
Department Informatik

Technical Reports / ISSN 2191-5008

Christoph P. Neumann, Florian Rampp, Richard Lenz

DEUS: Distributed Electronic Patient File Update System

Technical Report CS-2012-02

March 2012

Please cite as:

Christoph P. Neumann, Florian Rampp, Richard Lenz, "DEUS: Distributed Electronic Patient File Update System",
University of Erlangen, Dept. of Computer Science, Technical Reports, CS-2012-02, March 2012.

DEUS: Distributed Electronic Patient File Update System

Christoph P. Neumann, Florian Rampp, Richard Lenz
Institute of Computer Science 6 (Data Management)
Dept. of Computer Science, University of Erlangen, Germany
`christoph.neumann@cs.fau.de`

Abstract—Inadequate availability of patient information is a major cause for medical errors and affects costs in healthcare. Traditional approaches to information integration in healthcare do not solve the problem. Applying a document-oriented paradigm to systems integration enables inter-institutional information exchange in healthcare. The goal of the proposed architecture is to provide information exchange between strict autonomous healthcare institutions, bridging the gap between primary and secondary care.

In a long-term healthcare data distribution scenario, the patient has to maintain sovereignty over any personal health information. Thus, the traditional publish-subscribe architecture is extended by a phase of human mediation within the data flow. DEUS essentially decouples the roles of information author and information publisher into distinct actors, resulting in a triangular data flow. The interaction scenario will be motivated. The significance of human mediation will be discussed. DEUS provides a carefully distinguished actor and role model for mediated pub-sub. The data flow between the participants is factored into distinct phases of information interchange. The artefact model is decomposed into role-dependent constituent parts. Both a domain specific (healthcare) terminology and a generic terminology is provided. From a technical perspective, the system design is presented. The sublayer for network transfer will be highlighted as well as the subsystem for human-machine interaction.

Index Terms—Healthcare, information systems, inter-institutional, domain engineering, distributed applications, distributed data structures, document-orientation, publish-subscribe, human mediation

I. INTRODUCTION

In a systems analysis of adverse drug events, 18% of the medical errors were associated with inadequate availability of patient information [1]. The problem of inadequate availability of patient information as a major cause for medical errors is aggravated by the rise of healthcare networks and the increasing number of healthcare parties that are involved in a treatment: The

aging of western society affects the public health sector, chronic diseases and multimorbidity become the focus of interest, and the cost pressure rises. Chronic diseases and multimorbidity, like cancer, diabetes, asthma, and cardiac insufficiency, require more healthcare parties than common diseases. Coevally, the rapid advance in medicine leads to an advancing specialization of physicians that is an additional cause for the increasing number of involved parties regarding a single patient's treatment. For improving the treatment quality and in order to avoid unnecessary costs, an effective information and communication technology is vital for the support of inter-institutional patient treatment.

An IT infrastructure for healthcare networks must respect and consider the autonomy of preexisting systems in different institutions. At the same time some kind of integration is required, which helps to preserve inter-institutional data consistency. Closely integrated systems with a common database and terminological standards for database contents are unrealistic in this scenario. In particular, strict autonomy of the institutions requires the abdication of central infrastructure like joint databases, transaction monitors, and central context managers. Shared communication requires minimal standards avoiding full-fledged platform-specific middleware frameworks. Instead, a document-oriented publish-subscribe paradigm is favored, which supports loose coupling of systems at different sites.

In any case, semantic agreements like healthcare ontologies, terminologies, and clinical models evolve over time. Therefore, an adequate system design methodology for evolutionary information systems becomes imperative [2]. A layered approach for healthcare information systems decouples system design processes on different levels of abstraction, decreases complexity in each layer, and thereby supports system evolution and responsive infrastructures. In [3] such an approach with four layers is proposed: *generic framework* layer, *domain framework*

layer, *application layer*, and *custom layer*, with different people as key drivers for the different layers. This model, however is an idealized picture, which is not yet related to existing standards and frameworks. In this paper, the research focus is on domain framework solutions as extensions to existing healthcare information systems.

This paper describes a publish-subscribe architecture for healthcare supply chain scenarios. This technical report is an extended version of a previous publication [4]. There are two distinguishable features of the proposed solution: to apply document-orientation as instrument of inter-institutional integration and to allow patients to control information distribution. To put a mediated publish-subscribe architecture into practice requires a systematic distinction of actors, roles, phases, and responsibilities in the distribution scenario. The proposed architecture essentially decouples the roles of information author and information publisher into distinct actors.

II. SUPPLY CHAINS IN HEALTHCARE

A short overview of the domain participants is given: The focus of the medical supply chain in Germany are the patients who are treated by office-based physicians foremost, collectively described as the *primary care*. The *secondary care* adds hospitals, laboratories, pharmacies, and ancillary medical institutions as participants of the medical supply chain. Accompanying participants are the health insurance funds and the associations of statutory health insurance physicians.

The whole interaction and collaboration is liable to many technical, organizational, economic, and legal factors. The legal boundary conditions are critical for the information provision and availability because warranty of data protection is essential for patient-related data. The patient has to maintain sovereignty over any personal health information. This provides a basic motivation for the mediation approach that will be applied to the publish-subscribe pattern.

