

## Multidimensional Analysis: A Management Tool for Monitoring HIPAA Compliance and Departmental Performance

Robert M. Coleman, BS, Matthew D. Ralston, MD, Alexander Szafran, MS, and David M. Beaulieu, BS

**Most RIS and PACS systems include extensive auditing capabilities as part of their security model, but inspecting those audit logs to obtain useful information can be a daunting task. Manual analysis of audit trails, though cumbersome, is often resorted to because of the difficulty to construct queries to extract complex information from the audit logs. The approach proposed by the authors uses standard off-the-shelf multidimensional analysis software tools to assist the PACS/RIS administrator and/or security officer in analyzing those audit logs to identify and scrutinize suspicious events. Large amounts of data can be quickly reviewed and graphical analysis tools help explore system utilization. While additional efforts are required to fully satisfy the demands of the ever-increasing security and confidentiality pressures, multidimensional analysis tools are a practical step toward actually using the information that is already being captured in the systems' audit logs. In addition, once the work is performed to capture and manipulate the audit logs into a viable format for the multidimensional analysis tool, it is relatively easy to extend the system to incorporate other pertinent data, thereby enabling the ongoing analysis of other aspects of the department's workflow.**

**KEY WORDS: HIPAA, security, confidentiality, audit logs, multidimensional analysis, OLAP, PACS, picture archiving and communications systems, RIS**

### BACKGROUND

The Health Insurance Portability and Accountability Act of 1996 (HIPAA), among other things, has greatly raised the level of awareness regarding the security of confidential data in information systems throughout the health care enterprise. Among those systems, the Radiology Information System (RIS)

and Picture Archiving and Communications System (PACS) contain vast amounts of confidential information and are coming under increasing scrutiny by health care privacy officers. Moreover, access to these systems, especially web-based radiology image and report systems, is often granted to hundreds or even thousands of clinicians. While every attempt is made to limit access to a need-to-know basis, often such restrictions cannot be adequately manipulated to restrain access because it is nearly impossible to predict who will legitimately need access to a particular patient's record. Lock down the system too tight, and the smooth flow of clinical information will be hindered, thereby negatively impacting patient care. Open it up too far, and the risk of unauthorized use will increase.

One of the tools PACS and RIS vendors have used to help attack this problem is to build detailed access audit triggers into their database systems. Every user-initiated event (eg, login, logout, query, image view, report view) is captured in an audit log. In conjunction with strong username/password authentication, these audit logs can then be used to monitor the system for irregular user access or use. While not as valuable as preventing unauthorized access, quickly gathering detailed evidence of

---

*From the Maine Medical Center, Portland, ME, 04102, USA.*

*Correspondence to: Robert M. Coleman, BS; tel: 207-871-6534; fax: 207-871-3226; e-mail: colemr@mmc.org*

*Copyright © 2004 by SCAR (Society for Computer Applications in Radiology)*

*Online publication 15 July 2004*

*doi: 10.1007/s10278-004-1012-8*

suspicious use can go a long way toward halting unscrupulous users. Substantiated retaliation is an excellent deterrent.

The challenge, then, becomes one of how to analyze the audit logs to identify potential problems. At Maine Medical Center, the audit logs on the 1650-user web-based image distribution system grow by 25,000–30,000 records *per week*. This volume of data demands an innovative approach to quickly and easily distill the information down to a manageable subset of records that may warrant additional focus. A multidimensional analysis/online analytical processing (OLAP) tool proved to be a great fit for this application.

## METHODS

Classical “query and report” approaches common in relational database environments were initially used to analyze the audit logs generated by the PACS and RIS applications at Maine Medical Center. In fact, both the PACS and RIS vendors included standard “audit reports” in their product offerings. However, these standard reports, and reports created in a similar fashion, are limited in that they assume you know what you’re looking for when you create the report. When analyzing huge arrays of data for “suspicious activity,” the proverbial “needle in a haystack” syndrome applies.

