

Regional Healthcare Information Systems Architecture and Technology¹

William E. Johnston

Lawrence Berkeley Laboratory

Berkeley, CA 94720

wejohnston@lbl.gov, tel: 510-486-5014, fax: 510-486-6363

Technical Report: LBL-34770

1.0 Background

The National Information Infrastructure (NII) is seen as a key element for advancing the competitive position of our country in the 21st century. The NII is to be a shared infrastructure (network + services + access), together with a large collection of interoperating applications and a wealth of on-line data. The network provides the transport of everything from healthcare systems to manufacturing control information to multimedia tele-seminar activities in K-12 science and humanities education. A large base of interoperating applications supply the mechanisms for organizing, presenting, and analyzing data in support of increasingly complex social and technical activities. For example, success in advanced manufacturing can be laid to the increased flexibility and efficiency in organizing a complex environment through communication among information analysis systems.² A survey of abstracts from the Excerpta Medica database seem to indicate that similar flexibility and efficiency will be gained by the healthcare sector as it incorporates easily used, interworking information systems. The flexibility will lead to better access to healthcare (from, for example, wide spread and effective tele-medicine), and the efficiency will lead to lower costs and therefore more broad based socio-economic access.

The network infrastructure of the NII provides a common way to communicate among computer information systems. This common communication is provided by the widespread use of the ARPA Internet protocols in order to homogenize the underlying telecommunications networks, and the associated numerous, widely deployed supporting applications that provide human communication, directory services, and a range of information resources. Progress in the ability of independent applications to share data has been promoted by software architecture models like the ISO, Open Systems Interconnect (OSI). These models provide us the means to analyze and structure software architectures to identify, provide, and exploit points of common function³. It is these points of common function, together with standards for data and message formats, that allow us to build large scale distributed systems, and to standardize and reuse software components. The OSI model gave rise (directly or indirectly) to many of the data exchange standards (MAP / TOP, EDI, ODA, RDB, etc.)⁴ needed for information exchange, and the ARPA Internet protocols gave rise to the National Education and Research Network (NREN).⁵ For regional and national scale systems, this software and communications infrastructure has the added

1- Some of the ideas for the types of services and information flows in an integrated metropolitan healthcare information system were developed during discussions with the Kansas City STAR group. This group includes regional healthcare providers, Univ. of Kansas Medical Center, Sprint Corp., Ruff Corp., as well as others. A KC-STAR contact is Kim Galbraith of Sprint, 916-967-2158, fax: 916-967-2526.

2- Carnoy, Castells, Cohen, and Cardoso, *The New Global Economy in the Information Age* (Pennsylvania State University Press, 1993)

3- Jain and Agrawala, *Open Systems Interconnection: Its Architectures and Protocols* (McGraw-Hill, 1993)

4- Acronyms mean: Manufacturing Automation Protocol / Technical and Office Protocol; Electronic Data Interchange; Office Document Architecture; Remote Database Access

5- The NREN is a combination of the NSFNet, DOE's ESNet, the regional networks, and the thousands of "campus" (site area) networks, and in the past few years numerous purely commercial off-shoots of these networks.

W. E. Johnston, Lawrence Berkeley Laboratory

importance of providing us with the mechanisms to build large scale, scalable, fault tolerant systems for applications like healthcare information systems.

It is estimated that there are in excess of 15 million users with connections to the Internet, doing everything from scientific experiment control, to medical PACS systems, to automated manufacturing, to worldwide video teleconferencing. For most of these users, the issue is not how many people are connected to the Internet, but how much information is available on the Internet. For regional healthcare information systems, participating in the NII / NREN will mean access to data bases like those at the National Library of Medicine, genomic databases scattered around the world, e-mail and data transfer capability with almost every university in the U. S., and thousands of others all over the world. Recent examples of information availability include the Clinton administration's revised OMB policy requiring Federal agencies to provide network access to all public government information, and an experimental program at the Univ. of California's San Francisco Medical Center (UCSF) library to bring full text versions of all Springer-Verlag medical journals on line for Internet access on campus.

2.0 Regional Healthcare Information Systems Architecture

There is tremendous potential utility and synergy in connecting regional groups of physicians, nurses, medical researchers and students, to NII / NREN. The information infrastructure presented in this document is designed to build on the flexibility and access provided by twenty years of Internet experience and infrastructure building, and to address the specific needs of regional consortia of hospitals and educational institutions. The goal of the Regional Healthcare Information Systems Architecture is to provide significant new capabilities for the regional healthcare environment. For example:

- Homogeneous, universally available patient records
- Multimedia data as part of the CPR, to provide:
 - + audio and video clips from diagnostic and consulting sessions
 - + laboratory results, including instrument "trace" records
 - + X-ray, and other sorts of imagery
- access to lectures, workshops, and source material for on-going professional and patient education
- On-line reference source material such as reference image libraries for X-ray and pathology diagnosis
- Tele-presence type of access for emergency, traveling, outlying, and other remote and / or mobile users of all aspects of the regional facilities

In order to support these capabilities, the types of information addressed in this architecture include the full range of healthcare information: computerized patient records, including written commentary; laboratory test results; images of all types; multimedia material (video and audio); instrumentation records, etc. The architecture is intended to address a wide range of issues associated with the data apart from storage and access: The style of data management, inclusion of legacy systems, the all important issues of user interface, security and authentication, etc. The architecture is intended to accommodate a wide variety of solutions that encompass these issues, but does not specify the solutions. For an analysis of many of the data issues, see "White Paper on Computer-based Patient Record Systems"⁶, which is attached as Appendix A. Also see Velde⁷.

