Tools for Managing Image Flow in the Modality to Clinical-Image-Review Chain

Kenneth W. Clark, MS,¹ David L. Melson, BS,² Stephen M. Moore, MS,¹ G. James Blaine, DSc,¹ Ralph A. Moulton, MS,² William K. Clayton,² Colin S. Peterson, BS,² and Bruce A. Vendt, MBA²

Web-based clinical-image viewing is commonplace in large medical centers. As demands for product and performance escalate, physicians, sold on the concept of "any image, anytime, anywhere," fret when image studies cannot be viewed in a time frame to which they are accustomed. Image delivery pathways in large medical centers are oftentimes complicated by multiple networks, multiple picture archiving and communication systems (PACS), and multiple groups responsible for image acquisition and delivery to multiple destinations. When studies are delayed, it may be difficult to rapidly pinpoint bottlenecks. Described here are the tools used to monitor likely failure points in our modality to clinical-image-viewing chain and tools for reporting volume and throughput trends. Though perhaps unique to our environment, we believe that tools of this type are essential for understanding and monitoring image-study flow, reconfiguring resources to achieve better throughput, and planning for anticipated growth. Without such tools, quality clinical-image delivery may not be what it should.

KEY WORDS: Clinical-image viewing, imaging throughput, *Clinical Desktop*, *Imageweb*

cannot solve the problem, logs a trouble-ticket and summons the next level of PACS expertise. It may be tens of minutes before the clinician sees the images.

Complicating problem resolution in a large medical center are any or all of the following: imaging modalities from different vendors, multiple PACS, multiple information systems responsible for image acquisition and consumption, multiple and independently administered networks, and clinician groups with imaging needs quite different from those of diagnosticians. These groups differ in their assessments of PACS performance, and Nagy et al.⁴ suggest that their disparate metrics retard timely analysis and problem correction; furthermore, PACS administrators may not have full and ready access to good performance-assessment tools. In remedy, Nagy et al have pioneered an open-source toolset giving the PACS administrator centralized user-centric metrics.⁴

What follows is our approach to supporting clinical- and research-image viewing. Although

“WHY CAN’T I SEE MY IMAGES?”
This question is often posed by a clinician who has been promised image-viewing anytime and anywhere to picture archiving and communication systems (PACS) help desk staff, who must then sort through a mental checklist of probable trouble spots before mollifying the anxious physician. Is the network up? Is it saturated? Is the image server up? Have the images reached the image server? Is the Radiology Information System (RIS) working smoothly? Is there a user problem? And so on. Dialog ensues, symptoms are discussed, and the help desk, if it

¹From the Mallinckrodt Institute of Radiology, Electronic Radiology Laboratory, Washington University School of Medicine, 510 South Kingshighway, Saint Louis, Missouri 63110, USA.

²From the Mallinckrodt Institute of Radiology, Computer Section, Washington University School of Medicine, 510 South Kingshighway, St. Louis, Missouri 63110, USA.

Correspondence to: Kenneth W. Clark, MS, MIR/ERL Campus Box 8131, 510 South Kingshighway, Saint Louis, MO 63110; tel: 314-362-6965; fax: 314-362-6971; e-mail: clarkkk@mir.wustl.edu

Copyright © 2003 by SCAR (Society for Computer Applications in Radiology)

Online publication 15 December 2003

doi: 10.1007/s10278-003-1724-1

somewhat dependent on fluid PACS operation and administered by the PACS Group of our Radiology Department, the clinical-image service is a major operation of its own. Our approach is not a panacea, nor has it fully evolved. But it contains ideas that might prove useful to those charged with the management of 24 × 7 clinical-image viewing, yet who are often at the mercy of physical and human resources over which they exercise little control.