For example, suppose one wanted to determine if an administrative (ie, nonclinician) user of the PACS was inappropriately looking at a coworker’s radiology images or reports. Assuming there was a way to easily correlate user type with the user identifier captured in the audit logs (ie, which users were “administrative users”) and assuming there was a way to know which patients were employees of the hospital, it would be relatively straightforward to build a query that would check the audit logs looking for a match. The query would say something like, “give me all events performed by an administrative user where the patient type was ‘employee.’” Pretty straightforward, except that this entire example assumes that the PACS administrator or security officer *knew* that checking if users were looking at coworkers records was something to investigate. If the PACS administrator were simply staring at 800,000 audit records, would he/she they have thought to build a report to check for this possibility? This probably would not be the first query done by most system administrators.

Multidimensional analysis (aka online analytical processing or OLAP) tools, however, are designed to explore the data to search for hidden or previously unknown connections within and between the data elements. Multidimensional analysis tools partition the data to allow it to be viewed easily from any number of different perspectives. In addition, most multidimensional analysis tools provide excellent graphical tools to assist in the analysis of the data. A simple “double-click” on a high-level graph would

immediately reveal the detail used to derive the higher-level summarized data (ie, “drill-down” into the data). With another double-click, the user could further investigate an even smaller subset of data. Once the smaller subset of data has been examined, the graphical tool allows the user to easily return to the broader dataset (ie, “drill-up”). This power to quickly break a large dataset into smaller pieces and view it from every conceivable perspective facilitates analysis that would be very cumbersome with traditional methods. For example, identifying which users have reviewed studies for a given patient population is a trivial task with a properly configured multidimensional analysis tool.

Therefore, back to the example described above: Using a multidimensional analysis tool, the administrator would start with a summarized view of all the audited events captured by the system. A few double-clicks would drill down (ie, filter) to show events by user type. Table 1 gives an example of what the filtered crosstab view would look like. A few more double-clicks on the “Other Administration” row would show events for administrative users not in the radiology department, and a few more clicks would drill down to a particular administrative user. At that point, that user’s audited events could be quickly reviewed from any number of potentially revealing perspectives. Logins by date, for example, might reveal that the user logged into the system only a few times. A few more clicks might show that the user did several queries for the images and reports on a single day and all of those were for a single patient. Since the data has now been filtered down to the events of a single user with regard to a single patient, analyzing the pertinent information about the patient (eg, patient type, referring physician, exam type, etc) is quick and easy and might reveal why that PACS user might be so interested in looking at the single patient’s record (eg, the user and the patient have the same last name, something we would consider a “suspicious coincidence”).

While not foolproof, it’s easy to see how the capabilities of a multidimensional analysis system greatly enhance the ability of a PACS administrator (or security officer) to analyze audit logs for events that may require additional scrutiny. Effectively logging all auditable events (as required by HIPAA) is only of value if those logs can be (and actually are) reviewed to identify events that require further review. The application of a multidimensional analysis tool fills the niche created by the copious amounts of data captured in the various systems logs.

## RESULTS

At Maine Medical Center, we chose Cognos Incorporated’s *PowerPlay* as our multidimensional analysis tool. The process for loading data into a multidimensional analysis tool will vary depending on the product. With *PowerPlay*, the process involved extracting the data from the originating system into a simple spreadsheet or comma-delimited file and then simply importing that file into *PowerPlay*’s *Transformation Server* software. The initial load