6- Frank Olken and William Johnston, "White Paper on Computer-based Patient Record Systems", Lawrence Berkeley Laboratory technical report, LBL-34303.

7- Rudi Van de Velde, Hospital Information Systems - The Next Generation (Springer-Verlag, 1992)
Carnoy, Castells, Cohen, and Cardoso, *The New Global Economy in the Information Age* (Pennsylvania State University Press, 1993)

The architecture consists of several major elements: human resources, electronic information resources, the users (healthcare practitioners and patients, provider and payor organizations, and researchers), the computing, communications, and data storage infrastructure (information repository). The concept is illustrated in Figure 4: “Healthcare System Architecture - Mobile, Home, and Payor Sites”.

The architecture described here is a completely distributed information infrastructure, with levels of access depending primarily on the volume of data that must be moved among the communicating elements. In other words, everybody from the radiologist at a large hospital to the nurse-practitioner at a rural clinic, to the paramedic in the field, potentially has access to the same information (patient records, x-rays, lab results, etc.), but the bandwidth of access will vary by the circumstance. (Actual access is conditioned on, among other things, the “right” of access; i.e. the authorization of that person to access the information.)

Figure 2: “Healthcare System Architecture - HUB”, Figure 3: “Healthcare System Architecture - Lab, Telemedicine, and Clinic Sites”, and Figure 4: “Healthcare System Architecture - Mobile, Home, and Payor Sites” illustrate the key system elements of the infrastructure, and these elements are discussed below.

2.1 Regional Data Center / The Hub Information Services

The systems and communications infrastructure of regional healthcare information systems will support an variety of data, information, and medical services. Many of these services are characterized in the Hub concept⁸. The Hub is a logically central (though potentially geographically dispersed) facility that is a repository for, and gateway to: medical information in the form of numerous databases; medical consulting, educational material, etc., together with shared computing and data handling services for community healthcare providers. The network provides uniform access to these services, and the issues of physical centralization may be decided on the basis of economic efficiencies (i.e. economies of scale for hardware operation and maintenance), security, reliability, etc. A regional data center (RDC) may be the most efficient way to supply many data handling, communications, and computing facilities, while the Hub acts as an information center. The elements of the architecture shown in the figure provide functionality in the following ways.

(1) Large, secure, data archive functions:

The shared mass storage system (MSS) is a key element for many aspects of the information architecture. This system will provide a relatively large (50 - 100 gigabytes) fast access (cache) storage, backed by a very large (multiple terabyte) tertiary archival storage. The centralization of this facility can provide economy of scale in the hardware, software, and operational aspects of storage. Access control can be done at the application level, with the MSS server being protected by using secure client frontends as the only allowed access to data. This system may be physically dispersed as requirements for archive protection and disaster recovery are formalized. Basically all bulk storage needed to support the hospital consortium will be (administratively) centralized in this facility. Some of the functions supported by this MSS are:

i) Tele-radiology

8- The term “Hub” as applied to an information center that serves as an on-line library and consulting service for the medical community, is due to Dr. Ace Allan, University of Kansas Medical Center.

Probably the largest consumer of bulk storage, the cache and tertiary storage capacity will be sized to accommodate the digital radiology activity of the region.

ii) Multimedia record storage

The DBMS that handles the CPR system will likely not be able to deal efficiently with large audio, video, and instrumentation trace records that are, none-the-less, part of the patient record. These records will be stored on the MSS, and the individual CPR will likely contain pointers to these files on the MSS.

iii) Digital libraries

While many library functions will be provided via resources in the NII / NREN, locally developed material, material accessed with sufficient frequency as to warrant "local" storage, image based teaching materials, etc., may reside on the regional MSS.

iv) Backups

The CPR DBMS servers, as well as departmental computer systems, may choose to use the MSS as the repository for backups of local disks, databases, etc. This strategy can provide a level of reliability not generally possible in small data centers: The storage media can be used and archived in an environmentally controlled environment; storage media can monitored for use degradation, and can be routinely rotated (re-written); etc. All of this is possible in the site data center, but frequently is not done because of the expense. These aspects of the storage system are quite scalable, leading to economies in large operations.

v) CPR tertiary storage

Many DBMSs deal poorly with "house keeping" and organizing data records for archiving. This must be a criteria when selecting a DBMS, and it must be possible to use the MSS for the off-line, but still indexed DBMS storage. (One does not want to have to set up a parallel bulk storage facility just for the DBMS to use.)

(2) Shared Computing Resources:

Given high speed networks to provide access, then as experimental technologies like computer assisted radiation treatment planning, MRI 3D analysis and visualization, etc., move into medical practice, the computational requirements are very likely to exceed those available locally. The RDC could provide compute servers for these functions in a very transparent way. As with the MSS, the arguments for centralized hardware with remote access via high speed networks, are economy of scale in the hardware, software, and operational aspects.

