

# Integrating Clinical Trial Imaging Data Resources Using Service-Oriented Architecture and Grid Computing

Stefan Baumann El-Ghatta · Thierry Cladé ·  
Joshua C. Snyder

Published online: 7 May 2010

© The Author(s) 2010. This article is published with open access at Springerlink.com

**Abstract** Clinical trials which use imaging typically require data management and workflow integration across several parties. We identify opportunities for all parties involved to realize benefits with a modular interoperability model based on service-oriented architecture and grid computing principles. We discuss middleware products for implementation of this model, and propose caGrid as an ideal candidate due to its healthcare focus; free, open source license; and mature developer tools and support.

**Keywords** Service-oriented architecture · SOA · Grid computing · Clinical trial · Clinical imaging trial · caBIG® · caGrid

---

S. B. El-Ghatta (✉)  
Biomarker Development / Clinical Imaging,  
Novartis Institutes for BioMedical Research,  
Lichtstrasse 35,  
4056 Basel, Switzerland  
e-mail: stefan.baumann@novartis.com

T. Cladé  
IT & Automation, Novartis Institutes for BioMedical Research,  
Lichtstrasse 35,  
4056 Basel, Switzerland

J. C. Snyder  
Biomarker Development / Clinical Imaging,  
Novartis Institutes for BioMedical Research,  
220 Massachusetts Avenue,  
Cambridge, MA 02139, USA

## Background

Imaging plays an increasingly important role in the development of neuropharmacological drugs. 10 out of 106 New Drug Applications (NDAs) approved in the Division of Neuropharmacological Drug Products at the Food and Drug Administration (FDA) between 1995 and 2004 had imaging studies (Uppoor et al. 2008). Prominent examples for imaging used in neuroscience trials are brain volumetrics for Alzheimers disease, PET imaging of receptor occupancy, and Gadolinium-enhanced MRI for Multiple Sclerosis. Service-oriented architecture (SOA) and grid computing can bring a new level of efficiency to these clinical imaging trials. While academic studies are often performed by small, integrated teams, large-scale trials are characterized by distributed operations and oversight. A single clinical trial might involve the sponsor, such as a pharmaceutical or a contrast agent manufacturer; hundreds of imaging sites; a core lab or contract research organization (CRO) tasked with operational control; central radiologists tasked with data analysis; and a data center for image archiving. The sponsor alone might comprise several concerned organizational units, such as separate groups in charge of image and image-derived data management. Each of these parties uses its own systems and processes. There are no widely supported standards which enable, for example, decentralized data storage and access, study protocol definition, scanning sequence specification, or removal of patient-identifying information. Some of these capabilities are unavailable to imaging trial teams today. Others are offered by service providers, but are rarely productized so as to be configurable per trial or client. Even relatively straight forward workflows can suffer from lack

of transparent, configurable, trackable activities, such as quality control. An analysis application could have stringent input image requirements. If data aren't checked for compliance soon after acquisition, they may not be usable. The quality control offering of a CRO is often vaguely specified, implemented in a closed system, and unavailable for external queries. SOA is a broadly accepted approach for modularizing a landscape of such systems into discrete technical services, each with a public, discoverable interface, and composing those services into applications.

With an SOA framework, there would be further opportunities for efficiency improvements. For example, when datasets can be retrieved on demand from an external organization, each trial party can be considered a remote data storage resource. Federating access to the pool of these resources would allow all image transfers to happen as needed, preventing redundant transmission of bandwidth hungry datasets. This resource sharing paradigm is known generically as grid computing. caGrid<sup>1</sup> is an open source middleware product providing implementations of web service standards, as well as a set of grid-enabling core services. By adopting an SOA approach with caGrid, imaging trial parties can achieve a new level in the flexibility of their workflow integration, efficiency of their operations, and quality of their outputs.

## Opportunities

Four main roles are involved in a traditional clinical imaging trial: the company sponsor, the imaging CRO, the central reader or core labs, and the imaging data collection sites. Workflows typically span parties in a standard configuration as depicted in Fig. 1. The sponsor sets up a high level acquisition guideline with the CRO, and the sites generate the actual scans, then send the images back to the CRO. The CRO collates, de-identifies, and sends the images out to a central reader for analysis, or in some cases the roles of CRO and reader are taken by a single academic core lab. The reader sends analysis results to the sponsor for internal management and use.

