

EMERSE

Electronic Medical Record Search Engine

Technical Manual

project-emerse.org

David Hanauer <hanauer@umich.edu>

Last updated March 27, 2015

Important Note

This Technical Manual is a work-in-progress. We welcome any feedback to improve it, including corrections, additions, and other clarifications. If at any point you get stuck or run into problems, please contact us so we can help. The main point of contact is David Hanauer <hanauer@umich.edu>.

Future plans for EMERSE

For those wondering what is planned for EMERSE, here are a few items:

- Addition of an 'advanced' search option to allow power users to use regular expressions.
- Ability to search across all patients, without first defining a patient cohort
- Ability to search imaged/scanned documents
- Filters for reducing results returns (eg, only certain document types)
- Additional synonyms (they are already being added periodically)
- Improved attestation screens
- Ability to mark patients as being of interest or not, possibly to record other notes as well.
- New timeline data visualization
- Improved help screens for users
- If adoption beyond U of Michigan occurs, sharing of synonyms and search term bundles across instances
- More security, including the ability to prevent certain users from being able to look up patients (essentially limiting them to specific patient lists)
- Shared patient lists (to specific individuals, not to everyone like the bundles)

Still to do for this Manual:

- * Discuss synonyms and how to add/update
- * Discuss authenticating users with LDAP or others
- * How to add users
- * How to delete users
- * Sys admin features, if any
- * Queries to check on system stats

Background

EMERSE is the Electronic Medical Record Search Engine. It was designed to work with the free text (unstructured) clinical documents in an electronic health record (EHR) system. Most importantly, EMERSE was made for regular users, not those with IT or informatics expertise. It was developed at the University of Michigan in 2005 and has been continuously improved and updated since then. EMERSE has been used to support a variety of tasks, including clinical and translational research, internal quality improvement and quality assurance initiatives, as well as for hospital operational support tasks. It has been used very successfully by the billing and coding team for complex case reviews, improving reimbursement rates by nearly \$1 million per year. The software is used routinely by our compliance office, risk management, and infection control departments. It has also been used in the clinical setting, to support rapid data retrieval to answer clinical questions as well as to reduce the burden of prior authorizations, required by many insurance companies. EMERSE has also been used to support hundreds of research projects, resulting in many peer-reviewed publications.

EMERSE can integrate documents from multiple sources. At the University of Michigan, EMERSE has tens of millions of documents, including those from our original, homegrown EHR system, CareWeb, as well as newer documents from our replacement vendor EHR, Epic. EMERSE was designed to be vendor-neutral -- as long as you can get your documents out of the system, EMERSE should be able to search through them. We've even created an interface to move cohorts identified through the i2b2 Workbench directly into EMERSE for further searching of their free text notes.

EMERSE was created with support from the University of Michigan Comprehensive Cancer Center, the CTSA-supported Michigan Institute for Clinical and Health Research (MICHR), the University of Michigan Health System through the Medical Center Information Technology (MCIT) Department, and the University of Michigan Medical School through the Medical School Information Services (MSIS) Department. We are making EMERSE available at no cost for academic use, and for a reasonable fee to support ongoing development for other uses.

EMERSE provides features to help get your work done quickly and accurately. For example, the tool includes a large list of synonyms and related keywords to help you expand your search even if you don't know the right terms to use. This includes a large collection of trade and generic drug names, acronyms and abbreviations, and other wording variations including spelling errors.

EMERSE also provides intuitive and novel ways to visualize the search results, to help users focus on the information they need. For example, search results can be viewed as a 'heat map', showing the density of documents with a search hit, as well as what we call the 'mosaic view' which shows which specific terms appear for each patient in the search results.

Before you start

To be able to get EMERSE up and running there are several things you should know.

While the system was designed to be simple to use for non-technical people, installing the system requires significant technical expertise, and will likely need some degree of local customization. We expect those installing EMERSE will have expertise or experience in databases (SQL), Apache Lucene and Solr, and web servers such as Apache Tomcat.

Once a license has been signed (no cost for academic use) we can provide a fully operational system on a virtual machine that includes PubMed abstracts as an example data source. This is good for understanding how the system works, but is not ideal for operational use. We can also provide the system packed up as a .war file which can be placed in to the Tomcat `webapps` directory as described in the instruction. If you want the original source code, we can send an export of an Eclipse project. The source build is easiest if you use the brand of Eclipse that we use (spring source toolsuite), but it can also be built via Maven/command line.

It is also important to remember that if your institution will be using EMERSE to search clinical documents, it is necessary to ensure that all security and privacy regulations are followed. EMERSE should be installed by IT and informatics professionals who are familiar with the requirements for securing protected health information and maintaining the right infrastructure/IT environment to ensure that the data remain secure behind a firewall.

Hardware Requirements

Like most systems, EMERSE will always perform better with faster hardware. However, it should be possible to get a reasonably performing system without major investment. At the University of Michigan our production system is currently running on a 4-core Intel box with 12 GB RAM. Indices are store on a hard disk-based SAN, currently with 2 TB allocated. Solid state drive (SSD) storage would be expensive but would definitely improve the performance of the system. We estimate that performance would be increased by a factor of 2, potentially more.

With local hosting, we estimate an upfront hardware cost of \$10-20k, but that range could be very different depending on the local cost of hardware at your institution. Some of this will depend on how many users (especially number of concurrent users) you expect, and fewer users You could, of course use less expensive hardware, but please remember that *the user experience is vital for the success of EMERSE, and fast response times are vital to a good user experience.*

We recommend that the Oracle database that supports EMERSE be on a separate server from EMERSE itself, but they could both reside on the same server if needed.

It also possible to run everything on virtual machines but we generally do not recommend that approach. Depending on the VM configuration and host, performance is generally better when the server process connected to the disk is not through a virtualized storage layer. This is because the underlying Lucene indexes used by EMERSE should be as low latency/high throughput as is practical.

High Level Overview of Core Components

There are three inter-related but distinct components needed for a full installation of EMERSE. It is important to note that components (1) and (2) are necessary for EMERSE to run, but getting these two components up and running are independent of setting up EMERSE. We do provide guidance on the issues related to these two components. It is also worth noting that components (1) and (2) would likely be required for any type of search engine you would implement.

(1) Document repository or repositories

This is needed for Lucene to locate and index documents. This repository is not needed at run-time for EMERSE since it is not accessed by EMERSE. Rather, Lucene will create its own local copy when creating the indices. Lucene can be pointed to multiple sources, and it is not entirely necessary that you create a standalone document repository (you could, for example, connect Lucene directly to your EHR using whatever APIs are available). However, we recommend use of the repository since it will make it easier to handle and track instances when documents have changed (last updated, for example) as well as complete re-indexing (as opposed to incremental indexing) which may be required from time to time.

(2) Apache Lucene

EMERSE leverages the Apache Lucene library (jar files) to tokenize document texts in a format that enables fast retrieval. When users execute a search within EMERSE, the application code uses the Lucene API to open the index files and search for matching documents. The EMERSE application itself does not provide functionality to build the Lucene index files needed by the application. At UMHS, Apache SOLR is used to aid with the creation and maintenance of Lucene formatted indexes. It is a web based application that provides REST oriented API's that enable adding and removing content from Lucene indexes. At UMHS, as data from source systems changes, the changed documents are retrieved from the system and sent to the SOLR API for inclusion in EMERSE.

(3) EMERSE

The EMERSE software runs independently of, but is dependent upon, functional Lucene indices. At startup, EMERSE is pointed at the pre-existing Lucene indices, and then the magic happens.

Code Dependencies

EMERSE leverages multiple open source components, which are listed below.