**Table 1. Example Crosstab Showing Number of Web Events for Various Web User Types**

| User type                | Number of web events |                |                |                |                |                |                 |
|--------------------------|----------------------|----------------|----------------|----------------|----------------|----------------|-----------------|
|                          | Logins               | Logouts        | Login failures | Queries        | Studies viewed | Reports viewed | Studies deleted |
| Physicians               | 65,804               | 63,427         | 4,003          | 115,506        | 144,677        | 34,151         | 0               |
| Radiologists             | 1,191                | 1,096          | 111            | 2,744          | 4,154          | 792            | 0               |
| Residents                | 54,826               | 53,140         | 3,606          | 100,954        | 107,413        | 26,905         | 0               |
| PA's                     | 10,537               | 10,239         | 743            | 22,021         | 25,547         | 2,960          | 0               |
| Other clinicians         | 14,201               | 13,634         | 1,426          | 31,653         | 45,881         | 9,920          | 0               |
| Radiology technologists  | 11,076               | 10,118         | 598            | 43,564         | 56,592         | 1,712          | 0               |
| Radiology administration | 4,641                | 4,389          | 414            | 13,019         | 6,223          | 1,975          | 4,596           |
| Other administration     | 999                  | 968            | 162            | 1,910          | 1,143          | 944            | 0               |
| Disabled recounts        | 1,227                | 1,179          | 111            | 2,030          | 3,038          | 374            | 0               |
| <b>Totals</b>            | <b>164,502</b>       | <b>158,190</b> | <b>11,174</b>  | <b>333,401</b> | <b>394,668</b> | <b>79,733</b>  | <b>4,596</b>    |

of the data was very straightforward; we did not perform extensive manipulation of the data in order to import it into the OLAP tool. The data was basically a straight dump of the audit table from the PACS webserver.

For our first test, we loaded our multidimensional analysis tool with the image webserver's event table data. The initial dataset contained about three months of audited events, or about 340,000 records. Once inside the tool, we were able to create a quick *powercube*, OLAP's term for a multidimensional data structure used by the OLAP engine to index its data. The structure of the initial powercube was very basic; it was driven directly from the structure of the underlying event log file taken from the webserver. As described below, however, we subsequently developed more extensive datasets and corresponding powercubes to enhance our analysis capabilities. However, the overall effort to create the initial powercube took only about one day.

Within the first few moments of exploring the initial cube, we were able to discover several "suspicious" events hidden within the audit logs. Almost immediately, the power of the multidimensional analysis tool and its graphical user interface became apparent. In fact, the example used above was not fictitious; it was the exact scenario that we stumbled upon using the multidimensional analysis tool. After only a few moments of using the graphical analysis tool to drill into (ie, filter) and scrutinize the data, we were led to question the activities of several nonclinical users of the system. In particular, a few users who had been given access to PACS for testing, training, and demonstra-

tion purposes had reviewed the radiology images and reports of coworkers and relatives. Presented with the evidence, we were able to confront those individuals and halt the behavior. In fact, once the word got out about our ability to easily analyze our audit logs (ie, that we were actually looking at the information we had claimed we were capturing), the threat of retribution for violating the hospital's confidentiality and security policies became an effective deterrent. Incidents of questionable activity in the PACS by the administrative users have basically disappeared. We suspect that violations of our security and confidentiality policies exist throughout the health care arena, but identifying them without adequate audit logs and the tools to analyze those logs is difficult. The use of the multidimensional analysis tool allowed us to begin to address this issue.

As one might expect, the initial structure of the first powercube required refinement as we became more experienced with the multidimensional analysis tool, and, more importantly, more knowledgeable of the data we were attempting to analyze. This led to an iterative process whereby we would refine the format of the data being imported, then modify the structure of the powercube being created in order to maximize the value of the resulting dataset. Ultimately, we began reaching beyond the web audit log table and into other databases to pull other relevant information into the powercube. For example, we used the accession number stored in the web audit logs to access additional study information from the RIS database. Ultimately, we incorporated data from several hospital and departmental systems

into a mini-departmental data warehouse (see Fig. 1), then we used that warehouse to feed the data into the OLAP software. In Table 2, we depict the various databases and corresponding data elements used to populate the departmental warehouse.

Using the data from the additional data sources, we were able to construct a powercube with "dimensions" from a number of different perspectives. Examples of the various dimensions include user demographics (including where employed, specialty, role, etc), date of exam, location exam was performed, requesting, attending, and performing physicians (including their associated demographics), date and time of user access, and relevant patient demographics. Within each dimension, the data were categorized and structured into "tiers" to enable rollup into higher-level summary categories. Figure 2 shows an example of the structure for the *Web User Type* dimension. The users are categorized into a *User Type* tier and *User Sub-Type* tier. It is the definition of this data model that enables the multidimensional analysis tool to construct the powercube to provide quick real-time graphical analysis capabilities.