(3) Internet Router:

The Internet is a rich source of information, communication, collaboration, and innovation. It is also a fairly anarchistic environment in which clever students test their programming prowess by various annoying intrusions into operational environments. Mechanisms to limit access from the Internet into organizational environments are now well understood, and many commercial technology organizations impose various restrictions on access to their internal systems. For the most part, well maintained systems that have competent system managers, and no obviously "attractive" characteristics (e.g. being the subject of a movie or novel, or having political significance) are seldom bothered by hackers. It is up to each organization to weigh the benefits vs. the risks of open Internet access. For the most part the technology exists to implement the resulting policy.

In general, the desirable state of affairs is to maximize the information flow and minimize the risks. Healthcare professionals should be able to exchange large data sets, explore and retrieve information from important Internet resources like the National Library of Medicine and the various genetic and protein databases, teaching hospitals around the country and the world, participate in network based teleconferences, etc. At the same time, the systems responsible for the integrity of the CPR and other sensitive data obviously have to have fire-walls build around them. A fairly straight forward strategy for accomplishing open access from workstations, and at the same time protecting sensitive data and facilities, involves establishing a hierarchy of systems that permit varying degrees of free access to and from the Internet. Modern, versatile, high speed, routers provide many of the mechanisms to implement these levels of access.

(4) High Speed (ATM) Multimedia Data Network: The general architecture of the telecommunications system is that a high speed fiber based metropolitan ATM network provides the high speed inter-site data communication for large scale users. Within each large site, there are two primary ATM streams. One serves as the link level transport for the Internet protocol suites that carry the computer data between high speed systems like the MSS and the radiology workstations, of the multimedia file servers and the workstations. This stream is delivered to an Internet router that maps it to the various LAN protocols (Ethernet, FDDI, Appletalk, etc.) that deliver it to individual workstations. (Though in the future, the delivery to workstations may be ATM as well.) The second stream is a multimedia stream for high resolution video and audio. What this stream really does is to provide point-to-multipoint video routing capability. The primary use for this would be high resolution video display (e.g. auditorium projection and other real time, low compression applications). However, it is anticipated that most multimedia communication among workstations, or in the retrieval and display of records, will be done using Internet protocols, or be gatewayed into site LAN environments.

(5) Backup Communications: In addition to the primary access to the metropolitan area network, there would be a backup communication gateway. This gateway would connect to the same Internet router that is the primary gateway, but it would provide an alternate, low bandwidth access to the outside world. This gateway will provide access to non-local elements of the CPR when the primary communication is disabled. The telecommunications service to this backup gateway could be high speed dialup modems (50 kilobits/sec), basic rate ISDN (128 kilobits/sec) or primary rate ISDN (1544 kilobits/sec). The main feature of this facility is the independence of its telecommunications service from the metropolitan ATM network. Even the lowest backup data rates should be able to supply components of the CPR necessary for emergency medical services.

(6) Shared Accounting and Billing:

(7) Security Domain Server: Depending on the level of security, privacy, and authentication required for the various types of information, an independently administered and highly secure system is typically responsible for holding the information used to authenticate users, grant access permission, provide encryption keys, etc.

(8) Multimedia File Server: If libraries of multimedia material have to be delivered to a multiplicity of end users, then a system specifically designed for this purpose will be needed. (This system is not unlike the video file servers being developed for use in video-on-demand

commercial applications.

- (9) If the strategy for the common CPR database turns out to be a few central servers, the Hub can again provide hardware, software, and operational support.
- (10) Network control: The RDC can provide the additional service of network operations center (NOC). By treating the regional consortium as a single networking administrative unit, the day-to-day network operational issues may be relegated to a service provider.

2.2 The Metropolitan Hospital Sites

Apart from the communications services already mentioned, the hospital will have multimedia workstations (those capable of displaying multimedia records, teleconferencing with other remote workstations over the Internet, etc.), the traditional (existing) department systems (PCs and Macs), and perhaps a CPR DBMS server. Hub operational support for these systems might include backup, mail and file servers, etc. A multimedia gateway will allow some level of access to teleconferences and tele-medicine functions from personal computers (most of which currently do not communicate multimedia off of the local area network).

- (1) Internet router
The only difference between this router and the similar one at the RDC is that access to the Internet is through the router at the RDC. This has several advantages in that more network operations can be provided by the RDC, and the telecommunications connection to the Internet at the RDC can serve all of the consortia participants by routing over the ATM network.
- (2) Backup Communications: Same rationale as for the RDC.
- (3) High Speed (ATM) Multimedia Data Network:
- (4) Radiology Workstation: These systems are the primary high resolution digital image display and analysis systems. They are moderately fast systems with direct access to the MSS and the Multimedia File Server. Digital radiology image capture and hardcopy printing will probably be done via these workstations.
- (5) ATM Multimedia Gateway: This is primarily intended for high quality video / audio communications among the metropolitan hospital sites for educational, conferencing, and consulting. It will serve a function similar to commercial video teleconferencing facilities, but in a dedicated manner between the larger regional sites.
- (6) CPR DBMS Server: There are several strategies possible for the primary CPR database. One possible approach is that each regional site has their own database server that is the repository for the patient records of the people who visit that site as their primary source of healthcare. This system would be part of a distributed DBMS that served all of the sites, providing uniform access to patient records, but local physical storage will provide faster and more reliable access for local patients. The DBMS should be able to migrate records to other sites if the patient moves, or to the RDC in the case of inactive patients. In any event, uniform access to a distributed DBMS is the architectural goal. (Note that this does preclude an "administrative" owner of the records being identified. The DBMS should be flexible enough to support a variety of policies, independent of the physical location of the records.)
- (7) Department Systems: There will be many existing systems at the department and individual