Server side dependencies

groupid	artifactId	version	license	more info
org.springframework	spring-core	3.1.0.RELEASE	Apache 2	
org.springframework	spring-context	3.1.0.RELEASE	Apache 2	
org.springframework	spring-beans	3.1.0.RELEASE	Apache 2	
org.springframework	spring-aop	3.1.0.RELEASE	Apache 2	
org.springframework	spring-tx	3.1.0.RELEASE	Apache 2	
org.springframework	spring-webmvc	3.1.0.RELEASE	Apache 2	
org.springframework	spring-aspects	3.1.0.RELEASE	Apache 2	
org.springframework.security	spring-security-core	3.1.0.RELEASE	Apache 2	
org.springframework.security	spring-security-config	3.1.0.RELEASE	Apache 2	
org.springframework.security	spring-security-web	3.1.0.RELEASE	Apache 2	
org.springframework.security	spring-security-ldap	3.1.0.RELEASE	Apache 2	
org.springframework	spring-jms	3.1.0.RELEASE	Apache 2	
org.aspectj	aspectjrt	1.6.8	Eclipse 1.0	
org.aspectj	aspectjweaver	1.6.8	Eclipse 1.0	
org.springframework	spring-orm	3.1.0.RELEASE	Apache 2	
org.hibernate	hibernate-core	3.6.7.Final	LGPL 2.1	
org.hibernate	hibernate-entitymanager	3.6.7.Final	LGPL 2.1	
org.hibernate	hibernate-envers	3.6.7.Final	LGPL 2.1	
oracle	oracle-driver*	11.1.0.7.0	OTN	http://www.oracle.com/technetwork/licenses/distribution-license-152002.html
c3p0	c3p0	0.9.1.1	LGPL 2.0	
junit	junit	4.8.1	Eclipse 1.0	
org.codehaus.jackson	jackson-core-asl	1.8.5	Apache 2	
org.codehaus.jackson	jackson-mapper-asl	1.8.5	Apache 2	
org.springframework.webflow	spring-js	2.3.0.RELEASE	Apache 2	
org.slf4j	slf4j-log4j12	1.6.0	MIT	
commons-logging	commons-logging	1.1.3	Apache 2	
log4j	log4j	1.2.17	Apache 2	
commons-httpclient	commons-httpclient	3.1	Apache 2	
commons-fileupload	commons-fileupload	1.2	Apache 2	
commons-io	commons-io	1.3.2	Apache 2	
org.apache.activemq	activemq-client	5.8.0	Apache 2	
org.apache.activemq	activemq-broker	5.8.0	Apache 2	
net.sf.trove4j	trove4j	3.1a1	BSD	
org.springframework.batch	spring-batch-core	2.1.0.RELEASE	Apache 2	
org.springframework.integration	spring-integration-core	2.0.4.RELEASE	Apache 2	
org.springframework.integration	spring-integration-jms	2.0.4.RELEASE	Apache 2	
org.springframework.integration	spring-integration-stream	2.0.4.RELEASE	Apache 2	
org.springframework.integration	spring-integration-jmx	2.0.4.RELEASE	Apache 2	
org.apache.solr	solr-core	4.7.2	Apache 2	
joda-time	joda-time	1.6	Apache 2	

Client JavaScript dependencies

name	version	license	more info
jquery	1.7.1	MIT/GPL	
keyboard.js	0.2.2		https://github.com/RobertWHurst/KeyboardJS/blob/master/license.txt
date.js	1.0 Alpha-1	MIT	http://www.datejs.com/
json2.js		PUBLIC	http://www.JSON.org/json2.js
amplify.js		MIT	
knockout-2.1.0	2.1.0	MIT	
knockout-validation	1	MIT	
jquery-ui	1.8.16	MIT/GPL	
jquery-idletimer	0.8.092209	MIT	
jquery-json	2.3	MIT	
jquery.iframe	1.2.5	MIT	
jquery.fileupload	5.5.2	MIT/GPL	
jquery.balloon	3.0	MIT/GPL	
jquery.metadata		MIT/GPL	
jquery.rating	3.14	MIT/GPL	

Document Repository

While not entirely necessary, we recommend a document repository for Lucene to use. Documents do not need to be in a single repository, since it is possible to point Lucene to multiple repositories if needed. We recommend using a standard SQL database (Oracle, MySQL, PostgreSQL, etc) because of its support for backup/restore, etc. Also, we have found that it is easier to use such a database to keep track of document changes that occur after indexing. This allows us to more quickly identify documents that may need re-indexing. In theory it could be possible to point EMERSE directly at an electronic health record system if the right APIs are in place, but this might make it much harder to identify incremental changes in the data and thus might require much broader scanning of the entire repository each time for any potential changes.

It is also easier to manage data this way. When a new document source is to be added, one only needs to create a new table with the documents and metadata, and then populate the table with the source data. (Of course, one still will need to update Lucene to point to this new table, and to add the source into EMERSE as well.)

At the University of Michigan we use Oracle for our document stores (we have more than one), and the stores are organized by source systems. These stores include the documents themselves as well as metadata including medical record number, document date, last updated, as well as details that depend on the source such as clinical service, clinical provider name, etc. These metadata are used in the display for the users within EMERSE, although last updated is important for knowing if a document needs to be re-indexed.

Note that this document repository is not a core part of EMERSE and is not something that we set up or include in the EMERSE software. It is, however, a pre-requisite for Lucene to be able to retrieve and index the documents. It actually does not need to be running for EMERSE to function, but must be up and running when indexing by Lucene is occurring. Lucene makes a copy of the document it indexes, so Lucene serves up the documents at runtime for EMERSE, not the document repository.

Documents should ideally be in the form of plain ASCII text or HTML. At Michigan, since we also receive documents in RTF format, we use a commercial software package (Aspose.Words, <http://www.aspose.com/word-component-suite.aspx>) to perform this conversion from RTF to HTML for storage in our repository.

To capture clinical notes that reside in Epic, an HL7 based interface that emits HL7 messages containing the note in RTF format when notes are edited and signed has been configured. A java process then takes the content of the messages, and converts them to HTML using Aspose.Words, and stores it in the document repository. Available metadata in the HL7 message is also stored in the repository along with the note, such as encounter date, department, and edit date. We recommend that you try to only include finalized, or 'signed'

notes in the repository to prevent the need to continually monitor for document changes and frequent re-processing and re-indexing.

A small amount of Java code was written to retrieve the documents from the repository and other UMHS datasources via JDBC, and execute the SOLR API to update EMERSE's Lucene indexes. This code is executed periodically to keep the Lucene indexes up to date. This code is not part of the standard EMERSE release since the specific needs will vary for each institution.

For example, if PDF documents are stored, one could use a tool like Apache Tika to extract the text to present it to Solr.

EMERSE Deployment Guide

Deployment Overview

The high level items involved in deploying EMERSE include preparing the application server, initializing the database, and configuring the Lucene indices. These instructions will guide you through each of these topics. The diagram below illustrates the high-level architecture of EMERSE. Additional information related to customizing the Lucene indices to adapt EMERSE to your specific environment can be found later in this document.