Once the enhanced data model was constructed, we were able to analyze our PACS and web utilization using the power of the multidimensional analysis tool. The resulting powercube contained the information we needed to quickly and easily investigate security-related concerns. For example, whenever a resident or physician is no longer privileged at the hospital and no longer requires access to the web imaging system, our normal operational process would be to disable their web and/or PACS accounts. In addition, the credentialing system would pick up the change and indicate such in the data feed to the warehouse. By comparing web user activity on disabled accounts using the OLAP tool, we could ensure that all accounts that should have been disabled actually were disabled.

Other applications of the OLAP tool to the security and confidentiality dilemma included the ability to review all activity performed by a particular user or, conversely, all activity for a particular patient. For example, on several instances, we wanted to know if anyone had

inappropriately reviewed the radiology images or reports of a VIP (very important person). Using the OLAP tool, we simply filtered the data to exclude all audited events except those pertaining to the VIP patient and then reviewed all user activity to verify that all activity was legitimate.

## DISCUSSION

In general, our experience with the application of a multidimensional analysis tool to the problem of scrutinizing audit logs was very positive. As we developed our process for gathering the appropriate data from various sources, building an accurate data model, and utilizing the great user interface of the multidimensional analysis tool to review the data, we recognized a significant advantage of this approach compared with the traditional approach of querying the audit logs to search for particular events. We found things we were not even looking for, something that would probably never happen using a traditional approach.

There were, however, a number of challenges we discovered as a result of our work on this problem. First of all, the slice-and-dice analysis approach still requires human intervention. While we can automate the process of extracting the data from the various departmental systems and updating the powercube on a regular basis (daily, weekly, monthly), we have not found a way to automate the actual analysis. Ideally, we could build a system that would simply report all potential issues on a daily or weekly basis. As we continue to refine our process, we hope to minimize the amount of routine labor required to review the logs, but we do not believe we will ever reach a level of automation that may be possible with a rules-based custom-built database application. Such a system could potentially be constructed to identify and report suspicious events in a more automated fashion requiring little or no human intervention. However, as discussed above, the development of such a system would require some knowledge of what to look for, and thus a multidimensional approach to the initial analysis would be the perfect place to start.

The second challenge we noted was that identifying suspicious activity is not an easy

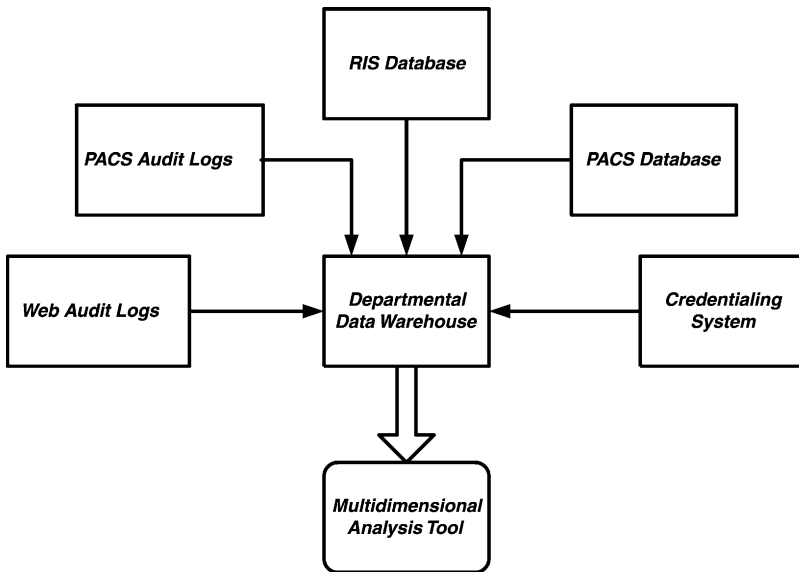


Fig 1. Radiology departmental data warehouse.

task. Even with the power of a fully developed multidimensional analysis tool, trying to identify events that are inappropriate can be very difficult. For a certain population of users, for example, administrative users that use the system only occasionally, reviewing their activity is pretty straightforward. However, for users that review dozens of images and reports every day, it is very difficult, if not impossible, to detect a single or small number of inappropriate activities. At Maine Medical Center, we have approximately 1650 users of the web-based imaging system, so analyzing the daily activity of all of them is not practical.