W. E. Johnston, Lawrence Berkeley Laboratory

level. The extent to which these systems are integrated into the regional architecture will depend on many cost and political factors. A good deal of integration is possible even with today's PC and Mac systems by incorporating the departmental file server into the regional system.

(8) CPR / Multimedia Workstation: This is the prototype of the primary access portal for the information system. This workstation will request, present, and interact with all aspects of the patient record, including the textual, image, instrumentation trace, video, and audio. This workstation will present formatted information summaries, do auxiliary analysis, and provide access to personal teleconferencing, multimedia resources, etc. There is nothing speculative about the technology of this workstation, several commercial offerings exist now, and most workstations will have all of these capabilities in the near future.

(9) Multimedia Gateway: The purpose of this system is to homogenize the formats of the various multimedia sources so that the existing systems will be able to participate. In the future, as standards for multimedia data are implemented the functions of this system will probably become less important.

2.3 Prototype Remote Laboratory

(1) Communications gateway

(2) Local storage

(3) Local record keeping

(4) Instrumentation systems

2.4 Tele-medicine (Remote) Sites

The basic idea at remote sites is to provide access to metropolitan area services, including two way video and audio, as well as remote instrumentation, in a uniform manner that allows access to the same information streams available in the metro area. The approach to doing this is to define a multimedia workstation that, in a single "box" will provide the communications interface, as well as audio, video, and instrumentation interfaces. This workstation would be a node on the Internet, and thus provide access to the CPR, digital library, etc. Most modern engineering workstations are capable of this level of service integration, the main issue will be the software and user interface design that provides a usable package. This system should flexibly provide all of the tele-medicine functions, as well as serving for tele-education, teleconferencing, etc.

(1) Multimedia workstation, communications gateway, and instrumentation interface

(2) local storage

(3) Sensors

(4) Video and audio I/O Communications gateway

2.5 Clinic

- (1) Communications gateway
- (2) CPR / multimedia workstations
- (3) Radiology workstations
- (4) CPR DBMS server
- (5) clinic systems

2.6 Mobile / Home Environment

- (1) Portable workstation, communications gateway, and sensor interfaces
- (2) Sensors
- (3) local storage
- (4) Image input / output capability

2.7 Payor

- (1) Communications gateway
- (2) Accounting and billing systems

3.0 Background for a Cost Model

These costs are estimated based on typical, commercially available hardware, at listed prices. The model is built up of specific manufacturer's products, but these are intended to represent a typical example of the product. There has not been a serious attempt to build an accurate model based on case studies of things like record size and usage patterns, but the model environment is based on estimates for an organization of average size, and probably just starting up in this environment. For example, the CPR system is sized for a few dozens of users per site, etc. The applications are assumed to be important, but not critical. Figures are for two start-up years for a metropolitan consortium of moderate size.

3.1 Computing and Communications Systems

For the HOST / OSL, NOTE:

- 1) Three situations are considered in the tables.
 - i) A full blown regional trial (essentially the scenario shown in the figures).
 - ii) A single site prototype (e.g. HOST, Open Systems Lab).
 - iii) A multi-site prototype (e.g. distributed OSL). Table column is per site.
(Stared (*) items indicate only one is needed for the distributed OSL)

2) No operational costs are included in this section, however hardware and software maintenance can be estimated at 5-8% of the total system cost per year.

3) Note that except for the MSS software, there is no DBMS or applications software included in this estimate.

TABLE 1. OSL Totals (from tables two and three)

ITEM	Regional Trial per site (\$K)	OSL (single site prototype)	OSL (multi-site prototype)
Total for single OSL site:		871	
Multiple OSL site configuration, one only item subtotal:			495
Multiple OSL site configuration, one-for-each-site subtotal:			186
Multiple OSL site configuration, total for three sites:			1053

3.1.1 The Metropolitan Hospital Sites

TABLE 2. Metropolitan Hospital Site

Item	Regional Trial per site (\$K)	OSL (single site prototype)	OSL (multi-site prototype)
Communications:			
ATM LAN switch / UNI: (The BISDN user-network interface, attached on the inside directly to: radiology workstations; Internet router; ATM multimedia gateway.)	50	=	=
Internet router (IP router with ATM, Ethernet, and FDDI interfaces, together with Appletalk gateway.)	57	57	57
Backup communications access (Router w/ multiple dialup interfaces and modems, and/or ISDN interface.)	18	=	=
Site preparation (for new high speed nets only) (Fiber installation for FDDI ring for CPR workstations and ATM point-to-point for radiology workstations.)	75	=	=
ATM multimedia gateway / codec (Video and audio to ATM, a la the Wash. U. MMX. For high quality multimedia distribution.)	15	=	=
sub totals			
CPR:			
Multimedia workstations (@ \$7.5K) (The basic CPR workstation. Capable of multimedia and instrument record playback)	30	30	15
CPR server (@ \$75K) (Sun SPARCserver 1000, 96 MBy memory, 10 GBy disk on 2 SCSI adaptors (5 GBy mirrored), FDDI, CDRom)	75	150	75
Hardcopy (@ \$4) (Laser printers, one 17 pg/min for each participating department. Individual printers - physician's offices - as needed. \$1.5K, quantity not estimated.)	20	4	4
UPS power system (UPS power for routers and CPR server.)	16	=	=
sub totals			
Digital Radiology:			