HISTORICAL PERSPECTIVE

In the mid 1990s, the BJC Health System (now BJC Healthcare) and the Washington University School of Medicine formed a technology alliance with IBM, Eastman Kodak, Southwestern Bell, and Motorola to develop an integrated clinical repository for BJC’s acute-care hospitals in Missouri and Illinois. The alliance was named *Project Spectrum*.¹ Access to the repository was supported with Web-based technology to deliver patient data, including radiology reports and images, to physician desktops. A pilot project deployed several hundred *Spectrum* workstations; but clinical data and image retrieval proved too slow, and corporate support waned, preventing *Project Spectrum* from reaching full acceptance. Nonetheless, a foundation had been laid for next-generation clinical and image data retrieval.

Concurrent with *Project Spectrum* development, our radiology research community demanded access to images of patients, anatomical specimens, and phantoms. Mounting research needs taxed modality devices and their technologists, who needed to accommodate nearly as many image-study destinations as the researchers asking for them. We developed a model by which all research images would be sent to a “distributor” machine that dealt studies to pre-configured destinations, including a research Web server from which investigators could pull their own studies. The advantages to this approach were: that (1) the modality technologists needed to know only one destination, (2) the modality device needed to be configured only once, and (3) researchers could pull their studies without PACS-department assistance.

Fig 1. Generalized image flow. Images acquired at the modality are sent to PACS, where they are auto-routed to an Image Relay Node and, from there, to the Clinical Image Server. Diagnostic radiologists retrieve their studies from PACS. *Imageweb* and *ClinDesk* image-viewing applications within a viewing-client PC retrieve studies from the server. The *ClinDesk* application is an electronic patient record; if radiology images are available, image retrieval is enabled via handshaking between the Clinical Data Repository and the Clinical Image Server. In reality, there are multiple PACS and a pool of six Image Relay Nodes. Radiology reports are available to the *Imageweb* client via the Broker and to *ClinDesk* via IDXrad and the Clinical Data Repository.

This model still serves our researchers and has evolved to serve the wider clinical-image viewing community.

INSTITUTIONAL SETTING

BJC HealthCare is one of the largest non-profit health care organizations in the United States, providing medical services through 13 facilities in the greater St. Louis area, southern Illinois, and mid-Missouri regions, with net revenues of \$2.2 billion; 140,000 admissions; 372,000 emergency-department visits; 4,300 licensed beds; 4,600 physicians; and 26,000 employees. Geographically co-located with BJC HealthCare’s flagship Barnes-Jewish Hospital (BJH), the Washington University School of Medicine (WUSM) provides significant professional services. This article focuses on those services, specifically clinical-image viewing, within the BJH-WUSM medical center.

Table 1. Annual Clinical Image Studies and Gigabytes by PACS and BJH Modality

PACS Vendor	MIR Section	Modality Devices (#)	Studies per Year	Gigabytes per Year	Studies Sent to Clinical Image Server
Siemens	CT, MR, ED, GI/GU, Ortho	CT (13); MR (8); CR (8); RF (10)	314,285	9,980	All
Philips	Chest, ICU	CR (2); DR (5)	53,812	1,081	All
MIR	Nuclear Med	NM (18); PT (2)	14,048	29	Some
ALI (McKesson)	Ultrasound	US (7)	14,401	274	All
GE	Mammography	MG (1)	0	0	None
Total			396,546	11,364	

Note. Projected from Actual Data for 9 months (August 2002 through April 2003).

METHODS

Clinical-Image Viewing

Radiology Services

Figure 1 depicts our medical center’s flow of clinical images. BJH owns and operates the radiology image-acquisition modalities. The WUSM Mallinckrodt Institute of Radiology (MIR) provides professional diagnostic interpretation, and its Computer Section manages the RIS, multiple PACS, and all clinical-image viewing. Clinical-image service operates quasi-independently of the diagnostic side of PACS for four main reasons.

1. Both diagnosticians and clinicians need rapid access to their studies and they prefer to avoid contentious study requests.
2. Clinicians needing to review studies from differing modalities might be required to query multiple PACS when they would prefer to retrieve studies from a single source.
3. Management of our several PACS is unique enough without the addition of servicing clinical-image viewing, which those PACS, in their current configuration, do not now support.
4. PACS-vendor viewing workstations are far more expensive than browser-enabled PCs.