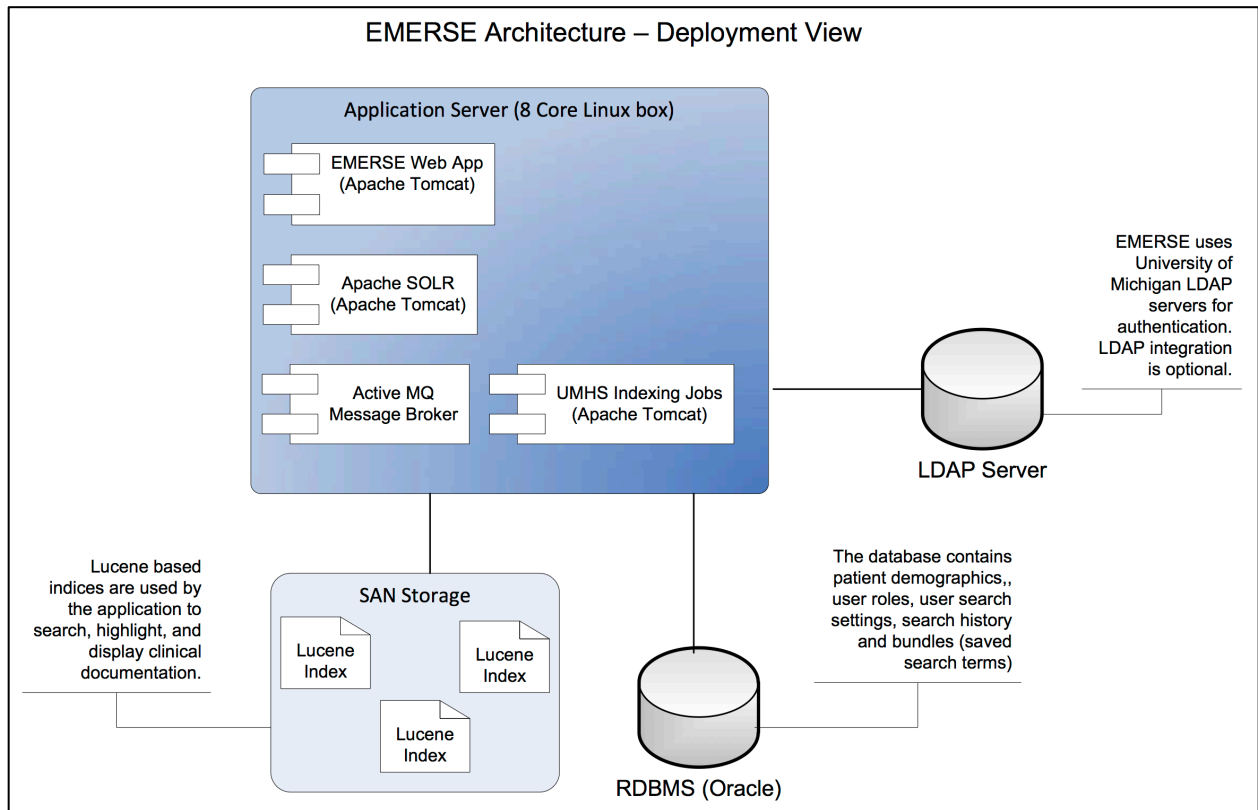


Figure. High-level architecture diagram.

Pre-Requisites

This guide assumes a few pre-requisite infrastructure items are already in place:

1. An Oracle server should be installed on a server and an account/schema created that allows full access to create common database objects, such as sequences, tables, indices etc. EMERSE doesn't place great demand on the Oracle database, so a relatively small server can be used with 10-50 gigs of storage allocated for user tablespaces. Currently at the University of Michigan we are using only 3 GB of space for production EMERSE with 600+ users and 2 years of data.
2. A Linux/Unix based server that will be used to install the application server and host the indexing services. This server should be connected to the highest speed storage available. Capacity is dependent on number of documents to be indexed. At UMHS 600 GB is in use to host approximately 80 million documents. If your documents are heavily formatted (such as RTF instead of TXT), storage requirements will be higher.

Application Server Installation

This section covers the process of installation and configuring of the application server where the EMERSE application will be deployed. An instance of Apache Active MQ, an EMERSE dependency will also be installed. The main pieces needing installation are:

1. Java Development Kit/SDK (JDK)
2. Apache Tomcat (Java Servlet Engine)
3. Apache ActiveMQ (A message broker)
4. EMERSE Web Archive File (WAR file) deployment and configuration

1. Java JDK

The first step in installing EMERSE is to download and install a Java Development Kit on the server. We recommend version 7 at this time.

Download Site:

<http://www.oracle.com/technetwork/java/javase/downloads/jdk7-downloads-1880260.html>

Choose the correct install that matches the target computer's operating system. If the target workstation is a 32 bit OS then choose the 32 bit download option, otherwise choose the appropriate 64 bit install of your system.

Linux:

Download tar binary package based on your Linux system into a desired directory. Untar the package for JDK to be installed using:

```
tar -xvf jdk-xxversion-linux-xnn.tar.gz
```

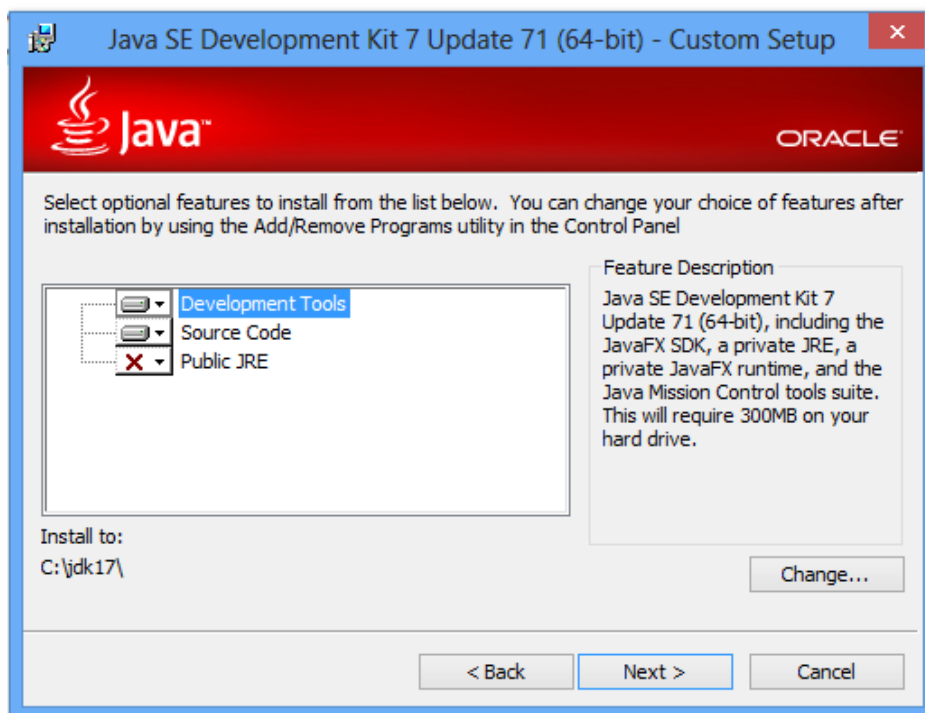
To install the JDK using the RPM binary file, download the file by accepting the license agreement. Install the package using:

```
rpm -ivh jdk-xxversion-linux-xnn.rpm
```

Please note that RPM files need to be installed as root user.

Windows:

After downloading, run the executable. If the target workstation has an existing java installation that you wish to maintain, unselect the “public JRE” option in the installer as shown below.



2. Apache Tomcat

EMERSE is packaged as a java based WAR file that will be deployed to the Tomcat web server. We have not tested EMERSE with other application servers, but it should run on other servers that support the J2EE servlet specification.

Version: Tomcat 7+

Download site: <http://tomcat.apache.org/download-70.cgi>

Windows:

Download the zip file from binary distributions under the “Core” section. Unzip this file in a desired directory. This will become the Tomcat installed directory. Edit the startup.bat file found under bin to point to the directory where the JDK was installed.

```
set JAVA_HOME=c:\path_to\jdk_install
```

Linux:

From the binary distributions listed on the page, choose the “Core” tar file. Move the tar file to a desired directory for installation. Extract the tar file using

```
tar zxvf apache-tomcat-7.0.nn.tar.gz
```

Edit the startup script (startup.sh) found at /path/to/tomcat/bin to point to Java installation directory by adding

```
export JAVA_HOME=/path/to/jdk_install
```

You can also modify JVM settings to use more RAM in the startup script using

```
export JAVA_OPTS="-XX:MaxPermSize=256m -XX:PermSize=128m -Xmx3072m -Xms1024m"
```

It has also been observed that Solr requires a higher limit on the number of open files to be higher than what is typically the default. This can be set using the ulimit command as below.

```
ulimit -v unlimited
```