Radiologist workstation (@ \$35K) (Sparc10, 5 GBy disk on 2 SCSI adaptors, FDDI, 2 high resolution gray scale displays)	105	35	35
Radiology film scanner (convert X-rays to digital images)	75	75	75*
Digital radiology film printer (convert digital X-rays to transparency film)	75	75	75*
Phosphor image plate reader (direct X-ray imaging to digital image)	75	=	=
sub totals			
(Not included: Multimedia gateway for existing PC and Mac systems; Mail servers and gateways for existing PC and Mac systems;)			
Total (approximate)			
Total (approximate) five sites			

3.1.2 Regional Data Center

TABLE 3. Regional Data Center

<u>ITEM</u>	Regional Trial per site (\$K)	OSL (single site prototype)	OSL (multi-site prototype)
<u>Mass Storage System:</u>			
MSS Server (Sun SPARCserver 2000 / SGI Challenge L, 50 GBy disk, on 4 SCSI adaptors on 4 SBus adaptors, FDDI, ATM, 256 MBy memory, CDROM, 8mm tape)	200	100	=
MSS Server (small version) (Sun SPARCserver 1000, 20 GBy disk, on 4 SCSI adaptors on 2 SBus adaptors, FDDI, ATM, 128 MBy memory, CDROM: S1000, 2proc, S1102-01-64-p66 = 47K second system board 1100A = 12K 64 MBy, 170A = 4.2K Sun Spectrum Silver (hardware maint.) = 0.5K 20 GBy disk.....)	=	85	85*
Tertiary Storage (2 x 1300 GBy, 5.25 write-once optical disk, jukebox, 2 drives)	700	=	=
Tertiary Storage (not suitable for long term radiology image storage) (1 x 460 GBy, 116x8mm tape, jukebox, 4 drives)	=	50	50*
MSS software	150	?? 75	? 75*
sub totals			
<u>Video / multimedia server - I</u>			
(2 x Sun SPARCserver 1000 / SGI Challenge L, 96 MBy memory, 10 GBy disk on 2 SCSI adaptors, FDDI, ATM, CDROM)	160	=	=
<u>Video / multimedia server - II</u>			
(SGI Challenge L, 2x150 MHz CPU, 96 MBy memory, 20 GBy disk on 2 SCSI adaptors, FDDI, ATM, CDROM)	=	110	110*
<u>CPR database server</u>			
(Sun SPARCserver 1000, 20 GBy disk, on 4 SCSI adaptors on 2 SBus adaptors (10 GBy mirrored), FDDI, ATM, 128 MBy memory, CDROM, 8mm tape)	85	(inc. in Tbl. 1)	(inc. in Tbl. 1)
<u>Security domain server</u>			
(SPARC 10/40 (S10EGX-40-32-p44), 2.5 GBy disk, FDDI, 2nd Ethernet) Sun Spectrum Silver (hardware maint.) = 0.17K/yr	25	25	25*
<u>Shared accounting and billing server</u>			
	??		
<u>Communications:</u>			
<u>ATM LAN switch / UNI:</u>			
(The BISDN user-network interface, attached on the inside directly to: radiology workstations; Internet router; ATM multimedia gateway.)	50	=	=
<u>Internet router</u>			
(IP router with ATM, Ethernet, and FDDI interfaces, together with Appletalk gateway.)	57	=	=
<u>Backup communications access</u>			
(Router w/ multiple dialup interfaces and modems, and/or ISDN interface.)	16	=	=

W. E. Johnston, Lawrence Berkeley Laboratory

<u>ATM multimedia gateway / codec</u> (Video and audio to ATM, a la the Wash. U. MMX. For high quality multimedia distribution.)	15		
<u>Multimedia workstations (@ \$10K)</u> (The basic CPR / information specialist workstation. Capable of multimedia and instrument record playback. Indies w/ATM or FDDI interfaces)	50	(inc. in Tbl. 1)	(inc. in Tbl. 1)
<u>CPR server</u> (For those sites that choose not to support a local DBMS server.) (Sun SPARCserver 1000, 96 MBy memory, 10 GBy disk on 2 SCSI adaptors (5 GBy mirrored), FDDI, CDRom)	72	(inc. in Tbl. 1)	(inc. in Tbl. 1)
<u>Hardcopy (5 X \$4)</u> (Laser printers, one 17 pg/min, individual printers - information specialists - as needed. \$1.5K.)	20	(inc. in Tbl. 1)	(inc. in Tbl. 1)
<u>Site preparation</u> (Fiber installation for FDDI ring and ATM point-to-point.)	15	(inc. in Tbl. 1)	(inc. in Tbl. 1)
<u>UPS power system</u> (UPS power for routers and servers)	30		
<u>Not included:</u>			
<u>Shared compute servers</u>			
<u>Multimedia gateway</u>			
<u>Any hardware and software needed for NOC functions</u>			
<u>Total (approx.)</u>	xxx		

3.1.3 Prototype Tele-Medicine Remote Site Workstation

TABLE 4.