Table 1 describes the various PACS and their associated modalities, together with annual numbers of studies and gigabytes involved. None of mammography studies and only a fraction of the nuclear-medicine studies are available for general clinical-image viewing; these are more the domains of specialized diagnosticians.

Radiology Information System (RIS)

IDXrad, our RIS, sends patient data, exam, and report information to a Mitra Broker, which provides work lists to imaging modalities. IDXrad also exchanges messages with a VMS (Virtual Memory System) machine, which, with additional input from PACS and Image Relay Nodes (below), monitors image-study movement and alerts PACS personnel of studies delayed on their way to intended destinations.

The image server application software is Mitra Exhibit 3.1. The image server hardware is a Dell PowerEdge 6450 with quad 700 MHz processors, 2 MB L2 cache, and 8 GB RAM. Operating system is Windows 2000 Advanced Server. Connection to SAN (Storage Area Network) storage is dual fiber channel Emulex LP9002L-E Optical Host Bus Adapters. The SAN controller is an EMC corporation Clarion CX600 with an 8 GB cache and 2.6TB usable storage. Fiber channel switches are Brocade 3802 2 GBps Silkworm.

We call the whole hardware and software configuration *Imageweb*. Image studies may be retrieved from the *Imageweb* server, via Web browser, to client computers for viewing. The Mitra Exhibit client-viewing software permits image manipulation but prohibits saving changes. Retrieval is either indirect, via BJC HealthCare’s electronic-patient-record application, called *Clinical Desktop* (*ClinDesk*), or it is direct, via so-called *Imageweb* workstations or MIR desktop machines. Access to both *ClinDesk* and *Imageweb* is via secure network or virtual private network (VPN) and per-application password. Radiology reports are available to *ClinDesk* via IDXrad and to *Imageweb* via the Broker.

ClinDesk is a patient-centric application in which a clinician opens a “chart” for a particular patient. The electronic chart mimics the paper chart with color-coded side tabs for orders, monitoring, laboratory, radiology, diagnostics, treatments, notes, correspondence, and management. *Imageweb* client viewing is more practice-centric, where a clinician chooses to review, for example, “today’s CT studies,” without having to sift through individual charts.

Clinical Desktop Client

BJC Healthcare has installed approximately 3,000 *ClinDesk* workstations. About 6,000 (of 11,000 with logon privileges) distinct users make 175,000 “chart-opens” per month. Requests for radiology images via *ClinDesk* number about 20,000 per month. Roughly 60% of these *ClinDesk* figures may be ascribed to the BJH-WUSM medical center, although both patients and physicians may move among BJC HealthCare facilities.

Imageweb Client

All radiologists and MIR staff may retrieve image studies to their desktop machines. In addition, specific *Imageweb* workstations with calibrated medium-resolution monitors have been deployed to several image-intense clinical practices (Lung Center, Neurosurgery, Orthopedic Surgery).² Viewing requires only a PC with 128 MB RAM and color monitor with 1280 × 1024 resolution. The application is called via standard Web browser with password access. Radiologists and specialty-service physicians (450 with logon privileges) request approximately 5,000 *Imageweb* studies per month.

Image Flow

Mostly, imaging modalities send studies directly to a PACS. The PACS, in turn, autoroutes to a pool of Image Relay Nodes (IRNs, below) that prioritizes the distribution of studies to the operational *Imageweb* server and to other destinations. Our Philips PACS cannot be configured to route inbound studies to multiple destinations; studies destined for the Philips PACS are sent from their modalities to the IRN pool, which can simultaneously route to the Philips PACS and to the *Imageweb* server.³

IRN Pool

The IRNs are UNIX machines whose task is to deliver images to their intended destinations as quickly as possible. They route studies according to certain queuing and priority schemes such that, for example, a wrist image from the Orthopedic Clinic would be routed before an MR study of a knee. There are five operational IRNs and one research IRN, and their queues and routing schemes may be manually adjusted for load balancing. On each IRN is a MySQL database (below) of information about each study and the time taken to receive the study (from modality or PACS) and to transmit to its destination(s). An image study passing through an IRN adds slightly to the overall image-delivery time, but the delay pales against the value of the data captured as input to activity monitoring.