The exception seen without the higher limit is below:

```
Caused by: java.io.IOException: Map failed
    at sun.nio.ch.FileChannelImpl.map(FileChannelImpl.java:849)
    at org.apache.lucene.store.MMapDirectory.map(MMapDirectory.java:283)
    at org.apache.lucene.store.MMapDirectory$MMapIndexInput.<init>(MMapDirectory.java:228)
    at org.apache.lucene.store.MMapDirectory.openInput(MMapDirectory.java:195)
    at
org.apache.lucene.codecs.compressing.CompressingTermVectorsReader.<init>(CompressingTermVectorsReader.java:118)
    at
org.apache.lucene.codecs.compressing.CompressingTermVectorsFormat.vectorsReader(CompressingTermVectorsFormat.java:85)
    at org.apache.lucene.index.SegmentCoreReaders.<init>(SegmentCoreReaders.java:132)
    at org.apache.lucene.index.SegmentReader.<init>(SegmentReader.java:96)
    at org.apache.lucene.index.StandardDirectoryReader$1.doBody(StandardDirectoryReader.java:63)
    at org.apache.lucene.index.SegmentInfos$FindSegmentsFile.run(SegmentInfos.java:843)
    at org.apache.lucene.index.StandardDirectoryReader.open(StandardDirectoryReader.java:53)
    at org.apache.lucene.index.DirectoryReader.open(DirectoryReader.java:66)
```

Start/Stop:

To start the server use:

```
/path/to/tomcat/bin/startup.sh
```

To stop the server use:

```
/path/to/tomcat/bin/shutdown.sh
```

3. Apache ActiveMQ

EMERSE search requires ActiveMQ for parallel processing of search results from Lucene indexes. Whenever EMERSE is running, ActiveMQ needs to be running in the background.

Version: ActiveMQ 5+

Download site: <http://activemq.apache.org/download.html>

Installation and Configuration:

Download and unzip binary distribution of ActiveMQ. After unpacking, edit the `activemq.bat` (windows) or `activemq.sh` (Linux) file to point to the directory where the JDK was installed.

Windows:

```
set JAVA_HOME=c:\path_to\jdk_install
```

Start the ActiveMQ broker by opening a terminal/command prompt and navigating to its “bin” directory. Type

```
activemq.bat start
```

Linux:

```
export JAVA_HOME=/path/to/jdk_install
```

To start the ActiveMQ broker :

```
cd /path/to/activemq_install/bin
```

```
./activemq start
```

Default port is 8161. Verify it is running by pointing a web browser to:

<http://hostname:8161/admin>

4. EMERSE Web Archive File (WAR file) deployment and configuration

Database Initialization

Provided with the distribution are a set of files, each containing SQL statements that create all needed database objects and sample data that will allow the EMERSE application to startup with a default set of database objects, and sample data in the patients, research studies, synonyms and tables. These files need to be executed in a SQL query tool in the following order:

1. create.sql
2. auditTables.sql
3. sqlToPutBackInModel.sql
4. synonymsCreate.sql
5. lookupData.sql
6. patientData.sql
7. indexData.sql
8. synonymIndexData.sql
9. synonymsData.sql

Index setup

Copy supplied files to a directory on the application server. The path to the default directory is

```
/app/indexes
```

(See the next section if you would like to change this default path)

EMERSE Deployment and Configuration

The next step in getting EMERSE up and running after initial installation of the application server and configuration of the database with default settings is to deploy the EMERSE WAR file. To deploy the file, first rename the supplied war file to `emerse.war`, then copy the war file to the `webapps` directory of the Tomcat server. If Tomcat is using default settings, the WAR file will be exploded into a number of files in a directory called `emerse`. This directory includes all the files needed to run the application. We need to make a change to the settings file to reflect the database that will be used. Inside `WEB-INF/classes` directory of the exploded war file, you will find a file called `project.properties`. This file contains the settings to connect to the database. Update the following values as appropriate for your Oracle database.

For example:

```
ds.username=emerse
ds.password=emersepassword
ds.url=jdbc:oracle:thin:@myhost.med.umich.edu:1521:hostname
ds.driver=oracle.jdbc.driver.OracleDriver
ds.maxPoolSize=4
```

To change path of the indexes directory, update the file `springPostprocessor.properties` found in `WEB-INF/classes` directory.

For example:

```
luceneSearcher.indexPath=/mydir/indexes
```

Once the file is saved, the application server will need to be restarted to reload the configuration to use the latest changes.

Running EMERSE

At this point EMERSE should be up and running. You can verify by pointing a browser to:

<http://hostname:port/emerse>

When the login screen appears provide the following credentials:

Demo user: emerse

Password: demouser

Next Steps

Please see the accompanying guide that provides information on how to integrate your institution's data into EMERSE. It provides tips on indexing data, and information on loading tables that are unique to your organization.

EMERSE Data and System Customization Guide

EMERSE Data

EMERSE stores its internal system data within an Oracle database. If necessary, it is possible to change the database to an open source one, although we do not recommend it at this time due to the effort it would take.

For the purposes of getting started, Oracle makes available a free “express edition” that is fully functional. This free edition of Oracle supports 1 core and up to 10 GB of disk space, which should be enough to support a few users in a demonstration version, or even for a low-powered production version.

<http://www.oracle.com/technetwork/database/database-technologies/express-edition/overview/index.html>

The primary data stored within this database includes a patient demographics table, audit logs, and user data including default settings for each user. These are described below. The large data stores for the documents and document indices are not stored within this database. Instead these are managed by Lucene in its own data store.

Patient Demographics Data

The EMERSE schema includes a patient table with medical record number (MRN), name, date of birth, and other demographic information which is displayed in the search results. Note that the demographics in this table reflect the same as those found within the i2b2 Workbench. In addition to displaying the name, the table is used to validate user-entered MRNs and to calculate current ages of the patients. It should be noted that currently only MRN, name, and `birth_date` are required and used by the system, and thus the other elements are not required.

PATIENT Table Details

Column name	Description	Required or Optional
<code>id</code>	Primary Key	Required
<code>external_id</code>	Medical Record Number	Required
<code>first_name</code>	First Name	Required
<code>middle_name</code>	Middle Name	Optional
<code>last_name</code>	Last Name	Required
<code>birth_date</code>	Birth Date -- used to calculate current age	Required
<code>sex_cd</code>	Sex	Optional
<code>language_cd</code>	Language	Optional
<code>race_cd</code>	Race	Optional
<code>marital_status_cd</code>	Marital Status	Optional
<code>religion_cd</code>	Religion	Optional
<code>zip_cd</code>	ZIP code	Optional

Research Studies Data

EMERSE is often used to aid in research studies. There is a provision for users conducting research to specify their study IRB number at the attestation page after successfully logging in. This enables EMERSE to link that particular session with the study. Loading your institutional IRB data into the `RESEARCH_STUDY` table will enable EMERSE to validate a user's access against the research studies. EMERSE checks to ensure that the study number is valid, that the study's expiration date is not earlier (older) than the current date, and the current study status (since only certain study statuses allow access). At the University of Michigan we have been using a commercial IRB tracking system, and we extract a subset of data from that system to bring into EMERSE for validation of users to studies. Examples of study statuses at Michigan include "Approved", "Terminated", "Pre Submission", "Expired", "Changes required by core staff", etc. Note: to use this feature you may have to customize the list of study statuses to match your local needs.

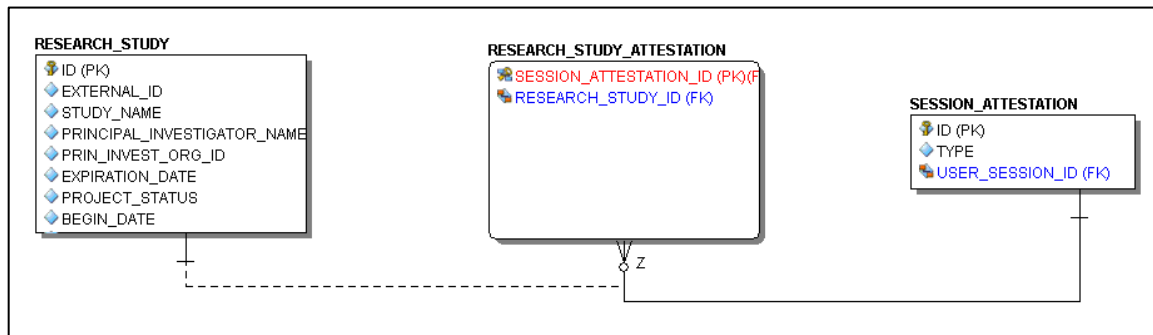


Figure. Entity relationship diagram of some tables related to capturing login attestations.