<u>Multimedia instrumentation workstation</u> <u>(SPARC 10 / SGI Indy, 2.5 GBy disk, comm. interface, instrument interfaces, personal video camera)</u>	<u>10 - 25</u>		
<u>Radiologist workstation</u> <u>(Sparc10, 5 GBy disk on 2 SCSI adaptors, FDDI, 2 high resolution gray scale displays)</u>	<u>35</u>		
<u>Radiology film scanner</u> <u>(convert X-rays to digital images)</u>	<u>75</u>		
<u>Total (approximate)</u>	<u>135</u>		

| **3.2 System Design, Software Integration, etc.**

| (The figures below cover design, initial configuration, and training the initial operations staff.)

3.3 Operations

(Estimated on-going operations effort, per year, for two years. The estimates are based on an initial configuration of five sites, plus Hub.)

4.0 Timetable - Computing and Local Communications Infrastructure

White Paper on Computer-based Patient Record Systems⁹

(Extramural Version)
Frank Olken and William E. Johnston
Information and Computing Sciences Div.
Lawrence Berkeley Laboratory
1 Cyclotron Road, Berkeley, CA 94720
June 18, 1993

1.0 Introduction

Here we provide a brief introduction to the various types of computerized patient records (CPR) systems and an overview of this paper.

• Types of computerized patient records systems

- Picture Archiving and Communications Systems (PACS) are essentially electronic systems for storing, retrieving, displaying, and manipulating images. Primarily targeted at radiology departments, such systems could be used to capture patient records as images of the handwritten (or printed) records.
- Free Text Systems permit medical personnel to enter (typically by typing) unstructured (e.g., narrative) text. These systems provide for text storage, retrieval and display. Hypertext systems (i.e., cross linked text fragments) fall into this category.
- Structured Records Systems encompass traditional data management systems which are the computer analog of detailed medical forms (e.g., for patient histories, test orders, clinical lab results).
- Multi-media systems encompass all of the above technologies (images, free text, and structured records).

• Rationale - Why deploy a CPR system?

• Issues - What are the central issues to be addressed in deploying a CPR system in an open network environment?

• Implementation scenarios - Some historical patterns of implementation.

• Potential collaborators - Institutions who might collaborate on a CPR proposal to ARPA and their

⁹. Issued as LBL Tech. Report LBL-34303. This work was partially supported by the Director, Office of Energy Research, Office of Basic Energy Sciences, Applied Mathematical Sciences Division, and partially supported by Director, Office of Health and Environmental Research, both of the U.S. Department of Energy under Contract DE-AC03-76SF00098. Authors electronic mail addresses: olken@ux5.lbl.gov, wejohnston@lbl.gov

expertise.

- Cost & Development Cycle - for CPR systems.
- Dual use technologies - for CPR
- Possible foci of the CPR proposal.
- LBL competencies relevant to the CPR proposal.
- Further reading

2.0 Rationale

The advantages of CPR systems depend on the type of system implemented. Any distributed CPR system offers the prospect of improved record availability, remote access, and tele-medicine.

However, most of the advantages of computerized patient records arise from the use of structured records i.e., computerized forms rather than unstructured narrative text. Structured records greatly facilitate both the review of records (for quality & cost control, research, etc.) and the integration of the CPR system with decision support systems. Such CPR systems are the most demanding in terms data entry requirements.

The principal benefits seem to be better care (e.g. due to better records, and decision support), rather than direct cost savings. Cost savings would primarily accrue from more effective care, rather than savings in record keeping costs.

A detailed accounting follows:

- PACS - picture (image) archiving and communications systems.

1) Better availability of records than manual systems. There is less likelihood that records will be in transit or lost. Results (of lab tests) can be transferred immediately to attending physicians.

2) More complete records, since it is easier to assemble records from multiple sites.

3) Remote access via telecommunication networks. Tele-presence of medical personnel.

4) Simultaneous access - CPR systems readily permit multiple persons to examine the same medical record simultaneously.

5) Image enhancement - A variety of image processing techniques (noise filtering, contrast enhancement) and image display techniques (pseudo-color) are easily implemented.

- Free narrative text

1) Advantages 1-4 of PACS systems (above).

2) Better legibility of typewritten text vs. handwriting.

3) Automated text searching capabilities, e.g., keyword searching.

4) Some prospect of machine processing of the narrative text to extract diagnoses, etc., for research or epidemiological studies.