Imageweb Server

The operational server (1) accepts new studies from the IRNs, (2) responds to client queries to retrieve studies, and (3) manages several months of clinical studies. The first task converts inbound digital imaging and communication in medicine (DICOM) images to wavelet representation and stores them. The second transmits to requesting clients, lossless or lossy, as directed by the client. The first and second tasks contend with each other and may somewhat retard the rate at which studies are accepted at the server and then become available for viewing. The third task, specifically the deletion of older studies not accessed for some time, is performed during late-night low-activity periods, within constraints of high/low watermarks. The server uses four processors to accommodate the load.

Other Destinations

Besides a production *Imageweb*, we maintain two other *Imagewebs*. One of these is used as a disaster-recovery back-up; it is older and not as fast as the production version and its on-line capacity can store only a few months of studies; all studies sent to the production *Imageweb* are also sent to this back-up system. The third is used to test new releases of the software, and studies are sent to it as needed for testing.

In addition to these three *Imageweb* systems, studies may also be electively sent to a research image server that provides no image-viewing service. Rather, a researcher logs into the research server, finds her/his study among a list of available studies, and retrieves the study to a workstation chosen from a list of approved destinations. Studies are retained on the research server for only 2 weeks.

Networks

The BJH and WUSM maintain separate secure networks. The hospital-owned modality devices are on the BJH network, and the IRNs and *Imageweb* servers are on the WUSM private and public networks. *Imageweb* clients may be on any of the three. Firewall rules between the networks afford smooth operation. Image-viewing clients on PCs outside of either secure network may gain access using a VPN client on that PC (Fig 2).

Image Flow Monitoring Tools

There are two classes of image-flow tools: real-time and historical. The real-time tools help understand evolving problems, and the historical tools help understand trends and are useful in planning load balancing and growth.

RIS, IRN, and Network Tools

A VMS machine receives RIS-level messages to assess timely functioning of modality work list provision. The VMS also receives two kinds of messages from the IRN pool. The first is node-specific; the scheme uses “pings” and “DICOM echoes” to report if any of the IRNs, image servers, or their network links appear to be malfunctioning. The second is study-specific and reports studies that do not reach their intended destinations at expected times, where expectation is based on study size, type of study, priority, and current workload. The VMS automatically pages on-call personnel when faults are detected, and IRN logging files are parsed to harvest key information, which are saved in a database (MySQL), against which a variety of management queries may be posed.

IRN MySQL Tools

Each IRN has its own MySQL database. Queries may be made to any combination of IRNs. A query is initiated from a Computer Section PC with a batch file. The batch job passes the query to the individual IRNs, and the concatenated stream of results is passed to a report-generating PC. Items saved for each study are listed in Table 2.

Fig 2. Network complexity. Firewalls exist between the BJH private network and the University networks. Multiple PACS are found on both university and hospital networks. Image-viewing clients on the university public network require VPN. The IRNs are found on all three networks.

Table 2. MySQL Database—Per Study Elements

Study Data	IRN Data	Patient Data	Date-Time Groups	Destination Data
Modality	IRN name	Patient name	DICOM study	Destination name
Accession number	IRN queue	Patient ID	Arrival of first image	Transmit start
Institution name			Arrival of last image	Transmit complete
Station name			Last image transmitted	
Source directory				
Number of images				
Number of kilobytes				

Three time-based workload measurements are made: (1) *receive-time* from the modality, defined as the difference in arrival times between the first and last images of a study; (2) *latent-time* within the IRN, defined as the time difference between the arrival of the last image in a study and the time the study was transmitted to its destination; and (3) *transmit-time* to *Imageweb*, defined as the time difference between the first and last images transmitted. Studies routed to multiple destinations may have differing transmit-times. The *receive-time* and *transmit-time* may be affected by network traffic, and *transmit-time* may additionally be affected by the rate at which servers accept studies. *Latent-time* may be affected by internal multitasking to handle numerous inbound and outbound studies, as well as by deliberate prioritizing schemes.