In this design, the sponsor sends out a high-level study protocol definition and receives tables of clinical trial results. Intermediate research steps essentially occur in a “closed box” under the purview of the CRO, as depicted in the figure below.

<sup>1</sup> This product includes software developed by the Ohio State University Research Foundation (“OSURF”), Argonne National Laboratory (“ANL”), SemanticBits LLC (“SemanticBits”), and Ekagra Software Technologies Ltd. (“Ekagra”) as described in the license accessible at the web page <http://cagrid.org/display/downloads/caGrid+1.3+License>.

This closed box workflow is not modular. It is difficult to include functions from multiple CROs and analysis service providers in the same trial, and to re-use integration infrastructure across trials. This can compromise research quality and result in inefficiencies in the conduct of imaging trials. Exposing service interfaces to the steps inside the closed box involves addressing a number of technology, standards, and regulatory gaps. Specific opportunities for progress include:

### Source Data Management

FDA regulations pertaining to clinical trials source data management were originally written in reference to a world of paper documentation, and have areas of ambiguity when it comes to electronic data. Increasing prevalence of electronic data in trials has motivated the Electronic Source Data Interchange (eSDI) Group, an initiative seeking to “investigate the use of electronic technology in the context of existing regulations for the collection of eSource [electronic source] data (including that from eDiaries, EHR, EDC) in clinical trials for regulatory submission” (CDISC 2006). eSDI and other recommendations do not address certain issues specific to image data, however, leaving gaps in our understanding of how best to ensure patient privacy and image readability while also maintaining source images unchanged.

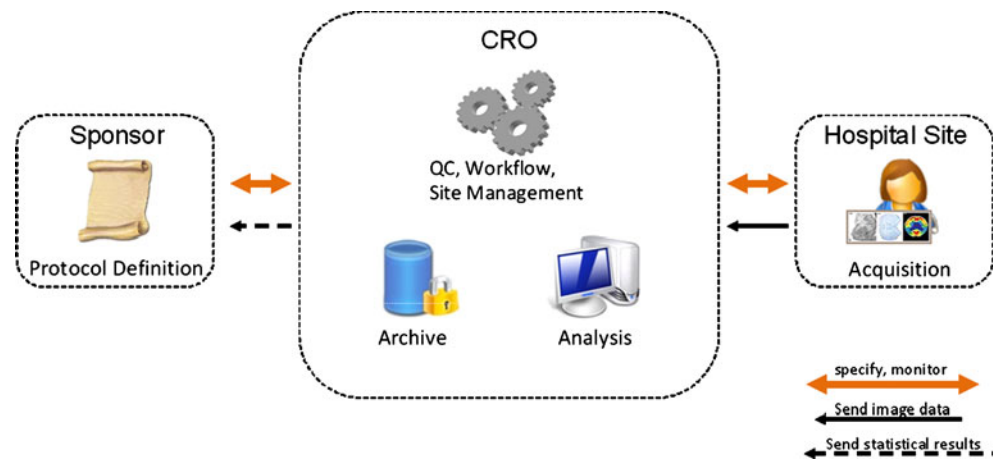
### Data Accessibility

There is no standard image transfer product designed to connect parties outside their respective intranets. Transfer from site to CRO by courier on physical media is a typical practice, causing delays and related QC problems. File Transfer Protocol (FTP) is sometimes used, but in itself is not compliant with FDA regulations as described in 21 CFR 11 regarding electronic records; it does not provide a fault-tolerant transfer mechanism complete with verification, user provisioning, and audit trail, nor is it easily integrated into automated workflows (FDA 1997).