RESEARCH_STUDY Table Details

Column name	Description	Required or Optional
id	Primary Key	Required
external_id	IRB study number -- used to link specific studies to usage, and is very helpful for tracking research usage	Required
principal_investigator_name	Name of the principal investigator.	Required
prin_invest_org_id	id of principal investigator. Not currently used by EMERSE.	Optional
expiration_date	Expiration date of study. Used to determine if a user should be allowed to proceed. If the expiration date is older than the current date, access should not be granted.	Required
project_status	Current project status. This is used by the IRB system to track where a study is in the review and approval process. Only certain study statuses allow access to EMERSE for research. The valid statuses are defined in the <code>VALID_RES_STUDY_STATUS</code> table	Required
begin_date	The date the study is allowed to begin. This may also be an approval date. Currently EMERSE does not use this date since being able to start the study is better captured in the status that is obtained from our electronic IRB tracking system.	Optional

SESSION_ATTESTATION Table Details

Column name	Description	Required or Optional
id	Primary Key	Required
type	This is the type of usage as defined by the UI of EMERSE, where a user selects their reason for using EMERSE. For example, research, decedent-only study, other. This is recorded for tracking purposes.	Required

VALID_RES_STUDY_STATUS Table Example

The following statuses are currently loaded into the *VALID_RES_STUDY_STATUS* table. These research study valid statuses are loaded up in the build script. These are unique to the University of Michigan (that is, they were developed locally and are implemented in our separate electronic IRB tracking system) and other implementations would have to have their own set of valid statuses if these were to be used to validate and approve usage for research.

STATUS
Exempt Approved - Initial
Approved
Not Regulated
Exempt Approved – Transitional

For non-research attestations, there is a lookup table called *OTHER_ATTESTATION_REASON* that lists available options:

USER_KEY	DESCRIPTION	DELETED_FLAG
FRETXT	Free Text Reason	0
RVPREPRES	Review Preparatory to Research	0
STDYDESC	Study involving only decedents (deceased patients)	0

The free text reason that users entered is stored in a table called *ATTESTATION_OTHER*. An example of what the table would look like with real data is:

SESSION_ATTESTATION_ID	FREE_TEXT_REASON	OTHER_ATTEST_REASON_KEY
50208	Testing out the system	FRETXT
52060	Testing out the system	FRETXT
46051	Looking up a patient in clinic	FRETXT
71052	infection control monitoring	FRETXT
74107	cancer registry operational work	FRETXT

Clinical Documents

EMERSE search is enabled by the indexing of clinical text documents by Apache Lucene. Documents in a clinical environment can come from a myriad of sources like transcription, Radiology, and Pathology, or from an electronic health record. Normally the structure, data, and metadata related to these documents from different sources varies considerably. Search results for each source are displayed in a separate tab in the UI.

Documents to Database mapping

Document Index: Each document source is a row in the `document_index` table. The EMERSE application searches and displays the results based on document source. Document sources normally differ in their format and metadata depending on the source of origin.

DOCUMENT_INDEX Table Details

Column name	Description	Required or Optional
<code>lucene_name</code>	the name of the document source. This field needs to be unique and search results are displayed on separate tabs for each source.	Required
<code>user_description</code>	is the description for the source of document. This field is used when printing individual documents obtained from a search, so that the document source is printed on the document.	Required
<code>compound_key_flag</code>	is a boolean flag that is set when a unique document identifier is a combination of more than one field in the document. This is needed if there is no other unique identifier for the documents. 0 = not needed; 1 = needed In other words, if the primary key of a document source consists of more than one field, this flag needs to be set to true. This flag is used internally to set document viewed flag when a user views a particular document. Regarding compound keys, EMERSE just uses the <code>RPT_ID</code> field when searching/fetching the documents from Lucene. The only time this compound flag is used is to track if the document has been viewed by a particular user. The keys are split up and stored in the <code>document_view</code> table to easily reference back to the source of documents if need be for auditing purposes. Lucene indexing requires that each document has a unique identifier. So, when indexing, the fields that uniquely identify a document need to be concatenated and indexed as <code>RPT_ID</code> . For example, when we indexed Radiology documents we used a combination of document Id and exam description to uniquely identify documents. These fields are concatenated using ' ' and then indexed as <code>RPT_ID</code> .	Required
<code>default_sort_column</code>	is the document field to use to sort the search results for each patient	Required
<code>display_name</code>	is the name displayed in the UI	Required
<code>display_prefix</code>	is the prefix used by UI components. This can be anything, but each source must have a unique <code>display_prefix</code> .	Required
<code>display_order</code>	is the order in which sources appear in the result summary tabs	Required

Shown here is a table with sample `document_index` table data containing with three different document sources:

<code>lucene_name</code>	<code>user_description</code>	<code>compound_key_flag</code>	<code>default_sort_column</code>	<code>display_name</code>	<code>display_prefix</code>	<code>display_order</code>
DMI	Central transcription document	0	Case Date	CareWeb	dmi	0
Radiology	Radiology Documents	0	Report Date	Radiology	rad	1
Pathology	Pathology Document	0	Last Updated	Pathology	path	2

EMR Intent

The `doc_field_emr_intent` lookup table describes the intent of the document fields. This table does not need to be edited by anyone and is used internally by the system to help map various sources and types of data to the intended uses of those data by the system. The table has a name and a description field. The values contained in the name field of this table are listed below. Note that the first 6 items are required for the Lucene indexer to work, the next two are optional, and the final one is no longer used.

Column name	Description	Required or Optional
MRN	patient medical record number, which is a unique patient identifier	Required
RPT_ID	Unique document identifier	Required
CLINICAL_DATE	Date when the clinical event occurred. Often this would be considered the "note date"	Required
LAST_UPDATED	Date when the document was last updated, since changes are sometimes made to documents	Required
RPT_TEXT	The actual text of the clinical document. This field is used by Lucene for lower-case indexing (case-insensitive searching).	Required
RPT_TEXT_NOIC	A copy of the document text to be indexed using a case-sensitive Lucene filter (NOIC = NO Ignore Case)	Required
TEXT	Any generic text field. Note that a document may have multiple of these types of generic text fields (e.g., clinical service, document type, clinician name, etc). This is useful when additional metadata are associated with the document and should be displayed.	Optional
DATE	Any generic date field, since a document may have more than one kind of date associated with it.	Optional
ENCOUNTER_ID	This is no longer used. It had been used for a time to search across all patients without limiting it to a set of medical record numbers. These kind of searches are now done by directly using the Solr API.	No longer used

The actual `doc_field_emr_intent` lookup table is as follows:

NAME	DESCRIPTION
MRN	Patient unique identifier
RPT_ID	Report unique identifier
CLINICAL_DATE	Relevant clinical date of this document
LAST_UPDATED	Last updated date field
RPT_TEXT	Report text case ignored
RPT_TEXT_NOIC	Report text
TEXT	Text field
DATE	Date field
ENCOUNTER_ID	Associated Encounter Id for this doc

Document Fields

All fields in from a document source that are available for display in EMERSE are enumerated in the `document_fields` table. This table consists of a column called `EMR_INTENT` that is linked to the name field of the `doc_field_emr_intent` mapping table. The column `DOC_INDEX_LUC_NAME` is linked to the `lucene_name` field of the `document_index` table.

