

Benefits of Using the DCM4CHE DICOM Archive

Max J. Warnock, Christopher Toland, Damien Evans, Bill Wallace, and Paul Nagy

KEY WORDS: Integrating Health Care Enterprise (IHE), Digital Imaging and Communications in Medicine (DICOM), digital image management, PACS administration, PACS DICOM IHE conformance

INTRODUCTION

With advances in digital imaging throughout the 1990s and the rapid adoption of Picture Archiving and Communication Systems (PACS) in the last decade, standards for information exchange have become crucial to effective communication, both within the radiology department and with the larger enterprise and outside institutions and agencies. The original Digital Imaging and Communications in Medicine¹ (DICOM) standard was introduced in 1993 and, in the intervening years, has been widely adopted. Interest in DICOM, as well as the continuous integration of new technologies and modalities into medical imaging, has resulted in a series of new and revised DICOM standards. Along with the growing complexity of these standards has come a need for tools that can manipulate and store DICOM information.

In 2000, responding to these needs, JDicom, a toolkit written in the Java programming language, was developed for manipulating DICOM. The popularity of this new tool kit persuaded its lead

developer to build a full-featured DICOM archive. The mission of the project was to produce a DICOM archive that was free, open-source, and cross-platform and that embraced new directions being drawn up by the Integrating the Health Care Enterprise (IHE) initiative. This ambitious goal attracted more developers, and the DCM4CHE² project was born. After 7 years of development, the DCM4CHE project has produced two generations of a DICOM archive. The current generation is the result of learned practical experience and reflects the old programmer's adage: "Build it twice, because you will anyway." This article reviews DCM4CHE and the DCM4CHE DICOM archive (DCM4CHEE) and evaluates their maturation as a viable platform for training, integration testing, and research.

ARCHITECTURE AND FEATURES

The DCM4CHEE architecture was designed to be modular (Fig. 1), allowing for independent use of the archive's services. Each service can be enabled/disabled using the DCM4CHEE JMX console (web interface). This modular design also lightens the load of maintaining old code and adding new features/services.

From the Diagnostic Radiology, University of Maryland School of Medicine, 22 South Greene Str., Baltimore, MD, 21201, USA.

Correspondence to: Max J. Warnock; e-mail: mwarnock@umm.edu

Copyright © 2007 by Society for Imaging Informatics in Medicine

Online publication 5 October 2007

doi: 10.1007/s10278-007-9064-1

¹The DICOM Standard web site. Available at <http://medical.nema.org/>. Accessed on June 6, 2007.

²The DCM4CHE Project History web site. Available at <http://www.dcm4che.org/confluence/display/proj/history>. Accessed on June 6, 2007.