As a result, in addition to a tool that allows us to easily review what transpired, we must also have some independent knowledge of what activity is legitimate and what is not. If we could somehow come up with a list of all users who had a legitimate need to see a particular patient's record, we could correlate that list with the audit logs to identify outliers. In fact, many information systems attempt to block initial access by limiting access to only those "physicians of record" associated with a particular patient (eg, ordering, attending, primary care). However, coming up with such a list, even after the fact, is a very difficult if not impossible endeavor.

From a clinical viewpoint, the problem of restricting access to the physician of record is

problematic. First, the PACS can know about only one (or at most two) clinicians attached to a particular patient examination. However, most patients, and in particular inpatients at MMC (a teaching facility), have multiple physicians that the PACS and RIS do not know about. Beyond that, there will inevitably be consultants, specialists, or even surgeons called in to help with the care of the patient, and neither the RIS nor PACS can know or predict which physicians (out of the 1000+ credentialed at MMC) will have a legitimate need to view this patient's images. In addition, physicians often have office staff that support the patient's care and have a legitimate need to review the patient's medical record. To arbitrarily restrict access to only those physicians that the PACS and RIS know about would be unacceptable, creating more bottlenecks than were present in the film-based environment.

As a result of the above factors, we have strongly advocated unrestricted access to PACS images, relying instead on robust auditing techniques and the enforcement tools (ie, the threat of loss of credentials or privileges) that MMC already has in place to deal with inappropriate behavior. As a condition of access to our system, all users are trained in accordance with the HIPAA regulations, §164.530 (b) 1: "A covered entity must train all members of its

**Table 2. Data Fields Extracted from Various Departmental Databases**

|                             |                        |                              |
|-----------------------------|------------------------|------------------------------|
| <b>RIS Database</b>         |                        |                              |
| MRN                         | Order Comments         | Patient Name                 |
| Accession Number            | Episode Type           | Patient Address              |
| Study Date                  | Work Area              | Patient DOB                  |
| Study Description           | Org Unit               | Financial Type               |
| Modality                    | Technologist           | Specimen Type                |
| Priority                    | Reading Radiologist    | Procedure Start Date/Time    |
| CPT Code                    | Verifying Radiologist  | Procedure Ordered Date/Time  |
| Admitting Diagnoses         | Verifying Date/Time    | Procedure Complete Date/Time |
| Order Reason                | Reading Resident       |                              |
| <b>PACS Database</b>        |                        |                              |
| Institution Name            | Station Name           | Speciality                   |
| Department                  | Number of Images       | Dictation Date/Time          |
| Study Location              | Study Size             | Transcribed Date/Time        |
| Source                      | Body Part              | Approval Date/Time           |
| <b>Web Audit Logs</b>       |                        |                              |
| Event Date & Time           | Modality               | Reason for Study             |
| Event Type                  | Study Comments         | Requesting Service           |
| User Name                   | Speciality             | Department                   |
| User Type                   | Body Part              | Patient Name                 |
| MRN                         | Study Description      | Patient Sex                  |
| Accession Number            | Requesting Physician   | Patient History              |
| Study Date & Time           | Referring Physician    |                              |
| <b>PACS Audit Logs</b>      |                        |                              |
| Event Date & Time           | Modality               | Reason for Study             |
| Event Type                  | Study Comments         | Requesting Service           |
| User Name                   | Speciality             | Department                   |
| User Type                   | Body Part              | Patient Name                 |
| MRN                         | Study Description      | Patient Sex                  |
| Accession Number            | Requesting Physician   | Patient History              |
| Study Date & Time           | Referring Physician    |                              |
| <b>Credentialing System</b> |                        |                              |
| User Name                   | Current Status         | Expertise                    |
| User Type                   | Degree                 | Secondary Expertise          |
| User SubType                | Privilege Level        | Home City, State, Zip        |
| Date Created                | Department Affiliation | Office Name                  |
| Date on Staff               | Section Affiliation    | Office City, State, Zip      |

workforce on the policies and procedures with respect to protected health information required by this subpart, as necessary and appropriate for the members of the workforce to carry out their function within the covered entity." In addition, each user must sign an agreement indicating they understand and will comply with all institutional security and confidentiality polices. Given these assurances, we grant access, and then follow up that access with review of audit logs to help ensure compliance.