- 5) Some improvement in research access - because the entire medical record is machine readable
- 6) Integration with bibliographic search capabilities, so that doctors can examine relevant medical literature.
- 7) Less duplication of tests, because the test records will be available from remote sites and easier to find in the patient record.
- Structured records (computerized forms)
 - 1) Usually entails standardized terminology which facilitates analysis and retrieval.
 - 2) Structured records are typically better organized, hence easier to review.
 - 3) More complete, because they are typically generated from a systematic questionnaire.
 - 4) Multiple "views" of the patient record are easily supported to accommodate requirements of different users (pharmacy, nursing, doctor, etc.)
 - 5) Structured records are much better suited to integration with decision support systems (i.e., diagnostic and treatment advisory systems).
 - 6) Structured records typically provide much simpler and more reliable research access to medical records. Natural language understanding requirements are greatly reduced.
 - 7) Structured records facilitate ongoing epidemiological surveillance (for new epidemics) because they provide more systematic recording of symptoms, diagnoses.
 - 8) Easier review of structured records makes quality control studies of medical care much easier.
 - 9) Structured records can be readily used as the input to automatic alarm generation programs which notify doctors/nurses of serious medical problems.
 - 10) Experience has shown that structured medical records and attendant decision support systems can reduce medical malpractice costs.
 - 11) Structured records facilitate automatic checking of drug interactions/contraindications and notification to the doctor.
 - 12) Automatic billing and report generation is much simpler from structured records than free text.
- Multi-media system
 - 1) Such systems support all of the above types of data and hence have all of the advantages described above.

3.0 Issues in Implementation

Here we identify the major issues involved in deploying a CPR system.

- Security is a central concern for CPR systems, especially for distributed CPR systems which rely on common carrier communications. Present systems do not appear to have adequately addressed this problem.

- Access control is needed to protect patient privacy and to prevent unauthorized entries to patient records, e.g., unauthorized medication orders.
- Medical records are legal documents and all medical personnel are required to sign each entry in a patient record. In a CPR system this probably entails the use of digital signatures.
- Encryption of data transmission will certainly be required when traversing common carrier communication facilities.
- Data acquisition (entry) has been another major problem for CPR systems.
 - Medical personnel (doctors, nurses, radiologists, etc.) have traditionally handwritten their remarks in the patient record. This is generally unsuitable for CPR systems.
 - + Alternative methods of text entry are: typing (but many doctors can not type quickly), real-time handwriting recognition (easier than off-line handwriting recognition), speech recognition (problematic due to large vocabularies), transcription (expensive, and sometimes error prone), or menu picks (attractive where possible).
 - + Two central concerns are: the portability of terminals (so that they can be used by doctors as they make rounds), and the ability to enter data with one hand while standing (e.g., in a clinical setting) - an advantage of writing over typing.
 - Instrumentation (clinical lab, ICU, radiology) are major sources of data. The primary issue here is standardization of data exchange formats.
 - Many hospitals have existing departmental systems (clinical lab, admissions, radiology, pharmacy) with which data must be exchanged (unless the departmental systems are to be replaced).
 - In a distributed setting (hospital consortia, hospital/clinic/private practitioner settings) data must be exchanged with external CPR systems.
- Standardization is clearly a key effort to the success of CPR systems and there are a number of standards efforts underway:
 - Data exchange standards (formatting, protocols) are needed between CPR systems, CPR and departmental systems, etc.
 - Data exchange standards are needed for instrument interfacing.
 - Nomenclature standardization (diagnoses, symptoms, etc.) are needed for the consolidation of medical records from multiple doctors, or hospitals, to facilitate research, epidemiological studies, decision support, and other uses of medical records.
- Legal Issues:
 - The legality of electronic records has not been established in many states.
 - Security and privacy policies and technologies are either presently or potentially subject to legislation. It has been suggested that privacy issues (esp.) should be supported by legislation which would make it a crime to make unauthorized disclosures. The legal status of digital signatures in a medical setting needs to be clarified.
 - Some states specify the retention period of patient records (e.g. 25 years). Ideally, one

would want to retain patient records for the lifetime of the patient - and they may be of use for diagnoses and treatment of inherited disorders in children.

- Liability for errors and failures of the CPR system and associated decision support systems needs legal clarification. It has been suggested that CPR systems will be subject to strict liability, while decision support systems will be subject to more traditional negligence criteria, because they are only advisory.
- The choice of the database management system (DBMS) to accommodate the diverse types of data, queries, access controls, etc. remains problematic. A full-blown multi-media CPR system will require support for:
 - Structured records, i.e., either relational or object oriented DBMSs,
 - Text, both storage of long text fields, indexing and query support (typically poorly supported in existing DBMSs),
 - Images - 2D, 3D, radiology, pathology - ranging in size up to 20 Megapixels (each 12 bits),
 - Voice data,
 - Time series (electroencephalograms, electrocardiograms),
 - Temporal data (the temporal evolution of the patient's health status)
 - Thesauri - automatic support for variant terminology
 - Schema evolution - Advances in medicine dictate changes in the content of the patient record - typically requiring modification of the database schema.
 - Tertiary storage management - Archival record keeping requirements suggest that the DBMS will have to manage data on tertiary storage (optical disk, magnetic tape) as well as secondary storage (magnetic disk).
- Text processing facilities will be needed: indexing, querying, natural language understanding
- The ability to abstract medical record information by summarizing incidents and treatments is needed.
- Convenient, comprehensible user interface design is needed (probably graphical interfaces) if the systems are to obtain acceptance. First generation systems were abandoned in part because medical personnel found them too rigid and cumbersome to use.
- Telecommunications capabilities are clearly needed to support distributed CPR systems:
 - Wireless (within a hospital or clinic, also paramedics)
 - LAN - hospital
 - MAN - metropolitan networks to support hospital consortia, hospital/clinic/private practitioners
 - WAN - regional, national networks to support consolidation of records from multiple sites, tele-medicine, etc.
- Storage requirements:
 - Adequate on-line storage is needed to provide rapid access (preferably less than one second)

to current patient data.

