

Enabling Clinical Decision Support in a Multi-Institutional Telemedicine Infrastructure

Laura A. Noiro, BS, Brent D. Launsby, BS, Thomas C. Bailey, MD

Division of Medical Informatics, Washington University School of Medicine

noiro@informatics.wustl.edu

INTRODUCTION

We have previously described a set of clinical decision support (CDS) applications that identify drug-dosing errors and drug-drug interactions^{1,2}. These systems had access to clinical information through point-to-point interfaces with the legacy applications that supplied the data once daily. Since real-time data and cross-institutional deployment of these applications were needed, we investigated reimplementing of these applications within a two-state, multi-institutional telemedicine infrastructure called Project Spectrum (PS)³. Due to limitations of the PS clinical data repository and a better understanding of the work distribution between data warehouses and databases for CDS, we designed an architecture that would allow us to take full advantage of PS and provide the best environment for CDS. The remainder of this paper will present the architecture chosen for building a real-time operational data store (ODS) for CDS and results of the current production system.

ARCHITECTURE

The Washington University School of Medicine CDS applications require data from the following clinical

domains: patient-specific demographics, encounters, vital signs, laboratory results, and pharmacy orders. All clinical data except pharmacy orders are available in the PS repository; facility-specific pharmacy HL7 messages flow from an existing real-time interface engine.

The overall architecture models the warehouse architecture proposed by Inman⁴. Our specific implementation of this architecture utilizes the PS repository as the primary data source and HL7 messages as a secondary data source. As depicted in Figure 1, these real-time data sources utilize tools to extract information from legacy systems, translate the data into consistent encoded data feeds, and send the data to the ODS environment via HL7 messages. This environment parses HL7 messages, performs additional translation on the data, and stores the information in the ODS⁵. The ODS supports the data needs for the CDS applications as well as daily queries.

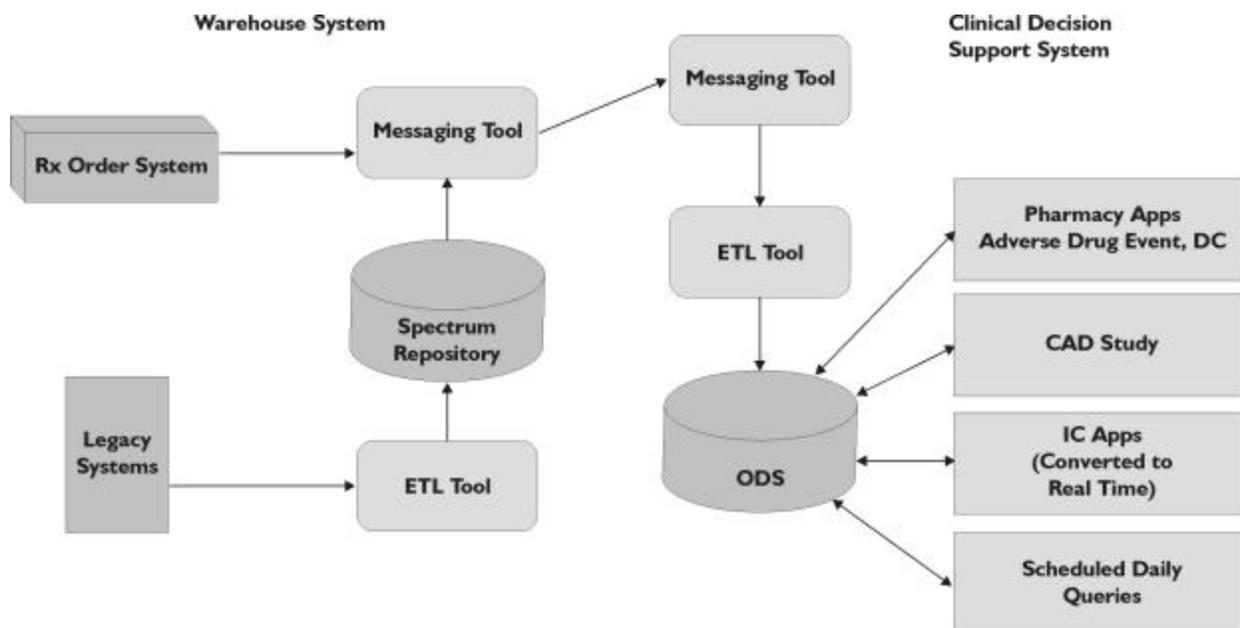


Figure 1. Real-Time Data Capture for the Clinical Decision Support Operational Data Store