For a search to be performed, it is necessary that each document source contains six fields defined within the `EMR_INTENT` column that gets linked to the six required fields as defined in the `doc_field_emr_intent` table. Additional metadata fields belonging to a document source can be specified using the generic `EMR_INTENT` options of `TEXT` or `DATE`. These additional metadata fields are used by EMERSE for display in the UI but are not used by Lucene.

Shown below is an example `document_fields` table for three different document sources:

LUCENE_NAME	DATA TYPE	DISPLAY_ORDER	DISPLAY_NAME	EMR_INTENT	DOC_INDEX_LUC_NAME	DISPLAY_FLAG	SUMMARY_DISPLAY_FLAG
MRN	Text	0	MRN	MRN	DMI	0	0
RPT_TEXT	Text	1	Report Text	RPT_TEXT	DMI	0	0
RPT_TEXT_NOIC	Text	2	Report Text	RPT_TEXT_NOIC	DMI	0	0
RPT_ID	Text	3	Report ID	RPT_ID	DMI	1	1
LAST_UPDATED	Date	4	Last Updated	LAST_UPDATED	DMI	1	0
CASE_DATE	Date	5	Case Date	CLINICAL_DATE	DMI	1	1
MRN	Text	0	MRN	MRN	PATHOLOGY	0	0
RPT_TEXT	Text	1	Report Text	RPT_TEXT	PATHOLOGY	0	0
RPT_TEXT_NOIC	Text	2	Report Text	RPT_TEXT_NOIC	PATHOLOGY	0	0
RPT_ID	Text	3	Report Id	RPT_ID	PATHOLOGY	1	1
LAST_UPDATED	Date	4	Last Updated	LAST_UPDATED	PATHOLOGY	1	1
DR_NUM	Text	5	Doctor Num	TEXT	PATHOLOGY	1	1
COLLECTION_DATE	Date	6	Collection Date	CLINICAL_DATE	PATHOLOGY	1	0
MRN	Text	0	MRN	MRN	RADIOLOGY	0	0
RPT_TEXT	Text	1	Report Text	RPT_TEXT	RADIOLOGY	0	0
RPT_TEXT_NOIC	Text	2	Report Text	RPT_TEXT_NOIC	RADIOLOGY	0	0
RPT_ID	Text	3	Report ID	RPT_ID	RADIOLOGY	1	1
LAST_UPDATED	Date	4	Last Updated	LAST_UPDATED	RADIOLOGY	1	0
SVC_CD	Text	5	Service Code	TEXT	RADIOLOGY	1	0
DR_NUM	Text	6	Doctor Num	TEXT	RADIOLOGY	1	0
RPT_DATE	Date	7	Report Date	CLINICAL_DATE	RADIOLOGY	1	1

A description of the columns names is as follows:

Column name	Description	Required or Optional
LUCENE_NAME	name of the document field that was indexed in Lucene. The names of the fields are specified in <code>schema.xml</code> file for Lucene to index documents and need to match in this table column	Required
DATATYPE	type of the document field	Required
DISPLAY_ORDER	order in which fields need to appear in the search results	Required
DISPLAY_NAME	name that appears in the UI	Required
EMR_INTENT	specifies the intent of the field. This refers to the fields defined in the <code>doc_field_emr_intent</code> table.	Required
DOC_INDEX_LUC_NAME	specifies the document type key from <code>document_index</code> table	Required
DISPLAY_FLAG	flag that controls if the field is displayed when document is displayed	Required
SUMMARY_DISPLAY_FLAG	flag that controls if the field is displayed in search results summary page	Required

Lucene Shards

A shard is an instance of a Lucene index. A large index can be broken into multiple pieces (shards) to make searching more efficient. All lucene indexes are listed in `lucene_shards` table. Each installation of EMERSE can have its own way of breaking up the indices, but we have found that it makes the most sense to have each document source in its own index, and larges indices are divided into multiple shards. In general we have created shards based on document sources as well as date ranges. For example, each year of Epic EHR documents is in its own Shard to keep their sizes 'manageable'. There is no specific or ideal size for a shard, but it likely depends on the available power and speed of your server hardware, the size of the documents, etc. Our shards our typically about 200 GB each.

EMERSE uses Lucene based indexes to search, highlight, and display documents. Depending on the number of documents each document type contains it is possible to split the documents into multiple Lucene indexes (shards or sub-indexes) for easy maintenance. Typically, splitting is done based on date ranges. This information is captured in a database table called `lucene_shards`

LUCENE_SHARDS Table Details

Column name	Description	Required or Optional
ID	The Lucene name of the index	Required
PARENT_DOC_INDEX	Specifies the document type key from <code>document_index</code> table	Required
START_DATETIME	Start date of clinical documents in this shard	Required
END_DATETIME	End date of clinical documents in this shard	Required

A sample `LUCENE_SHARDS` table is shown below

ID	PARENT_DOC_INDEX	START_DATETIME	END_DATETIME
rpta	DMI	01.02.2008 00:00:00	31.12.2011 00:00:00
rptb	RADIOLOGY	01.01.2011 00:00:00	31.12.2011 00:00:00
rptc	PATHOLOGY	01.01.2011 00:00:00	31.12.2012 00:00:00

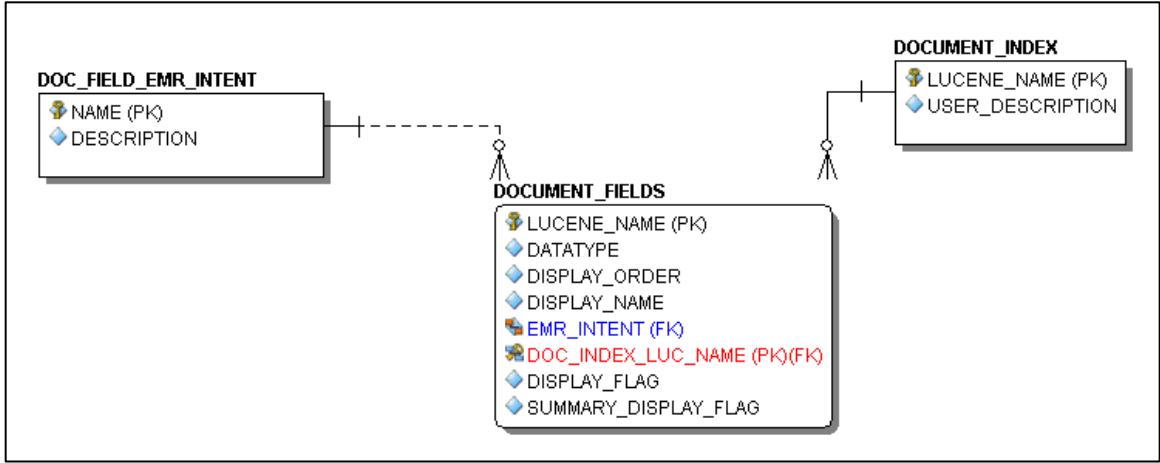


Figure. Entity relationship diagram showing how the three tables above are related.

Indexing documents with Apache Solr

EMERSE requires that all text documents be indexed in the Lucene format. Indexing can be done in several ways. We use Apache Solr which is a standalone full text server built using Lucene to index documents. Solr provide a lot of functionality on top of the Lucene libraries. In the case of indexing Solr provides:

1. A REST-based API
2. Remote Java access via SolrJ

When using the REST API, the JSON/XML/CSV type of data structures are posted via HTTP calls to the Solr server. We use SolrJ, where binary java objects are transmitted, so mapping to these data structures is unnecessary.

A simplified code example of indexing with SolrJ can be found here:

<http://www.solrtutorial.com/solrj-tutorial.html>

Please refer to the installation section of this manual for Solr install instructions. It is assumed that all documents that need to be indexed are available in a data store with at least the four pieces of metadata information required for each document: MRN, RPT_ID, CLINICAL_DATE, and LAST_UPDATED along with clinical document text.