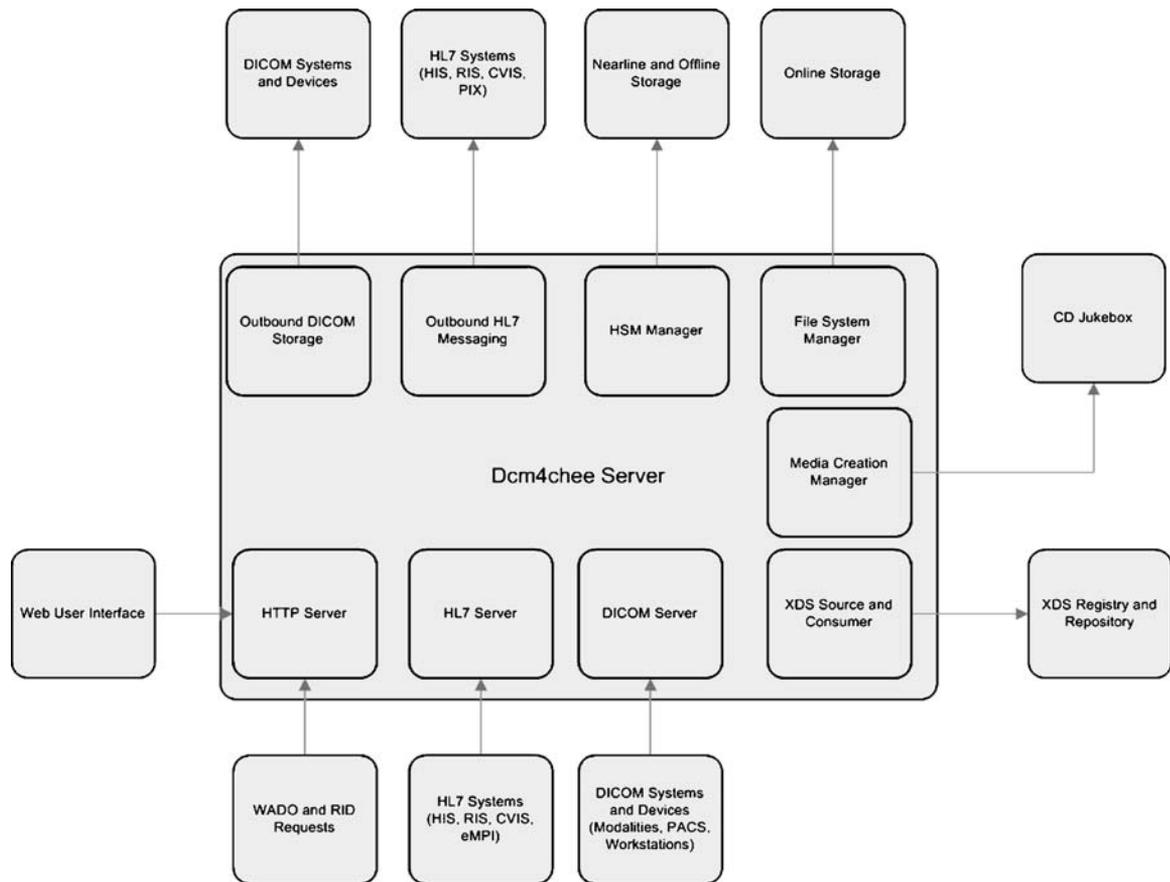


Fig 1. Graphical representation of the modular DCM4CHEE architecture.

As a result of this architecture, DCM4CHEE features have grown up in pillars, encompassing the standards [DICOM, Health Level 7 (HL7), etc.] with which they interact. Individual features of the archive are listed under their corresponding standards in Table 1. Each of these services represents a part of routine workflow in a radiology department. DCM4CHEE exposes and uses these services as outlined in IHE.³

DCM4CHE as a Training Platform

Perhaps the most daunting task for any new PACS administrator is understanding and gaining experience with the DICOM standard. DICOM is a binary run-length encoded standard developed

solely for medical imaging and can require hexadecimal editors merely to read. DICOM is perhaps the most important single area that a PACS administrator must master, and this mastery increases in difficulty as the DICOM standards grow in numbers and complexity. The most effective way to learn DICOM is through a practical, hands-on approach, which DCM4CHEE provides.

DCM4CHEE is easily set up and includes comprehensive logging of DICOM transactions. The archive comes in a precompiled (binary) version designed for two open-source and readily available databases, MySQL and POSTGRES. The package available on sourceforge.net⁴ comes with directions for setting up the archive on multiple

³A web site detailing the IHE profiles supported by DCM4CHEE. Available at <http://www.dcm4che.org/confluence/display/ee2/IHE+Integration+Statement>. Accessed on June 6, 2007.

⁴The DCM4CHE project file archive. Available at <http://sourceforge.net/projects/dcm4che/>. Accessed on June 6, 2007.

Table 1. DCM4CHEE Standards and Features

Standard Supported	Feature
DICOM	C-ECHO SCP/SCU (ping) C-STORE SCP/SCU (storage) C-FIND SCP (query) C-MOVE SCP (retrieve) N-ACTION SCP/SCU (storage commitment) Routing of images based on any DICOM tag General Purpose Worklist SCP Modality Worklist SCP Modality Performed Procedure Step SCP Hanging Protocol SCP Instance Availability Notification SCU Teaching File and Clinical Trial Export Manager
HL7	Inbound ADT (patient demographics) Inbound ORM (imaging exam orders) Inbound ORU (diagnostic results) PIX Query (patient identifier resolution) Prefetching of DICOM images based on ORM messages HL7 message routing
Storage	Online: capable of multiple file systems with MD5 verification of content Nearline: integration with an HSM or offline storage system CD media: media creation manager for writing to CD DICOM Dir or date/time-based directory structure
Web-based services	HTTP access to DICOM objects and reports (via WADO and RID services) Web-based administration user interface Search for studies Browse worklists Configure DICOM destinations Configuration View Audit logs
Cross-enterprise document and image sharing (XDS)	Image information source Document repository
Auditing	Key object generation source Audit logging for all transactions (according to IHE) ARR

ADT = admission/discharge transfer order; ORM = order management; ORU = observation results unsolicited

operating systems.⁵ The depth of the installation documentation provides one example of the maturity of DCM4CHEE. Its web-based graphical user interface (GUI) provides the ability to search the

archive and offers a number of other operations that can be performed without the need of in-depth DICOM knowledge. The log files allow the user to view the DICOM transactions produced by the GUI. By viewing the log of the GUI transactions, even a novice can glean the required information for DICOM transactions and, in the process, gain insight into how the standard works. Setting up and using the archive and experimenting with the toolkit is a productive way to gain experience with the DICOM standard.