Sponsor access to CRO-managed images is equally restricted. Inaccessibility to CRO-maintained primary imaging data can be a major research impediment. With images traditionally trapped in data silos at CROs, a sponsor faces overhead in achieving a number of goals requiring fast image query and retrieval:

- Provision of data for health authority review
- In-house analysis for methodology development or validation
- Visual review of images and annotations by CROs and sponsors (e.g. for consensus reading)
- Provision to secondary analysis services
- Cost efficient archiving

**Fig. 1** Standard workflow for a clinical trial using imaging



### De-identification

A medical imaging data standard, Digital Imaging and Communications in Medicine (DICOM®),<sup>2</sup> states the need for image de-identification in its draft Supplement 142: “In clinical trials, images are often acquired during the course of clinical care, in which case the patient’s individually identifiable information needs to be removed to protect the patient’s privacy. In addition, there is often a need to remove other information not directly related to the patient’s identity per se, but which might assist in recovering their identity or bias the image interpretation in some way. Conversely, it is important to preserve certain specific information for quality control and analysis that is essential to the conduct of the clinical trial, which might otherwise be removed.” Imaging sites may lack tools to perform the de-identification, pushing the responsibility onto the CROs by necessity. Lacking standard de-identification profiles, the sites risk removing too little or too much information.

### QC

CROs are often unable to perform real time QC checks. This commonly occurs when imaging data are sent via courier. Days pass between the image scan acquisition and the QC process at the CRO, and if the image data are found to be unusable, it may be too late for the CRO to ask the imaging site to correct the errors (e.g. if a re-scan of the subject is required). Along with improved image transfer, automated assessment and feedback mechanisms to avert this scenario are needed.

Furthermore, a lack of standards and tools for describing and enforcing quality requirements often limits sponsors

from conducting in-house or secondary analyses. Newly available, highly automated analysis algorithms can have stringent requirements on input image quality. These are difficult to enforce solely by contract if the CRO performing QC is not also responsible for all aspects of the analysis.

### Analysis

Specialized analysis services are often available only from third parties who do not also provide trial operational support. For example, a sponsor might need to employ a proprietary image analysis service offered by a specific provider, and also a CRO to manage the trial operations and the analysis workflow. In this case, developing and validating interfaces to integrate the analysis service into the CRO workflow could cause months of delay for a one-off solution which could be useless in the context of other trials.

This section has begun to demonstrate some of the opportunities for information technology to enhance imaging trial workflow. The next section will discuss specific implementation recommendations.

### Recommendations

Software applications for clinical imaging trials should become service-oriented. The benefit of composable, reusable services to all parties is ease of integration. Sites benefit from a faster, more user-friendly process for submitting images to a clinical trial; CROs gain a wider market for their services by selling them for use in trials they aren’t otherwise supporting; software providers similarly can sell their products as plug-ins for trial workflows; and sponsors enjoy flexibility in mixing and matching services from different parties. Using a grid computing approach to optimize data sharing further minimizes time

<sup>2</sup> DICOM® is the registered trademark of the National Electrical Manufacturers Association for its standards publications relating to digital communications of medical information.

consuming, bandwidth heavy image transfer operations between all parties.

Specific benefits can be realized in several areas:

#### Mix and Match Workflows

Taking full advantage of widespread SOA adoption, a sponsor should be able to freely mix and match the following services from CROs, manufacturers, and software companies:

- operations (site initiation and management, monitoring)
- scan acquisition definition
- image de-identification and transfer
- image tracking
- image QC (on-site or central)
- analysis/processing
- reporting
- image viewing (both remotely or locally)
- long-term image archiving

By freely composing these into workflows, a sponsor may support significantly more complex scenarios, such as:

##### *Imaging based central safety read:*

Fast image transport would be possible between all study partners. This would enable real time QC and fast reads, necessary in a variety of scenarios, e.g. when safety signals are read by a secondary service provider.

##### *Small exploratory study involving post-processing activity without an imaging CRO:*

Direct transfer of images to the sponsor's image management system would be possible with minimal human intervention once challenges of sufficient de-identification, correct data formatting, correct association of study metadata with images (i.e. subject identifier, time point description, etc.) and reliable transfer are addressed. In-house processing capabilities could then be leveraged, provided that such use is specifically described in the study protocol and the signed patient Informed Consent (IC).<sup>3</sup>

##### *Central reader use of third party applications:*

In a single trial, an analysis service provider would analyze images by using a specialized third party application owned by the sponsor or a software provider without undertaking expensive integration efforts.

##### *Harmonized image normalization across studies:*

Different studies would use different analysis providers but share the same pre-processing pipelines for image normalization.

Figure 2 demonstrates the flexibility of SOA based workflows by comparison to the process depicted in Fig. 1.