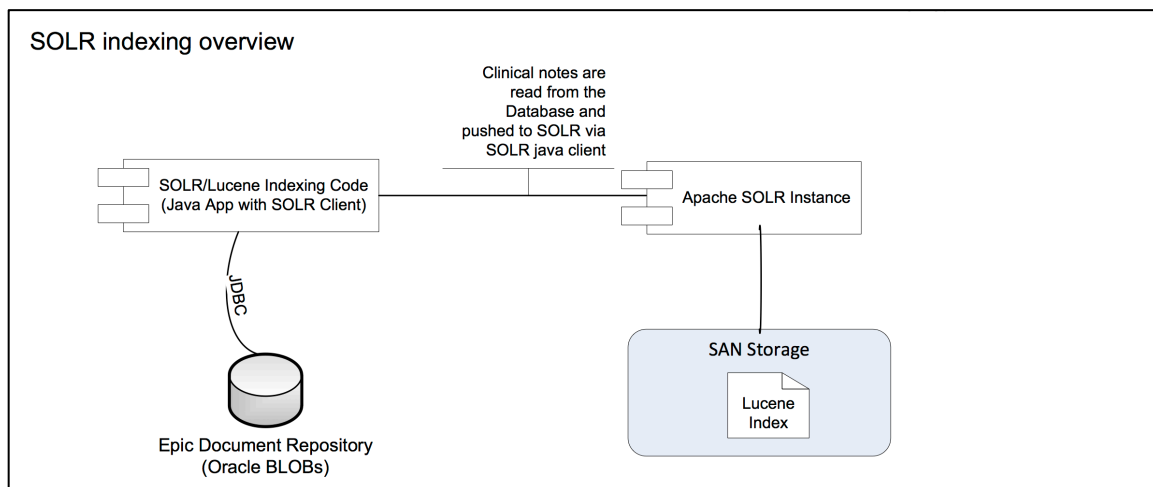


Figure. High-level overview of indexing using Apache Solr.

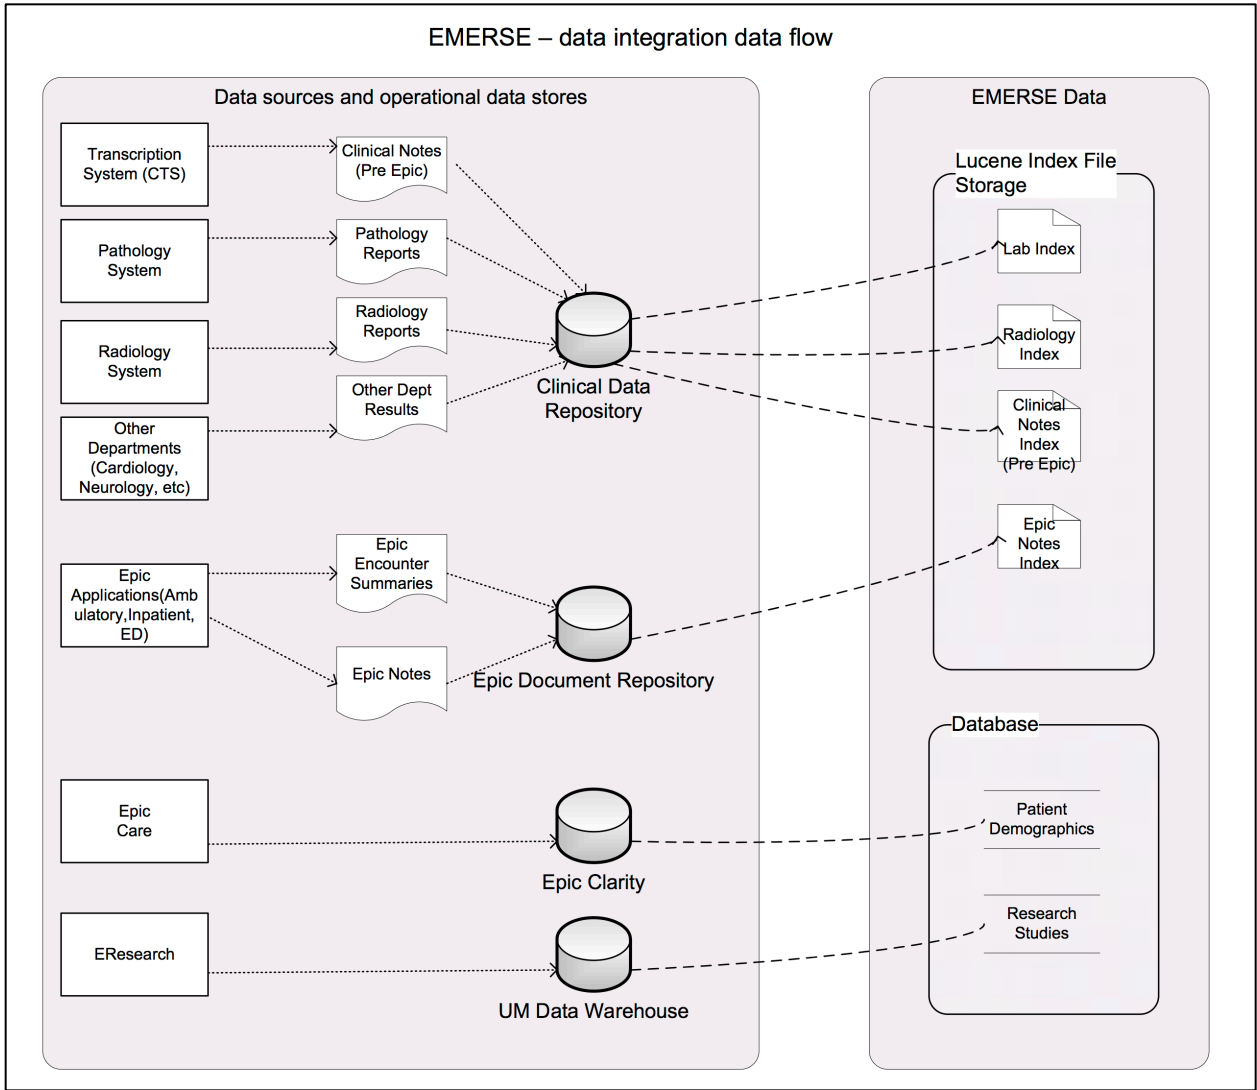


Figure. High-level overview of how data are integrated from various source systems into EMERSE at the University of Michigan. Each institution will likely have its own unique set of sources and integration issues.

Indexing using Solr DIH (Data Import Handler)

Reference: <https://wiki.apache.org/solr/DataImportHandler>

The easiest way to index documents stored in your local document store database is to use the DataImportHandler (DIH) that ships with Solr. Solr at startup needs a home variable (`solr.home`) defined where it would look for the various configuration files. DIH requires the following two jars to be placed in the lib directory of Solr installation:

1. `solr-dataimporthandler-4.x.x.jar`
2. `solr-dataimporthandler-extras-4.x.x.jar`

Database configuration information and the queries needed to retrieve documents are specified in DIH `dataconfig.xml` file. A sample configuration with mandatory fields required by EMERSE is shown below:

```
<dataConfig>
  <dataSource name="XE" driver="oracle.jdbc.driver.OracleDriver" url="jdbc:oracle:thin:@localhost:1521:XE"
    user="system" password="abc" />
  <document name="rpta">
    <entity name="rptsa" pk="RPT_ID" query="select rpt_id,mrn,cast(rpt_date as date) rpt_date,rpt_text,
      cast(clinic_date as date) clinic_date from emerse.reports_ctsa where rpt_type=1"
      transformer="ClobTransformer,DateFormatTransformer">
      <field column="RPT_ID" name="RPT_ID" />
      <field column="RPT_TEXT" name="RPT_TEXT" clob="true"/>
      <field column="MRN" name="MRN" />
      <field column="RPT_DATE" name="RPT_DATE" dateTimeFormat="yyyy-MM-dd HH:mm:ss.S" locale="en" />
      <field column="CLINIC_DATE" name="CLINIC_DATE" dateTimeFormat="yyyy-MM-dd HH:mm:ss.S" locale="en"/>
    </entity>
  </document>
</dataConfig>
```