This architecture uses standards-based, commercially available software. Specifically, the architecture consists of a commercial platform-

independent tool for fail-safe HL7 message delivery from the repository to the ODS, a relational database as the ODS, and a commercially available

tool for HL7 message parsing. The translate-and-load application as well as the monitoring applications were custom built using Unix script, C, and SQL. Since most of the translation and data encoding occur before the data are stored into the PS repository, the ODS translation procedures are fairly simple, with the exception of merge and unmerge processing. Merge and unmerge processing is not handled by the PS repository, which creates a burden for the ODS.

The monitoring processes are a crucial aspect of our ODS. These processes ensure that data are flowing between the two systems, attempt to self-correct problems as they occur, and log potential problems for followup by staff members. These processes drastically reduce the amount of staffing required to maintain the data feeds.

RESULTS

The ODS hardware is a Sun Enterprise Server 250 with two 300 Mhz UltraSPARC II CPUs, 1 G memory, and a 100 Mhz Ethernet card.

We are currently collecting demographic, encounter, laboratory, vital signs, merge, and unmerge data from eight hospitals. These data have been collected since August 1999. We are collecting drug orders from two hospitals—Barnes-Jewish Hospital since August 1999 and Christian Hospital since March 2000. The database is approximately 6 GB. We will begin purging old data soon, as the CDS applications typically need less than 1 year's worth of data to perform their operations.

We have six data feeds—demographics, encounter, laboratory and vital signs, merge, unmerge, and pharmacy. In a typical month (August 2000), these feeds provided 325,960, 348,495, 703,347, 15,083, 3, and 287,419 messages, respectively, for a total of 1,680,307. We plan to support outpatient registration (910,110 messages/month) and an additional facility pharmacy data feed (64,000 messages/month) without performance degradation to the system. It is difficult to accurately determine the time between order entry and receipt of information in the ODS; however, we are confident that data will arrive in less than 1 minute when all systems are functioning normally. We dedicate .50 of a full-time employee (1,040 hours/year) to supporting the data feeds and the ODS. This allows time for some new development; however, most of this time is spent testing changes to the existing data interfaces. Additional resources are provided by BJC Health System to maintain the link between the legacy systems and the PS repository. BJC also provides resources for investigating and correcting problems that track back to the source systems. Over the past 6 months, we have had approximately 12 after-hours notifications relating to the data feeds. Generally, the problem is that a source system has stopped sending data.

CONCLUSION

We leveraged telemedicine technologies to successfully build a real-time, relational database that supports CDS applications and daily queries. This database is called an operational data store. The ODS and data feeds were built using standards-based commercial tools with some Unix script, C, and SQL source. Therefore, it is generally applicable to any health care organization that has a clinical data repository or HL7 messaging.

The ongoing maintenance of the system is reasonable; the biggest time commitment is testing new and changing data feeds and tracking data inconsistencies.

ACKNOWLEDGMENT

This project was funded by National Library of Medicine Contract No. N01-LM-6-3550.

REFERENCES

1. McMullin ST, Reichley RM, Kahn MG, Dunagan WC, Bailey TC. Use of an expert system to identify potential dosing errors at a large university hospital. *Am J Healthcare Pharm* 1997;54:545-549.
2. McMullin ST, Reichley RM, Watson LA, Steib SA, Kahn MG, Bailey TC. "A Web-Based Clinical Information System To Detect Dangerous Drug Combinations at a Large University Hospital." Presentation at the Missouri Society of Hospital Pharmacists 1997 Annual Meeting, April 9-11, 1997.
3. Kahn MG. Computerized decision support systems for health providers. NLM Grant Application, 1996.
4. Inman WH. The Corporate Information Factory and Enterprise Infrastructure, February 2000. http://www.billinmon.com/library/library_frame.html
5. McNamee LA, Launsby BD, Frisse ME, Lehmann R, Ebker K. Scaling an expert system data mart: more facilities in real-time. *Proc AMIA Symp* 1998;498-502.