#### Source Data Management

The eSDI provides broadly accepted recommendations for managing electronic source data. We want to apply these to imaging scenarios. In contrast with most electronic clinical data, images are produced in the structurally complex DICOM<sup>®</sup> format. The use of DICOM<sup>®</sup> gives rise to three issues:

1) *Messaging protocol:* DICOM<sup>®</sup> defines a messaging protocol as well as a file format. Images can be encapsulated by messages, and the content of these messages includes encoding and messaging specific information which is not relevant to understanding the image itself, and which is dependent on the implementation context of the supporting software. When DICOM<sup>®</sup> images are transmitted using these messages, it is impossible to confirm that they remain unchanged in accordance with the relevant regulations (ICH 1996).

2) *Structural complexity:* The 2008 version of the DICOM<sup>®</sup> standard is over 3600 pages long with additional supplements and corrections added since publication. The format is so complex and flexible that different implementations often produce images with structurally non-conforming aspects. These images need to be “cleaned” so that they flawlessly conform to the DICOM<sup>®</sup> standard, or else they may be unreadable and unusable in the clinical trial.

3) *De-identification of images to meet trial requirements:* Images collected for clinical trials are also part of the general care workflow of hospitals. As such, DICOM<sup>®</sup> images potentially contain hundreds of different items (e.g. patient names and addresses) that must be removed before they meet clinical trial de-identification requirements set forth by HIPAA and corresponding European regulations.

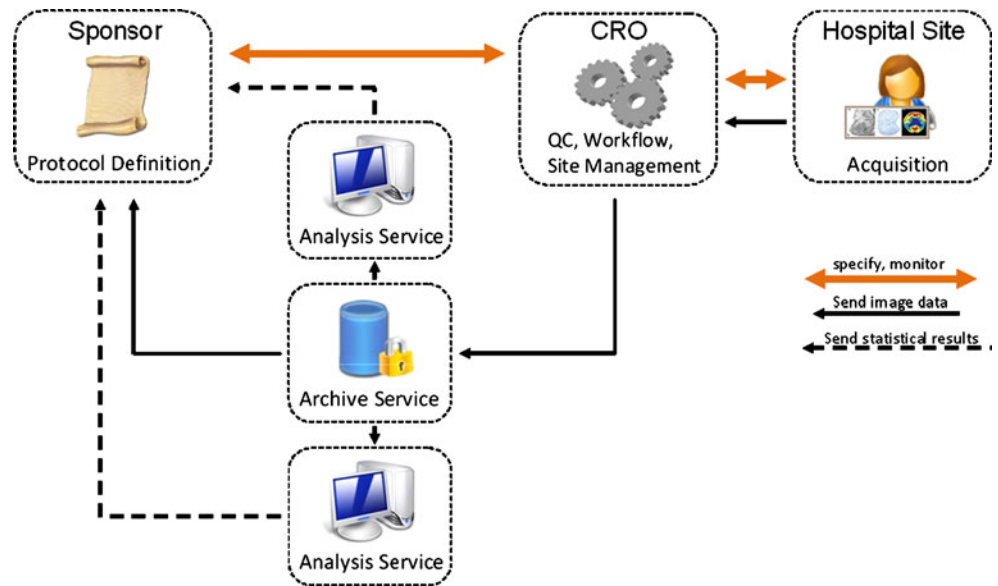
A site image submission service compatible with source data regulations must be able to provide solutions for those three issues.

Issue 1 may be resolved by using the DICOM<sup>®</sup> file format exclusively in order to “freeze” an image into a persistent byte-stream whose contents are static, and which can be copied across systems and remain unchanged.

To resolve issues 2 and 3, we reference the fourth eSDI scenario, “Extraction and Investigator Verification (Electronic

<sup>3</sup> The IC cannot simply allow for image use in unspecified or unknown contexts (ICH 1996).

**Fig. 2** Potential data flow and distribution of services in an SOA enabled workflow



Health Records),” as the best match for the hospital imaging workflow. Since the DICOM® images are not initially compatible with clinical trial protocol requirements, and the scanner itself as a medical device does not fall under the computerized system validation requirements in 21 CFR 11, the original images can be viewed as being part of an electronic health record (EHR) environment rather than a clinical trial electronic data capture (EDC) system. They may be cleaned and de-identified to remedy issues 2 and 3, verified by the investigator, e.g. with electronic signature, and are only thereafter considered source data (CDISC 2006) (Fig. 3).