Mapping of the database columns to the Lucene format and how they need to be indexed is specified in `schema.xml`. A snippet of `schema.xml` for the fields specified above in the `dataconfig.xml` above is shown here:

```
<field type="string" name="RPT_ID" indexed="true" stored="true" />
<field type="string" name="MRN" indexed="true" stored="true" />
<field type="date" name="CLINIC_DATE" indexed="true" stored="true" />
<field type="date" name="RPT_DATE" indexed="true" stored="true" />
<field type="text_general_nostopwords" name="RPT_TEXT" indexed="true" stored="true" termVectors="true"
  termPositions="true" termOffsets="true" />
<field name="RPT_TEXT_NOIC" type="text_general_nostopwords_nolowcase" indexed="true" stored="true"
  termVectors="true" termPositions="true" termOffsets="true" />
<field name="_version_" type="long" indexed="true" stored="true"/>
```

A few things to note:

1. There are multiple ways to create indices depending on which 'analyzer' is used to tokenize the text. Tokenization refers, in part, to the process of how text should be broken up into individual words, and considers properties such as hyphens between words.
2. The Lucene field `RPT_TEXT_NOIC` does not exist in the database query output. The `Copyfield` command of Solr is utilized to make a copy using `RPT_TEXT`. The only difference between these two fields is that text in `RPT_TEXT_NOIC` is tokenized and indexed 'as is' without applying a lowercase filter `<copyField source="RPT_TEXT" dest="RPT_TEXT_NOIC" />`
3. The field type `text_general_nostopwords` is an extension of Solr `text_general` wherein index and query analyzer does not use any stopwords filter
4. The field type `text_general_nostopwords_nolowcase` is an extension of `text_general_nostopwords` where lowercase filter is skipped in addition to stopwords

Indexing programmatically with SolrJ

Solrj is a Java client that can be used to add/update a Solr index. This may provide better data indexing speed because multiple threads can be used to load the data.

Reference: http://wiki.apache.org/solr/Solrj#Adding_Data_to_Solr

Solr connects to the document repository using SQL/JDBC. We are using a Unix Cron job to tell the indexing code to start and run once per night via an HTTP call to the running app.

Interactions with an EHR like Epic

The University of Michigan Health System (UMHS) EHR is based on Epic. EMERSE receives encounter summaries and notes from Epic via an Epic Bridges outbound HL7 interface. Details about this interface can be found here:

<http://open.epic.com/Content/specs/OutgoingDocumentationInterfaceTechnicalSpecification.pdf>

The emitted HL7 messages include the clinical notes in RTF format, which preserves the formatting and structure of the notes. These files are converted to HTML by ASPOSE Words API, and stored for nightly indexing into Lucene format.

Reference: <http://www.aspose.com/java/word-component.aspx>

EMERSE is designed to run stand alone, so there are no real-time Web Service calls to Epic. It is also possible to receive plain text notes either through HL7 or through Epic Clarity database, however, the display of the notes will not look as good as an RTF version since the original document formatting may be essential for understanding/interpreting the document once displayed within EMERSE. Epic Clarity currently does not preserve the formatting, line breaks, etc. but could be considered a source of documents if that were to change.

If there is a desire to use Clarity to obtain notes, we suggest looking into the following tables:

HNO_INFO	(note metadata)
HNO_ENC_INFO	(notes /encounter linkage)
HNO_NOTE_TEXT	(actual note text)

These tables don't seem to be documented in Epic's Clarity ambulatory care documentation, even though outpatient notes are stored in the same tables along with the inpatient notes as of Epic 2012. However, The inpatient Clarity "training companion" does have an overview of these tables, and they are mentioned in the Clarity data dictionary, so between these two references you should be able to build a query that collects Epic notes.

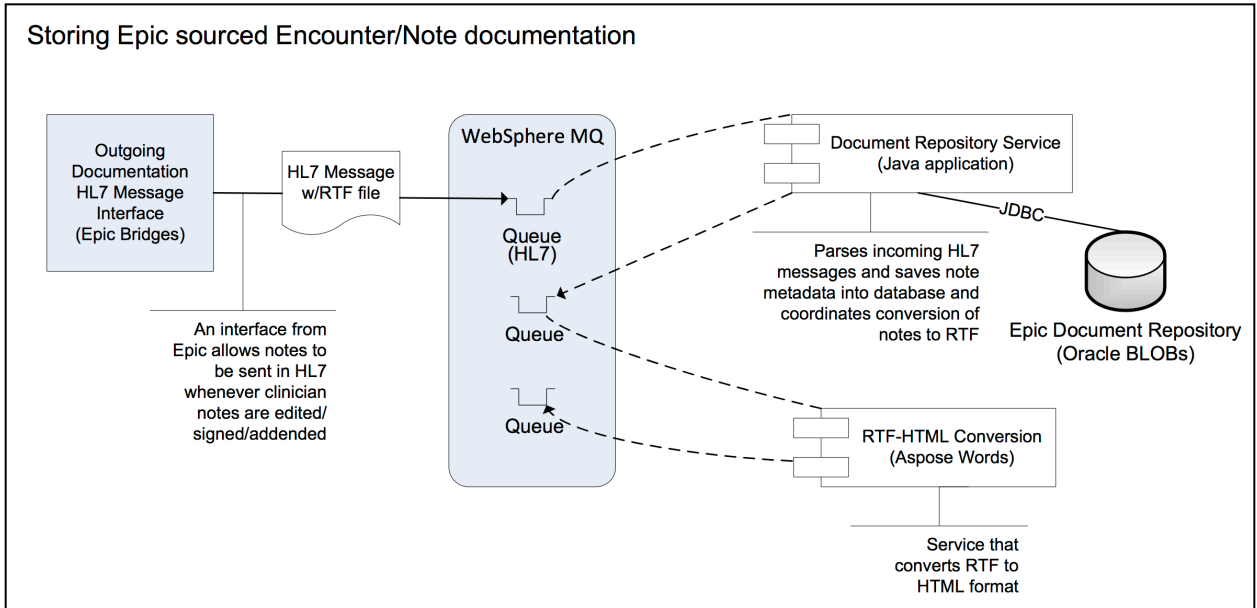


Figure. The approach used at the University of Michigan to obtain notes from the Epic EHR and make them searchable within EMERSE.

Known Bugs and Issues

Because EMERSE relies on other open source software, we have identified bugs that we are not able to readily fix. These currently include:

[March 2015]

The 'Advanced Search' feature sends a standard Lucene query without any interpretation or modification by EMERSE. We have found that for complex Boolean queries the system does retrieve the documents correctly, but the separate Highlighter component does not properly highlight based on the Boolean query. Rather it will highlight any terms it finds in the query regardless of the Boolean operation.

The screen shot below demonstrates this. The document was retrieved because 'hypertension' and 'diabetes' were found together in the document. Note that the term 'cardiac' does not appear in the document, but 'back pain' was still highlighted.

The screenshot displays the EMERSE search interface. At the top, the search terms are "(hypertension AND diabetes) OR (cardiac AND \"back pain\")". Below the search bar, there are buttons for "Search" and "Overview", along with navigation arrows and the text "Patient 1 of 11". A horizontal menu below the search bar includes "MiChart", "Careweb", "Radiology", "Pathology", "Other", and "Imaged Doc". The main content area shows a "Summary" section with the following text: "...a history of chronic back pain and migraines who presents... ...shortness of breath, or back pain. She denies any nasal... ...letal: Negative for back pain and gait problem. ... Back pain, chronic Diabetes Hypertension ...a history of chronic back pain and migraines who presents...". The terms "back pain", "Diabetes", and "Hypertension" are highlighted in yellow.