One novel approach suggested by our Chief Information Officer would allow unhindered access to those patient's records where there is a predetermined connection between the patient and the user requesting access. In the instances

where there is no predetermined connection, the user would be reminded of our security policy and asked to provide some type of real-time reason why he/she need access. Using this "break the glass" approach, immediate access would be granted but the event (and the user's response to the access challenge) would be captured in the audit logs and later reviewed to determine legitimacy.

There are a couple of challenges with this approach. First, we must still be able to pre-determine the access needs for a high percentage of users (in order to avoid constant challenges and subsequent impact to streamlined patient care), which, for the reasons mentioned above, is a very difficult problem. Second, such an approach would require our PACS vendor to

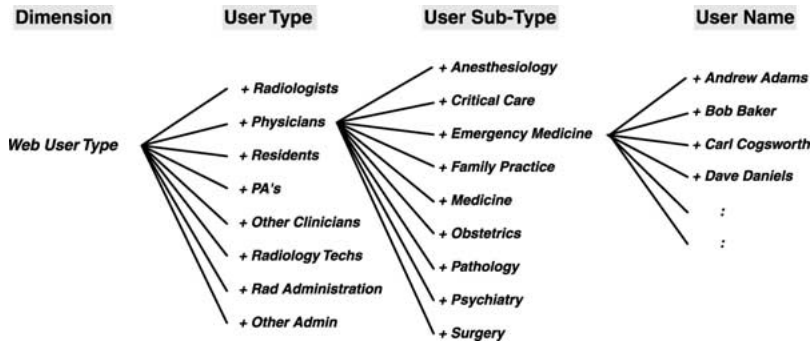


Fig 2. Category tree for *Web User Type* dimension.

offer this capability as part of the system's auditing capabilities, something we have not yet been successful in convincing the vendor to do. Hence, at least for now, attempting to limit access to a particular patient's record based on a predetermined knowledge of who needs that access is not practical and would certainly negatively impact patient care. Instead, we must assume a user requesting access has a legitimate need-to-know and grant that access to facilitate superior patient care.

One final note regarding security and confidentiality of patient data: Why would we go through all this trouble to attempt to identify inappropriate activity? What harm is there if, for example, a resident (not involved in the case) reviews the chest image of a high-profile accident victim just because he is curious? Why would we care? First of all, HIPAA regulations require we do what we can to safeguard a patient's information: "A covered entity must reasonably safeguard protected health information from any intentional or unintentional use or disclosure that is in violation of the standards, implementation specifications or other requirements of this subpart." ( §164.530 (c) 2). Second, even in the absence of HIPAA, there are other concerns that could arise that would suggest we review the audit logs for questionable activities. For example, it is conceivable that a user could do marketing research by identifying all exams that were performed for a competing practice and then attempt to create a competitive advantage from that analysis. While clearly not appropriate, this type of activity would not bode well with most hospital administrators.

Overall, the multidimensional analysis approach has greatly enhanced our ability to take advantage of the data being captured in our audit logs to improve the security and confidentiality of patient data at Maine Medical Center. While limitations of the system exist, we plan on continuing to work toward addressing those constraints. In the end, we believe multidimensional analysis will be one of the key tools used to ensure compliance with our HIPAA-backed policies and procedures with respect to protecting health information.