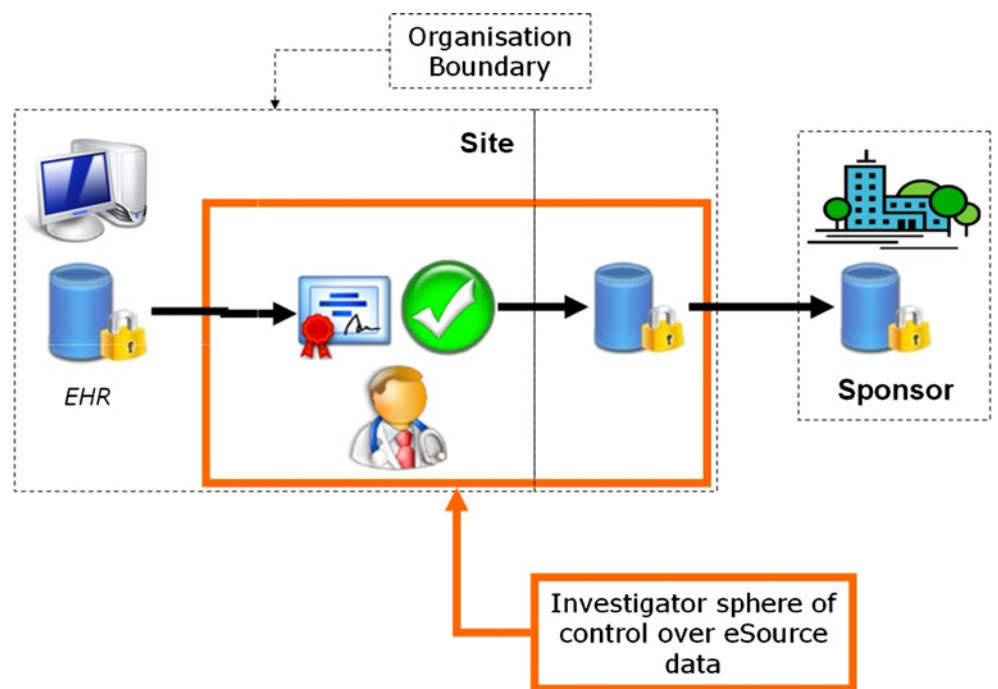
Sites frequently lack the proper tools to enact this workflow, providing sponsors and CROs a clear opportu-

nity to enable a means of optimal regulatory compliance by using site submission tools meeting the following requirements: 1) image cleaning and de-identification, 2) means for the investigator to be able to verify the source data, 3) image copy from hospital site to target in a persisting format, and 4) persistence of a local copy, or access provision to an online copy.

**Data Accessibility**

Partners who need to exchange imaging data can provide or consume image transfer services, which will depend on the XML definition of clinical trial image collections. For

**Fig. 3** “The fourth eSDI scenario for extraction of data from electronic health records with verification by the investigator” (CDISC 2006)





example, a CRO and a sponsor might each deploy “Submit Image Object” and “Retrieve Image Object” services built on top of the caGrid platform (which provides a means of transmitting the binary image data as well as implementing service interfaces); sites and analysis service providers could connect to these using standard image submission clients.

Storage and bandwidth capacity at a sponsor with many imaging trials may be too limited to allow collection and retention of all images for all trials. SOA enables convenient access to distributed archives of images at CROs, using image query and transmission services for retrieval as needed.

### De-identification

Achieving better control of de-identification at both CROs and imaging sites is largely a matter of providing standards and tools. Standardizing on DICOM<sup>®</sup> for image representation gives a frame of reference against which de-identification requirements can be defined. The standard requirements themselves must still be defined, and there must be flexible mechanisms for enforcing them. DICOM<sup>®</sup> Working Group 18 has produced a draft supplement to the DICOM<sup>®</sup> standard providing a number of confidentiality profiles for this purpose. An XML format can be used to provide a machine readable specification of the standard profiles, and also case-specific customizations, e.g. the retention of a private tag which contains coil information necessary for image contrast quantification. This definition can be shared, and the sponsor and other partners can build supporting tools for use at hospital sites.