As the user's PACS administration and testing needs grow more complex, the DICOM toolkit on which DCM4CHEE is built becomes a valuable asset. Knowing the DICOM terminology and having an understanding of how and when messages are sent opens new perspectives on ways to validate processes and test the PACS environment. Although the current documentation for the tool kit may present some challenges to the novice administrator, those with DICOM experience will find it a useful asset. Gaining experience with the tool kit will not only help in learning DICOM but can be a critical tool for troubleshooting systems.

Despite the fact that DICOM is a widely adopted standard for medical imaging, it is not an all-encompassing solution for integration and transaction issues in all health care environments. DCM4CHEE also supports HL7 admission/discharge transfer order, observation results unsolicited, and order management messaging, as well as syslog audit record messages and ebXML in cross-document sharing of the IHE. DCM4CHEE's compliance with the IHE profiles helps to round out a PACS administrator's vision of clinical workflow in medicine. Having an IHE-compliant archive with which to experiment provides valuable hands-on experience. DCM4CHEE makes examining each step of the major processes of a radiology practice straightforward with verbose logging, as well as an audit record repository (ARR). These features aid in knowing what should be happening at each step of clinical workflow and thereby assist in securing faster diagnoses and more timely resolutions of future systems failures.

A Workbench for Testing

Because it is designed around IHE, DCM4CHEE also serves as a platform for testing new and/or existing applications in clinical systems. When

⁵The DCM4CHE Project History web site. Available at <http://www.dcm4che.org/confluence/display/proj/history>. Accessed on June 6, 2007.

adding a new modality in the private practice or hospital setting, DCM4CHEE can be used to test modality worklist, one of the IHE profiles, outside of the routine PACS workflow. The Modality Performed Procedure Step, Cross Patient Identifier (IHE PIX), and HL7 servers allow for testing of all the integration features of imaging modalities. As new PACS vendors appear and established PACS vendors release new versions, DCM4CHEE can be used to quickly test for IHE compliance. DCM4CHEE also comes with an ARR (enabled by default). An ARR is a log of all transactions between the archive and other actors, including modalities, viewers, and other archives. This repository is viewed and made searchable by the DCM4CHEE's web GUI and is a good source of information on how DICOM transactions are working in the system. It records detailed information about the sources exchanging data, as well as metadata on the information being sent. When using the archive, the ARR is a wealth of information about system utilization and also serves as an excellent health insurance portability and accountability compliance log.

As IHE becomes more pervasive in the health care industry and other initiatives work to provide new standards for sharing patient data across the enterprise, IHE cross-enterprise document sharing (XDS) profiles will become more important. DCM4CHEE has already incorporated the XDS-I IHE profile for image sharing into the archive. As new XDS repositories and sources become available on the network, these can be tested using the DCM4CHEE GUI tools for XDS-I. The user can simply select a study to be sent to the XDS-I source, select application entity title from the drop-down list, and click the XDS-I icon. Again, by examining the log file, the interaction between the two servers can be seen, as well as any errors that might have occurred, so that additional configuration can be performed as needed before the XDS repository is put into production.

Most administrators fear changes in their systems. "If it's not broke, don't fix it" is the general attitude and understandably so. PACS is no longer a novelty but a mission-critical driver in medical imaging workflow. Administrators do not have the luxury of letting the system fail today and worrying about fixes tomorrow. A key part of all proactive change management is proper testing. DCM4CHEE can be used for a substantial portion

of the testing required to ensure stability of clinical systems.

Teaching Archive and Research Platform

In many departments, research and teaching activities are second-tier funding considerations, budgeted separately from (and much less liberally than) clinical systems. Because DCM4CHEE is both free and stable, it can be used as an archive at the cost of the hardware on which it runs. The true value of today's growing number of electronic teaching/research files lies not in the DICOM study alone but also in the annotations, key images, and presentation state. IHE has defined the Teaching File and Clinical Export (TCE) profile, which includes a standardized way for PACS systems to share this information. DCM4CHEE supports importing DICOM, key object notes, annotations, and associated documents as outlined in the TCE profile. A radiologist can annotate a study in a PACS that supports the TCE profile and export it directly to a DCM4CHEE research database, retaining all the pertinent information – as well as customized teaching notes – provided by the radiologist. Moreover, the research archive's database can be queried to gain additional insight into and identify patterns in the metadata associated with research studies. An added benefit of using a research archive is that there is no degradation in the performance of the clinical PACS when researchers retrieve studies or perform complex database queries. Segregating studies into a research archive ensures good performance and provides a test bed for faster and easier research.