#### ADDITIONAL REWARDS

Beyond the benefit of the application of multidimensional analysis tools to the problem of analyzing audit logs, a perhaps even greater benefit can be realized when the same tools are applied to other aspects of the radiology department's workflow. At Maine Medical, we quickly moved beyond just analyzing the audit log data for security purposes and began looking at many other aspects of our system that had previously been undiscovered. Once the data from the departmental data warehouse had been loaded into the multidimensional analysis system, we had created a powerful tool for analyzing the department's workflow.

For example, from the webserver data we could now easily analyze and graph web activity by a number of independent dimensions including web user, user type (eg, physician, resident, technologist, clinician), user privilege level (eg, radiologist, administration), modality, study location, study date, web access date, type of access, ordering physician, requesting ser-

**Table 3. Percentage of Exams Transcribed within a Certain Time, Grouped by Modality**

| Modality         | Time from Ordered to Transcribed |               |               |               |               |              |              |
|------------------|----------------------------------|---------------|---------------|---------------|---------------|--------------|--------------|
|                  | <1 h                             | 1-2 h         | 2-4 h         | 4-8 h         | 8-24 h        | 24-48 h      | >48 h        |
| CT               | 28.0%                            | 13.3          | 14.0          | 10.8          | 21.3          | 8.4          | 4.3          |
| MR               | 12.2                             | 18.9          | 13.1          | 11.0          | 23.1          | 12.7         | 9.0          |
| AN               | 1.5                              | 11.7          | 24.8          | 16.6          | 21.6          | 13.4         | 10.4         |
| US               | 26.2                             | 14.8          | 13.0          | 12.4          | 19.9          | 9.3          | 4.5          |
| MM               | 81.1                             | 16.0          | 2.5           | 0.3           | 0.1           | 0.1          | 0.0          |
| NM               | 4.6                              | 9.6           | 23.1          | 20.4          | 20.1          | 15.2         | 7.0          |
| RF               | 14.9                             | 21.7          | 16.8          | 14.0          | 18.0          | 9.8          | 4.9          |
| DX               | 30.9                             | 11.1          | 9.9           | 10.2          | 24.7          | 9.5          | 3.8          |
| <b>All exams</b> | <b>31.57%</b>                    | <b>13.09%</b> | <b>11.31%</b> | <b>10.20%</b> | <b>20.80%</b> | <b>8.87%</b> | <b>4.17%</b> |

**Table 4. Percentage of Exams Transcribed Within a Certain Time, Grouped by Day of Week**

| Day             | Time from Ordered to Transcribed |               |               |               |               |              |              |
|-----------------|----------------------------------|---------------|---------------|---------------|---------------|--------------|--------------|
|                 | <1 h                             | 1-2 h         | 2-4 h         | 4-8 h         | 8-24 h        | 24-48 h      | >48 h        |
| Sunday          | 25.3%                            | 9.5           | 6.6           | 6.0           | 31.6          | 17.3         | 3.7          |
| Monday          | 32.1                             | 13.6          | 12.1          | 11.4          | 19.7          | 8.4          | 2.7          |
| Tuesday         | 31.9                             | 14.7          | 13.4          | 10.4          | 20.5          | 7.1          | 2.0          |
| Wednesday       | 34.1                             | 13.3          | 12.9          | 10.5          | 20.1          | 6.6          | 2.5          |
| Thursday        | 31.6                             | 13.4          | 11.2          | 11.3          | 22.4          | 7.1          | 3.0          |
| Friday          | 32.4                             | 12.9          | 10.6          | 11.0          | 16.4          | 8.0          | 8.6          |
| Saturday        | 26.2                             | 10.0          | 6.3           | 4.7           | 23.9          | 19.6         | 9.3          |
| <b>All days</b> | <b>31.57%</b>                    | <b>13.09%</b> | <b>11.31%</b> | <b>10.20%</b> | <b>20.80%</b> | <b>8.87%</b> | <b>4.17%</b> |

vice, and even patient name. Esoteric questions like "In the month of June, for which procedure codes were the orthopedists most likely to look at both the MR images and the associated report, and which procedures were they most likely to only look at the images" were now easily answered. We were able to start studying utilization of our image server in a way that had never before been possible.