Sometimes identifying information is “burned into” the actual image, meaning that part of the image itself is overwritten with a patient identifier. This is more difficult to automatically correct like the metadata, but can be checked during QC, as described below.

### QC

More generalized QC requirements can be described with XML in the same way as de-identification rules. A patient name field may be slotted for removal, but an acquisition date field may be flagged as requiring content, or an “echo time” field for MRI data may be required to have a specific value range. DICOM<sup>®</sup> compliance and image compatibility with processing software can be improved significantly by enforcing such rules.

Using a standard QC profile and a data exchange service, a CRO can implement real-time QC procedures. A completely automated workflow may be defined to receive images through a web accessible service, apply a QC profile upon receipt and even check for image acquisition quality issues such as artifacts, and accept or reject the submission with appropriate notification back to the sender. A process that

previously took days is reduced to minutes, and sites can receive feedback soon enough to take remedial action if needed, e.g. a patient re-scan. Additionally, the site itself can take on the QC burden using a supporting toolset.

Some QC measures can be performed on the image bitmap itself, such as checking for burned in identifying information, evaluating the contrast, or checking that the correct anatomy has been imaged. These can be integrated into the workflow in much the same way as analysis algorithms, as described below.

### Analysis

With SOA, analysis tools can be wrapped in public interfaces, making them discrete and re-usable components of the trial workflow just like other services. Still, analysis tools have widely varying semantics in terms of what inputs they require (a single image, two separate images acquired with different scanning sequences, images combined with specific clinical data, etc.) and what outputs they produce. Harmonizing and ultimately standardizing the semantics of analysis algorithm interfaces is a necessary effort above and beyond SOA.

### Implementation

An SOA depends on services, and their interoperability depends on common syntax and semantics. A growing number of providers have already begun to offer componentized solutions. For example, image analysis algorithms can be packaged and sold as fully automated applications. So far, however, the providers have been given neither incentive nor technology standards for building service interfaces to their tools. Sponsors can drive the adoption of SOA for imaging trials by defining messaging and data standards, deploying their own applications as services, and incentivizing their adoption through open source application development and contractual stipulations.

DICOM<sup>®</sup> comprises an imaging specific messaging protocol that supports many radiology department workflows, such as image transfer from scanner to picture archiving and communication system (PACS), image querying, printing, hanging protocol, graphical annotations, teaching files, and image based reporting. However, these capabilities were developed for use between local radiology sites via secure hospital LANs. DICOM<sup>®</sup> is ill equipped as a messaging and image transfer protocol over the Internet due to lack of security, fault tolerance, discoverability, and adaptability; from a regulatory perspective, it’s problematic for clinical trial workflows simply because no established audit trail can verify image preservation. XML based specifications, such as Simple Object Access Protocol (SOAP), do not suffer from these

drawbacks. They are widely supported as service messaging standards due to flexibility and broad tool support. They form the backbone of a group of standards known collectively as web services, which are widely implemented in the SOA solution space.

The messaging standards define a syntax for interoperability, but a semantic model for describing the data is also needed. This is a pre-requisite for sharing quality requirements and correctly performing image reads. While DICOM® is unsuitable in this case as a messaging protocol, it does provide a widely supported image file format and semantic model. Presently most imaging devices can output in DICOM®, making it an ideal choice for source data capture. Downstream analysis applications often require other formats, such as NIfTI-1.1<sup>4</sup> (a specification designed to facilitate file-based interoperability between functional MRI analysis tools) or Nrrd (a file format and toolset for representing and processing multi-dimensional image data),<sup>5</sup> but in such cases a conversion has to be performed regardless, and many tools exist for the purpose. Maintaining the original image data in their DICOM® file format is a simple approach which allows for maximum flexibility. Beyond the images, additional trial metadata must be described, to define both the interrelationships of image files (i.e. which DICOM® image instances or DICOM® series comprise a discrete unit of input for downstream analysis) and the associations to trial time points, subject identifiers, etc. These may be thought of as “metadata” in a DICOM® and imaging world, but they are integral to the primary data in clinical trials. They need to be managed in databases and transmitted via standard database integration techniques. XML is again the de facto web services standard format for serializing database objects.