- Archival data facilities are needed to retain patient record data for at least 25 years, preferably for the patient's lifetime.
- Images will account for the bulk of the storage requirements. They typically require 15-20 megapixels at 12 bits per pixel (for radiological images).
- Essentially one hundred percent availability is essential if the CPR system is to replace paper records. Medical emergencies can not accommodate scheduled downtime or record unavailability due to disk failures.
 - + Some sort of storage replication will clearly be necessary, either mirroring or RAID technology.
 - + Physical dispersal of replicated copies of the patient record data is desirable to prevent data loss during fires, riots, natural disasters.
- The ability to support heterogeneous DBMSs is needed to accommodate:
 - Existing, disparate departmental computing systems - e.g., clinical lab, radiology, pharmacy, billing, etc.
 - Disparate CPR systems at multiple sites - hospitals, clinics, private practitioners' offices, HMOs.

4.0 Implementation Scenarios

Five medical functions have proven the most attractive targets for computerized patient records to date:

- Billing
- Clinical Laboratory Records - Sample tracking, recording lab results, fast communication of results, reduced transcription errors
- Radiology - digital acquisition, storage, processing, retrieval and communication of images. Necessary for CAT scans, MRI, ultrasonic scanning.
- Pharmacy - detailed records requirements, automatic checking for drug interactions.
- Self-administered patient histories - Computer controlled branching questionnaires for patient history taking. This saves time of medical personnel, helps assure completeness, easier for patients to describe sensitive topics (sexuality). Such systems can also be readily designed to accommodate non-English (Spanish, Chinese) speaking patients.

Typically, hospitals will already have systems in place for these functions. Such departmental systems will have to be integrated in the omnibus CPR system (or replaced). Given that these functions have proven the most attractive to date, one would expect that a CPR system would commence with these departments.

Another attractive target for CPR are emergency rooms, where requirements for fast medical treatment place a premium on fast access to patient records. Decision support systems have also proven useful in emergency rooms in reducing malpractice problems.

W. E. Johnston, Lawrence Berkeley Laboratory

5.0 Issues to Address

- Security - access control / digital signatures. Presently a weak link in existing CPR systems. Access control, digital signature, and encryption technology exist which should be able to address most of the CPR concerns.
- Data entry - wireless portable terminals / handwriting recog./ speech recog.
- Heterogeneous DB systems - to accommodate departmental systems and external CPR systems. LBL has expertise in this area.
- Distributed replicated storage - e.g., the image server

6.0 Reference Experts at LBL

- Heterogeneous DB
 - Schema integration - V. Markowitz
 - Record matching - A. Segev & F. Olken
- Temporal DB - A. Segev & A. Shoshani
- Image DB - D. Rotem
- Laboratory Information Management Systems - V. Markowitz, I.M. Chen, F. Olken
- Imaging - W. Johnston
- Distributed Image Servers - W. Johnston, D. Rotem
- Distributed Systems Integration - W. Johnston
- Replicated Data - A. Segev, D. Rotem
- Time Series Data - A. Segev, F. Olken
- Sampling from DB (for quality control / epidemiology) - F. Olken, D. Rotem
- Thesauri and Metadata - J. McCarthy

7.0 Cost, Development Cycle

CPR systems thus far have been expensive and time consuming to develop and deploy. Costs of \$2-6M for a medium size hospital to deploy a CPR system have been reported (> \$1K per bed). To develop a new CPR system from scratch has been estimated at up to \$10M and 5-10 years. Development costs and times should decline as better components (DBMSs, distributed systems software, imaging software) becomes more readily available. Deployment costs should fall with declining hardware costs.

8.0 Dual Use Technologies

CPR systems are technologies with dual civilian and military use. Even in peacetime the DOD is among the largest healthcare providers in the country (for 1.5M armed services personnel and their dependents). War places even greater demands on DOD for healthcare.

Furthermore distributed CPR would provide an excellent test of the networking, imaging, and distributed systems technologies for which DOD has sponsored the development. In particular CPR systems place stringent requirements on access control, and record verification (digital signatures) which are increasingly important considerations for a National Information Infrastructure.

9.0 Further Reading

See [1,2, and 3] and the references cited therein.

- [1] Ball, M., and M. Collen, editors. Aspects of the Computer-based Patient Record. Springer-Verlag, 1991.
- [2] Committee on Improving the Patient Record, Institute of Medicine. The Computer-based Patient Record: An Essential Technology for Healthcare. National Academy Press, 1991.
- [3] Government Accounting Office. Medical ADP Systems: Automated Medical Records Hold Promise to Improve Patient Care. Technical report GAO/IMTEC-91-5, Jan. 1991.