errors, resolve data issues, and continue to move the vision of the exchange forward. And as the INPC has grown, so too has its need for more staff members to effectively monitor all of its members relations and data connections.

For example, we recently had a lab send us an HL7 message using unexpected units (up/mL). The INPC exception processor detected the anomaly (unrecognized units), which resulted in 26,000 records being dumped into an exception queue for analysis. Turns out the lab system had an embedded typo (the units should have been ug/mL), and the problem was resolved after a phone call and a few emails.

A more common problem we face is reporting of units in any other field, usually the notes field, except the appropriate HL7 units field (OBX-6). This is a problem common to all the labs from which we receive data, and it is a recurring problem for newly created tests.

Exceptions require human intervention as subtle differences between common clinical concepts are difficult for computers to resolve, despite several attempts in the past [8][9].Given the need for regular human intervention, we employ 2-3 FTEs to constantly monitor and troubleshoot the more than 150 message streams from the major hospital systems in Indianapolis, regional referral laboratories, specialty providers, several rural providers throughout the state of Indiana. We predict the need to add trained personnel in the future as the INPC continues to grow.

Our human resources also drive innovation. For years we have benefited from the talents of National Library of Medicine (NLM) informatics fellows typically post-doctoral physicians. These individuals have been key players in designing, creating, testing, and evaluating innovative components of the INPC infrastructure, including add-on programs such as CHICA^{TM} and PHESS^{TM} that extend the INPC beyond clinical messaging. Many of these fellows have stayed on as faculty at Regenstrief and the Indiana University School of Medicine, continuing to enhance the INPC and mentoring new fellows.

Sustainability

The ability of senior leadership to repeatedly make a clear, evidence-based business case for the INPC has contributed significantly to its sustainability. Initial funding for the INPC came from the NLMs high-performance computing and communication initiative. Subsequent funding has come from the NLM, the Agency for Healthcare Research and Quality (AHRQ), the Health Resources and Services Administration (HRSA), the National Cancer Institute (NCI), the Indiana Genomics Initiative, and the Indiana Twenty-First-Century Fund [2]. Each grant supporting a portion of the INPCs development enabled Regenstrief to measure clinical, financial, and community outcomes. These data provided support to the INPCs business case, which allowed the network to secure additional funding for expansion of existing services and development of new ones.

In addition, the networks interoperable, flexible design supports a variety of clinical and research activities. Applications such as CareWeb^{TM}, DOCS4DOCS, the Shared Pathology Information Network (SPIN) anonymous query tool, and the Public Health Emergency Surveillance System (PHESS^{TM}) for syndromic surveillance build on the INPCs core infrastructure. However, each creates a unique service for all or specific network members. An innovative, legally separate organization, the Indiana Health Information Exchange (IHIE), has also capitalized on the INPC infrastructure, creating a highly reliable, customer-oriented organization to support care delivery organizations using Regenstrief technology. By vertically expanding the INPC in this way, senior leadership has successfully created new resource opportunities for the INPC to grow and improve.

Conclusion

Reliable, up-to-date clinical data at the point of care remain the key to improving both the quality and safety of health care. Successful exchange of clinical data occurs only when all participating providers and organizations share not only data, but an understanding of what those data mean. Incremental change and growth are key to the success of data exchange networks. Over time, effective networks tend to expand the types of data they carry, as well as the applications for those data. Standards and reusable components help HIEs to maximize their efficiency through shorter development time and lower costs creating opportunities for integration with new systems and organizations. If other exchanges are as successful as the INPC, valuable improvements in care will be achieved in many communities.

References

[1]   Walker J, Pan E, Johnston D, Adler-Milstein J, Bates DW, and Middleton B. The value of health care information exchange and interoperability. *Health Affairs*. 2005

[cited 2006 May 1]. Available from: http://content.healthaffairs.org/cgi/content/abstract/hthaff.w5.10

[2] [2] Biondich PG and Grannis SJ. The Indiana network for patient care: An integrated clinical information system informed by over thirty years of experience. *J Public Health Mgmt*. 2004 Nov; 10 Suppl 1:S81-6.

[3] McDonald CM, Overhage JM, Barnes M, Schadow G, Blevins L, Dexter PR, Mamlin B, and the INPC Management Committee. The Indiana network for patient care: A working local health information infrastructure. *Health Affairs*. 2005; 24(5): 1214-1220.

[4] Sears C, Prescott VM, and McDonald CJ. The Indiana network for patient care: A case study of a successful healthcare data sharing agreement. *ABA Health eSource* [serial on the Internet]. 2005 [cited 2006 Nov 29]; 2(1). Available from: http://www.abanet.org/health/esource/vol2no1/Sears.pdf.

[5] McDonald CJ, Overhage JM, Tierney WM, Dexter PR, Martin DK, Suico JG, Zafar A, Schadow G, Blevins L, Glazener T, Meeks-Johnson J, Lemmon L, Warvel J, Porterfield B, Warvel J, Cassidy P, Lindbergh D, Belsito A, Tucker M, Williams B, Wodniak C. The Regenstrief Medical Record System: a quarter century experience. *Int J Med Inform*. 1999 Jun; 54(3):225-53.

[6] Waguespack L and Schiano WT. Component-based IS architecture. *Information Systems Management*. 2004; 21(3): 53-60.

[7] Madanmohan TR and Rahul D. Open source reuse in commercial firms. *IEEE Software*. 2004 Nov-Dec; 21(6): 62-69.

[8] Schadow G, McDonald CJ, Suico JG, Fohring U, and Tolxdorff T. Units of measure in clinical information systems. *J Am Med Inform Assoc*. 1999 Mar-Apr; 6(2): 151-62.

[9] Friedman C and Johnson SB. Natural language and text processing in biomedicine. In: Shortliffe EH and Cimino JJ, eds. Biomedical Informatics. 3rd ed. New York: Springer, 2006.

Address for correspondence

Brian E. Dixon, M.P.A.
Regenstrief Institute, Inc.
410 West 10th Street, Suite 2000
Indianapolis, IN 46202
U.S.A.
(317) 423-5582
bdixon@regenstrief.org