With interoperability standards defined, the next step is implementing and deploying services. The software infrastructure enabling interactions between these services is known generically as “middleware.” Web services middleware comprises implementations of standards such as SOAP. A clinical imaging trial SOA requires a middleware that supports partners with widely varying IT proficiency and budget; provides out of the box support for web service standards for messaging, security, and binary data transport; enjoys widespread adoption; and complies with relevant healthcare standards. caGrid fully satisfies these requirements, supporting “the requirements associated with discovery, analysis and integration of large scale data, and coordinated studies [in] biomedical fields” (Oster et al. 2008). It was developed as the software infrastructure underlying the cancer Biomedical Informatics Grid® (caBIG®), an initiative

of the National Cancer Institute to enable sharing of data and research tools across institutions (Oster et al. 2008). Its particular advantages include: 1) a robust and consistent platform dedicated to service implementation, 2) an open source license and free support resources that can be used by all partners, including small partners unable to employ more expensive middleware, 3) a library of existing services specific to healthcare and resource sharing, 4) a strong reception and increasing adoption by the biomedical research communities, and 5) coherence with all important standardization initiatives (HL7, DICOM®, CDISC). A substantial technical overview as well as information about its utilization by the healthcare community may be found at the caGrid website.<sup>6</sup> As an example of a specific implementation for real world studies, the CardioVascular Research Grid (CVRG) is linked with eight driving biomedical projects as described on its project website.<sup>7</sup>

Competing commercial middleware, such as TIBCO’s offerings, Oracle Fusion Middleware, or IBM WebSphere are expensive products with no special healthcare focus. Free and open source middleware is available, such as Apache Axis for SOAP, but again lacks a healthcare focus, as well as the toolkits and support resources of the commercial products. caGrid strikes a balance between these. It is open source, free, and also provides commercial-grade developer support with features such as a service development GUI. It is built on top of other free middleware such as Axis, providing additional, not competing, features. Mature products built with it are already available, such as the In Vivo Imaging Middleware (IVIM). These too are open source, offering a significant value addition for new service development with overlapping requirements.

Other resource sharing and interoperability initiatives exist, but none whose goals are so well aligned with the infrastructure needs of clinical trials with imaging endpoints, in particular the provision and marketing of a service oriented middleware. For example, the Biomedical Informatics Research Network (BIRN) is a collaborative effort between several research institutions to provide data-sharing infrastructure. It offers a grid of shared data and analytical resources, but lacks a service-oriented architecture and a middleware product. Integrating the Healthcare Enterprise (IHE) is an organization focused on issuing standards-based specifications of healthcare tasks which can be implemented to achieve interoperability between systems. It provides a number of profiles for common healthcare use cases. Of particular interest is its Cross Enterprise Document Sharing for Imaging (XDS-I.b) Integration Profile (Seifert 2009). This defines a set of roles and

<sup>4</sup> <http://nifti.nimh.nih.gov/nifti-1>

<sup>5</sup> <http://teem.sourceforge.net/nrrd/>

<sup>6</sup> <http://www.cagrid.org/display/knowledgebase/caGrid+1.3+Technical+Overview> and <http://cagrid.org/display/community/Home>

<sup>7</sup> <http://www.cvrgrid.org/?q=content/driving-biomedical-projects>

transactions, in terms of web service standards, which can be implemented to enable the kind of on demand, distributed image storage and retrieval available with a caGrid enabled SOA. XDS-I.b provides concrete, tested specifications, but no implementations, middleware, or tools. Also, it satisfies just one use case of several for imaging trial enhancement. caGrid provides a complete implementation middleware with all the underlying messaging, image transfer, and security mechanisms necessary to tie services from different partners together into trial workflows. As an open source framework, caGrid does not include the same level of customer support and comprehensive documentation as some commercial alternatives, resulting in a steeper learning curve for adopters. On the other hand, innovators can help flatten this curve with their own open source, caGrid-enabled reference tools.