Reporting Tools

MySQL query result files are sent to a report-generating machine in tab-delimited format where they are filed in a

Table 3. Number of Image Studies and Gigabytes, by Modality, July 2002 through April 2003

Modality	Studies (%)	Gigabytes (%)
CR	195,043 (57)	2,755 (28)
CT	73,206 (22)	4,807 (49)
MR	33,567 (10)	1,780 (18)
DR, DS, NM, OT, PT,RF, RG, SC, US, XA	39,090 (11)	488 (5)
Total	340,906 (100)	9,830 (100)

Microsoft (MS) Access database. Queries posed to the Access database produce reports and charts that may be imported to MS Excel spreadsheet pivot tables from which any number of summary plots may be made. Reports typically summarize volume and/or throughput for a given period (yesterday, last week, last month, etc.).

Fig 3. Monthly CR, CT, MR median receive, latent, transmit times, August 2002 through April 2003.

Image Server Tools

Imageweb's server software has some useful tools for assessing performance, although access to them is independent of the image-viewing portion and requires a separate browser window with only administrator-level access. Furthermore, the harvesting of internal logging information for external consumption is not easily understood. We are, however, able to extract a time-stamped stream of each logon, logoff, and study viewed. Weekly, this stream is saved; and the number of requests to view studies, by source, is posted to an Excel spreadsheet and charted. Sources are ClinDesk users, Neurosurgery, Orthopedic Surgery, and various MIR sections.

Tool Integration and Management

An internal Web site is under construction to allow our Computer Section management team to easily view (1)

current status of the RIS, Broker, the IRNs, and the Imageweb server; (2) recent and historical volume and throughput reports. The second item is further along, and the user merely specifies the date or date range over which a report is desired.

The MIR Computer Section (CS), the support arm of our medical center's radiology imaging services, manages the following activities: the radiology information system (RIS), all PACS operations except hardware maintenance, the Radiology Department's desktop-computing infrastructure, an MIR network, and the Imageweb clinical-image viewing service. A CS Help Desk triages reported problems to appropriate levels of expertise. Many of the 24 full-time CS employees work within more than one of these areas and are cross-trained to do so.

The Imageweb service requires solid relationships with the application vendor (Mitra), the hardware vendor (Dell), the storage vendor (EMC), and the development and management teams of BJC's Clinical Desktop electronic patient record application. The "home-made" image-distributor

Fig 4. Hourly median latent times (CR, CT, and MR), Monday through Friday, August 2002 through April 2003. Computed radiography (CR) latency remains steady in the 5-minute range. The CT and MR latencies build from late morning before falling early evening. These delays are a function of both volume and the rate at which the image server can accept and process that volume. The high CT and MR latencies in late evening are due to (1) the image servers shutting down their inbound queues while databases are updated, (2) researchers performing large studies when instruments are available to them, and (3) the Radiology Help Desk pulling prior studies from PACS for the next day's Neurosurgery Clinic. These priors date beyond the clinical image server's stock of current studies.

mechanism, which feeds image studies to *Imageweb*, is managed and developed by the CS. On average, *Imageweb* support requires approximately 3.5 FTEs. The CS director and an Executive Physician Group of radiologists provide significant oversight.

RESULTS

Table 3 shows volumes and megabytes, by modality, July 2002 through April 2003, passing through the IRN pool to the *Imageweb* server.

Figure 3 shows CR, CT, and MR median receive, latent, and transmit times by month. The slight upswings noted in April are due, in great part, to our testing of a new release of the *Imageweb* application for which the IRN pool added another destination. Both computed tomography (CT) and magnetic resonance (MR) latent time dropped dramatically in October 2002 when we added two IRNs and redistributed internal queues (Fig 3d).

Figure 4 shows median receive, latent, and transmit times by hour, Monday through Friday, August 2002 through April 2003.