Simple study segregation and querying of the archive are not the only types of research that can be performed using DCM4CHEE. The archive can be used as a foundation for PACS-based product development. One example of a project using DCM4CHEE as a base for testing and development is Xero, a truly thin-client viewer for medical imaging.⁶ The primary goal of Xero is to have a full-featured PACS viewer that works in any major web browser without additional plug-ins or programs installed on the client. Xero utilizes web access for DICOM objects (WADO) and web

⁶The web site for the Xero project. Available at <http://www.dcm4che.org/confluence/display/ee2/Xero>. Accessed on June 6, 2007.

technologies such as asynchronous JavaScript and XML (AJAX). WADO is a standardized way of sending DICOM objects over hypertext transfer protocol (HTTP). The standard allows DICOM images or parts of images to be sent upon request. By using AJAX to make HTTP requests without refreshing the web page, Xero provides native, in-browser, window/level, pan/zoom, and stacking features for viewing images. Xero is being developed using DCM4CHEE as a DICOM archive and uses the DCM4CHE toolkit for many of its internal tasks. However, the product is not tied to DCM4CHEE because it integrates using the WADO standard. Any IHE-compliant archive with a WADO cache can be the source for the Xero viewer.

Using DCM4CHE as a base server for IHE-driven health care products is a major benefit to researchers and developers. It allows them to focus on their products rather than being forced to build proprietary archives for testing, an approach that can lead to optimizations that make the initial product dependent on the ad hoc archives and stray from IHE initiatives.

As a result of open-source tools and intense community-based interest in the relevant standards, many projects are utilizing DCM4CHE.⁷ These ancillary communities ensure the continued success of DCM4CHE by building up the community of users, contributing to documentation, and creating experts who can teach the next generation of PACS administrators and researchers.

A Community of Support

As with all open source projects, the success and adoption of DCM4CHE depend on the community surrounding the code. The DCM4CHE community has an active core team of developers and continues to gain support in the form of bug fixes, wiki pages, and mailing list posts.

The wiki⁸ is a good source of information for getting started with DCM4CHE and in understanding the intricacies of the archive and toolkit. However, when looking for help, it is important

⁷Listing of DCM4CHE based projects and users. Available at <http://www.dcm4che.org/confluence/pages/viewpage.action?pageId=393>. Accessed on June 6, 2007.

⁸The DCM4CHE Wiki. Available at <http://www.dcm4che.org/confluence/display/proj/The+Project>. Accessed on June 6, 2007.

to remember that the wiki is neither a tutorial nor a forum. It does not have a simple sitemap or navigation but, instead, is driven primarily by the search feature. When troubleshooting or looking for answers, it may be most helpful to consult the wiki first.

For more active support on points not well documented, the user should turn to the mailing list.⁹ This electronic list is quite active, and queries usually receive quick responses. One benefit of having a worldwide community is that assistance is available at any time, day or night. Subscribing to the mailing list also allows the user to contribute to the community by helping others with problems. By following the mailing list, users may also find answers to problems that have not yet been addressed in their own systems. Users may also participate by creating new pages on the DCM4CHE wiki with their own solutions for common problems. Code is not the only contribution that is made in the open-source community. The synergistic act of simply participating as a user in the DCM4CHE project is one method of support, and this support can be extended by contributing advice based on experience in the form of wiki pages and assistance to less experienced users on the mailing list.

CONCLUSION

The DCM4CHE project has survived and prospered to produce a powerful, stable, and feature-rich archive. It is a tool with which both PACS administrators and researchers should be familiar. It is a major building block with which the health care community can work to fix problems, learn standards, and challenge vendors to become better integrated and more open. The future focus of DCM4CHE will be shaped entirely by the experience and feedback of its users. This participatory growth is likely to produce an even more useful and flexible platform.

ACKNOWLEDGMENT

I would like to thank Dr. Nancy Knight from the University of Maryland School of Medicine for her assistance in preparing this manuscript.

⁹The DCM4CHE mailing list subscription web site. Available at http://sourceforge.net/mail/?group_id=37982. Accessed on June 6, 2007.