Novartis has done this with the release of ImagEDC,<sup>8</sup> an image de-identification, cleaning, and submission software. It consists of a client tool which can be deployed at hospital sites, enabling them to produce images compatible with trial requirements without additional processing required from the CRO; a “Submit Image Object” web service which consumes trial data and images, storing them in a local repository; and a tracking service which records workflow events such as image submission and receipt. The client is designed with the aforementioned eSDI scenario in mind. The inputs may be considered EHR data, and the outputs include local and remote DICOM image files, plus an electronic record of the operator’s user name, which can be used to signify verification that they constitute the source data for a clinical trial. The service is defined using web service standards, and built using caGrid. The client or the service business logic can be re-used as they are, or re-implemented according to the needs of a specific organization or trial, without changing the service definition. For example, ImagEDC could be deployed to connect a CRO and a sponsor. The reference “Submit Image Object” operation could be used to transfer the image files immediately to the sponsor, or another implementation could be written which would only transmit the identifiers. The sponsor could use these to track image availability at the CRO, and issue queries to retrieve the image data on-demand. The site could use the reference implementation of the client tool in either case.

ImagEDC demonstrates the benefits that moving toward an SOA can bring to clinical imaging trials. Fully realizing this goal would require the implementation of a broad range of services. ImagEDC is a single example covering the “Submit Image Object” scenario. As a proof of concept, it also shows the suitability of caGrid as a service middleware that doesn’t

require extensive IT resources for implementation. At the time of this article’s publication, ImagEDC is the product of two to four weeks full time equivalent (FTE) development effort, with a heavy reliance on caGrid, but no prior expertise with it. Substantial organizational challenges involved in implementing an SOA across different partners remain, but these are present for any infrastructure integration efforts. Sponsors can address them with contractual obligations for CROs; with sites, it is less straight forward, since they do not typically have infrastructure and IT resources to implement custom integrations or complicated deployments. This can be addressed by providing them with tools with minimum barriers to adoption. For example, ImagEDC uses Java Web Start technology to allow a user to download, install, and run the tool on his or her local machine simple by visiting a URL in a web browser.

## Conclusion

There is opportunity to enhance the conduct of imaging trials through SOA and grid computing approaches. Improved interoperability brings clear benefits to trial sponsors in terms of flexibility, efficiency, and quality. Service providers also gain a wider market for their products, and everyone benefits from lighter integration efforts.

While there is much useful work being done already on healthcare interoperability standards such as IHE’s XDS-I. b, CROs and sponsors have yet to implement these as part of a service-oriented framework optimized for resource sharing. The missing ingredient is middleware. caGrid neatly fills this gap, providing not only a web services stack, but a rich set of core services and developer tools.

## Information Sharing Statement

All software, including source code, referenced in this article may be accessed by the general public at the URLs given within the text. The associated websites contain detailed information about licensing, source code acquisition, compiling, deployment, and user communities.

**Acknowledgments** We would like to thank David Tuch for his helpful comments and advice during the creation of this article.

**Open Access** This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

## References

Clinical Data Interchange Standards Consortium, Electronic Source Data Interchange (eSDI) Group. (2006). Leveraging the CDISC

<sup>8</sup> Software source code, open source licensing information, and documentation, including instructions for deployment and use, are publicly available at <http://code.google.com/p/imageduc/>



- Standards to Facilitate the Use of Electronic Source Data within Clinical Trials. Version 1.0.
- FDA. (1997). 21 CFR Part 11: Electronic Records; Electronic Signatures; Final Rule, Federal Register, Washington, DC.
- International Conference on Harmonisation of Technical Requirements for Registration of Pharmaceuticals for Human Use (ICH). (1996). ICH Harmonised Tripartite Guideline, Guideline for Good Clinical Practice E6(R1).
- Oster, S., Langella, S., Hastings, S., Ervin, D., Madduri, R., Phillips, J. et al. (2008). caGrid 1.0: An enterprise grid infrastructure for biomedical research. *Journal of the American Medical Informatics Association*, 15, 138–149.
- Seifert, P. (2009). IHE Radiology Technical Framework Supplement—Cross-Enterprise Document Sharing for Imaging (XDS-I.b) Integration Profile, Trial Implementation, Rev. 1.1 - 2009-06-21, IHE International.
- Uppoor, R. S., Mummaneni, P., Cooper, E., Pien, H. H., Sorensen, A. G., Collins, J. et al. (2008). The use of imaging in the early development of neuropharmacological drugs: a survey of approved NDAs. *Clinical Pharmacology & Therapeutics*, 84, 69–74.