Figure 5 shows *Imageweb* requests to view images, weekly, from 30 December 2001 through April 2003. About 75% of these requests are from *ClinDesk* users. Requests ap-

Fig 5. Weekly clinical-image study requests. "Clinician" requests showed a steady rise in the 2nd half 2002, into 2003. The MIR sections, Neurosurgery, and Orthopedic Surgery group have been fairly steady. Overall, requests seem to have leveled in Spring 2003.

pear to have leveled in 2003. The remaining 25% (a steady 1000/week) are Radiology, Orthopedic Surgery, and Neurosurgery users.

DISCUSSION

The results are but a snapshot of our dynamic environment. A new *Imageweb* (Agfa Web 1000) was being tested during the second quarter of 2003; thus, the IRN pool was feeding the same image studies to three *Imageweb* servers (production, back-up, test). Provided this additional burden to the IRNs does not impact clinicians' perceptions of timely image delivery, we shall continue to feed the three systems. If delivery slows, we have the ability within the IRN pool to delay sending to the test system until lighter-traffic hours.

The results suggest an acceptable stability over a 9-month period recently completed. This is a tribute to the queue-restructuring flexibility we have built into our model and to the dedication of those charged with keeping these systems and networks operating smoothly. Success has not come without challenges. First, we had a major hardware failure of the operational *Imageweb*; our backup system took over until the primary was repaired and its database restored. Second, there have been many changes to the BJH and WUSM networks as well as IRN deployment shifts across the networks.

The tools described above have proven their mettle on three fronts. First, rapid growth in clinical-image viewing suggested we add two IRNs, which we did in October 2002. Second, we found we could improve overall throughput

by redistributing several queues within the IRNs. Third, by adding additional processors to the *Imageweb* server, we could greatly improve both the rate at which studies would be available for viewing once they arrived at the server and the rate at which the IRN pool could transmit studies to server. Without our toolset, we would find it difficult to gracefully react to short-term challenges and plan for longer-term growth.

Admittedly, the monitoring environment described here is highly customized and has evolved over years as needs arose and with limited financial and personnel resources. Given what was learned and knowing that roughly equivalent commercial alternatives are now available, it would be reasonable to pursue an alternative, provided the out-of-pocket and intangible costs, both up-front and long-term, make business sense. Regardless, we hope that this report provides significant ideas and insight to those charged with the acquisition and management of clinical-imaging provision.

CONCLUSIONS

We have found these tools essential for understanding our enterprise’s clinical-image viewing complexities, measuring and monitoring activity, reacting responsibly to bottlenecking problems as they arise, re-configuring resources to achieve better load-balancing, and planning for anticipated growth. Without these,

or some equivalent workload-characterization tools, quality image delivery may rarely meet clinicians’ expectations.

Each medical-center enterprise has its own unique complexities, and our approach to monitoring clinical-image delivery may not universally apply. However, we do believe that segmenting issues, problems, solutions, and monitoring tools into the right sizes will make image delivery to clinicians better understood and more effective.

ACKNOWLEDGMENT

Funding for this study was provided by Washington University School of Medicine, Mallinckrodt Institute of Radiology.

REFERENCES

1. Blaine GJ, Jost RG, Martin L, et al: Information and image integration: Project Spectrum. Medical Imaging 1988, PACS Design and Evaluation: Engineering and Clinical Issues, Horii SC, Blaine GJ (eds). Proc SPIE 3339:430-439, 1998
2. Clark KW, Francis S, Blaine GJ, et al: Planning and implementing image delivery to outpatient specialty clinical practices in a large medical center. J Digit Imaging 15(Suppl 1):156-161, 2002
3. Melson DL, Moore SM, Blaine GJ, et al: Challenges in image acquisition and distribution for clinical image service. J Digit Imaging 15(Suppl 1):144-150, 2002
4. Nagy PG, Daly M, Warnock M, et al: PACSPulse: a Web-based DICOM network traffic monitor and analysis tool. RadioGraphics 23:795-801, 2003