One of the most significant endeavors we tackled was the question of report turnaround time (ie, total time from order placement to final imaging report availability), one of the key benchmarks of the imaging department's overall performance. By extracting data from the various systems, we were able to build a powercube that revealed the duration of each of the steps contributing to the overall report turnaround time. We captured a number of time stamps as part of this analysis:

Ordered in RIS

Completed by Technologists in RIS

First Image Arrived in PACS

Final Image Arrived in PACS

Dictation Started

Dictation Completed

Preliminary Report Transcribed

Final Approval by Radiologist

From these time stamps, we easily computed the durations between each of the steps. Once the powercube was constructed, we were able to graph the average duration for each of these steps. Further refinement of the powercube allowed us to build histograms to more clearly reveal how well the department was performing. Tables 3 and 4 give examples of the type of information we were able to derive from this analysis.

It is interesting to note that an average turnaround time benchmark had been measured and reported in the department for years. However, never before had we been able to report it with such accuracy or with a more meaningful indicator such as percentage of cases transcribed within 24 hours. Likewise, what were once just a few "high-level" summary items could now be studied using the power of the multidimensional analysis tool. We could, for example, review transcription turnaround time for outpatient CT studies done on Tuesday and dictated by a



particular radiologist. There was almost no limit to the way in which we could analyze the data. We could now better understand what factors were most significantly contributing to our overall report turnaround time.

It is also interesting to point out that, in the past, there was always some suspicion as to the accuracy of the summary data being reported, and rightly so. Given the volume of data being processed, and given the difficulty in accurately computing the desired benchmarks, proving the accuracy of the reports was difficult if not impossible. With the multidimensional analysis system, though, errors in the data become very apparent. As one drills down into the data so that a smaller and smaller subset of data is reviewed, hidden idiosyncrasies or even errors jump out very quickly. For example, one of the factors that was contributing to the inaccuracies of our overall report turnaround time was the routine practice of recreating orders for exams that had been done but for which the original order had been canceled (for some reason). In these instances, the time between order creation and order completion was often a huge negative value (sometimes many days). The contribution of this to the overall average report turnaround time was significant. With the multidimensional analysis system, these types of errors became readily apparent and could be accounted for in preprocessing of the data and the construction of the powercube.

Overall, we identified several of these types of problems in our prior approach. Another example showed that we had a one-hour discrepancy in some of our reported time stamps. As it turned out, we found a bug in some of the extraction software that had not properly accounted for daylight savings time. It is quite conceivable that this error would have gone undetected for some time had we not had the power of a multidimensional analysis system at our fingertips.

The report turnaround time study was just one of the many things we have reviewed with the multidimensional analysis system. Other things such as the growth of the average studies

(ie, average number of images per study) or even network performance are relatively easy to review with the system. In the future, we hope to add financial data to monitor departmental fiscal performance, and even build a powercube of detailed database events to aid in ongoing system performance tuning. The extensive power of the system is probably best summarized by a statement from our administrative director when he said, "This tool has a lot of answers looking for questions!"

## CONCLUSION

Classical reports extracted from conventional relational databases cannot keep up with the demands of the modern day department administrator because those reports are bounded by one's ability to properly formulate a question in the context of a query. Audit logs and other large datasets must be analyzed using a different approach. OLAP allows the end user to interact with the data in a way that allows real-time movement from one question to another that's triggered by the answer to the first. Using multidimensional analysis, extracting real information from millions of records of data in departmental databases has become a tangible goal. In just a few weeks, multidimensional analysis of Maine Medical Center's PACS and RIS data yielded exceptional value. Using a multidimensional analysis OLAP tool costing less \$15,000 (less than what we pay for two weeks of support from our PACS vendor), we were able to quickly and easily answer some crucial questions regarding our systems. The original goal of finding a way to improve the process of analyzing our audit logs was met, and quickly surpassed, as we made new and exciting discoveries regarding the power of multidimensional data analysis when applied to the departmental databases. As we move to operationalize the multidimensional approach to our departmental reporting, we hope to add additional datasets (ie, dimensions and powercubes) to further pursue the veiled information obscured in our vast data arrays.