A. Inter-Institutional Problems

Considering the comprehensive medical supply chain, functional integration and process integration between the autonomous information systems of the several participants is still unsolved in organizational and diagnostic-therapeutic processes [5]. Particularly chronic diseases like diabetes, asthma, and cardiac insufficiency require a long-cyclic exchange of patient information between all involved healthcare professionals from different institutions [6]. At present, this is done by paper documents like discharge letters or by repeated

anamnesis interviews with the patient. The sticking point is that the discharge letters are frequently missing or are insufficient in detail [7], either because they are not written by the physicians at all or because they are not available to all involved parties [8]. Repeated anamnesis interviews are no alternative for document interchange between healthcare professionals. Redundantly applied diagnostic methods by each distinct institution are the norm. As simple as order entry and result reporting may seem, it is still one of the most important issues in healthcare information systems [9]. The data integration of healthcare information systems will be addressed by a document-oriented approach in the proposed solution.

In order to foster the continuity of care [10], the inter-institutional cooperation needs to bridge the current gap between institutions of the primary and secondary care [11]. Such effort must not instrument regional standards, as it is done in *regional healthcare information networks* (RHIN) [12], but *transregional standards*. Accomplishing information exchange in distributed healthcare scenarios requires the integration of heterogeneous and *strict autonomous IT systems*. To allow for *inter-institutional cooperation* the support of distributed and seamless flow of information is required, thus changing paradigms from closed and hegemonic to open and distributed architectures. The proposed architecture adheres to these boundary conditions.

B. From Bilateral Information Exchange to Information Distribution

Information interchange by letters is the way of traditional cooperation – a referral from one institution to another delegates responsibility and liability of diagnosis or therapy to the other institution. Yet, genuine physician teams from different institutions are upcoming [13]. For some years now, in Germany, the treatment of breast cancer is organized by accredited in-station breast cancer treatment centers cooperating with manifold accredited partners like oncologists, radiologists, and the post-operative care [14]. Collaborative treatment scenarios can be described as distributed medical treatment processes with physician teams from different institutions interacting closely meshed for treating complex chronic diseases and multimorbidity.

Such scenarios change the requirements for the availability of patient information and deserve extended information exchange models, still adhering to the strict autonomy of the involved IT systems. Although comprehensive IT support for such closely meshed treatment scenarios in general has unsolved legal boundary condi-

tions, this paper outlines an architecture for distributed patient files as a technical foundation. It is based on digital information cards that are yielded from institutional *electronic health records* (EHR), e.g. [15], into a distributed publish-subscribe system that allows the patient to govern data interchange.

C. Fundamentals of Healthcare Data Models

Exchange of patient information among institutions requires data compatibility. *Data integration* is the process that achieves such data compatibility, either by common standards or by data transformation. Data integration for medical processes requires standards for medical terminology that have to deal with volatile medical concepts [18]. Over the intervening years numerous standards for medical ontologies have been created on *type level* for system implementers at design-time and on *instance level* primarily for end-users at run-time.

At instance level, standards like ICD [19], SNOMED [20], and LOINC [21] exist which unremittingly evolve over time. The HL7¹ v2 [22] is a well-established standard for clinical message specification. It is a standard on type level, and incorporates arbitrary coding schemes and terminologies on instance level. Dealing with inherent volatility of reference terminologies by information system design in a general matter actually is another unsolved scientific issue. Despite many attempts, there is no stable unique and comprehensive ontology of the medical domain in sight.

The HL7 v2 standard allows for the specification of arbitrary self-defined messages, which has lead to incompatibilities. The relatively new HL7 v3 standard is based on the HL7 v3 *reference information model* (RIM) and is radically different from the v2 standard: It allows for new types to be derived from a limited number of core classes, enabling RIM-based systems to handle even unknown message-types in a generic way. Furthermore, a conceptual change from messages to documents is provided by the HL7 v3 *clinical document architecture* (CDA) substandard [23]. CDA allows for XML-structured medical documents. Based on HL7 v3 CDA, in Germany, the SCIPHOX² [24] working group tries to develop specific document content types for German healthcare; for example, referral vouchers and discharge letters³.

¹Health Level 7, <http://www.hl7.org>

²Standardized Communication of Information Systems in Physician Offices and Hospitals using XML

³Particularly “eArztbrief SCIPHOX CDA R1” and its advancement “eArztbrief VHitG CDA R2” [25]

Any new standards should respect the ones already in practice for backwards-compatibility and to achieve and maximize acceptance. In conclusion, the event type of the proposed mediated publish-subscribe system will be based on HL7 v3 CDA.

D. Objectives

There are two prime objectives of the proposed solution. The first is the abdication of any central server. The second is the application of document-oriented integration with lightweight interfaces instead of service-oriented integration with semantically rich interfaces. A subsequent design objective is to aim for minimal standards in order to yield minimal requirements to the participating systems. Favoring local autonomy over central hegemony requires, for example, that distribution of information will not be enforced, but is voluntary and process participation can be supplemented on demand. Platform independence and the avoidance of vendor lock-ins require that the basic architecture is decoupled from any specifically instrumented middleware and components off-the-shelf. A sophisticated system modularization has been a major design goal.

III. APPLYING DOCUMENTS AS EVENT TYPE

Integration in healthcare is traditionally based on interface-orientation. Three-tier network-based architectures with remote procedure calls are yet the dominant style of information systems. The most common technological occurrence of remote invocations is based on SOAP with WSDL, augmented by several WS-* frameworks. The *interface-oriented integration* focuses on available functionality, and the integration method affects semantically rich interfaces. An invocation uses parameters to detail its synchronous service request to a target system. In interface-oriented integration the information being passed is not necessarily viable on its own but often in the context of the service request only.

Even if a service is triggered event-oriented using asynchronous messaging, like it is done in HL7 v2-based systems, such parameters or messages essentially represent transient fine-grained information that is assimilated by the targeted system. The three main problems in information integration projects, including healthcare systems, are insufficient synchronization of redundant data, problems with data consistency, and functional overlapping [26]. Therefore, interface-oriented and message-oriented integration between distinct institutions is complex and custom-designed.

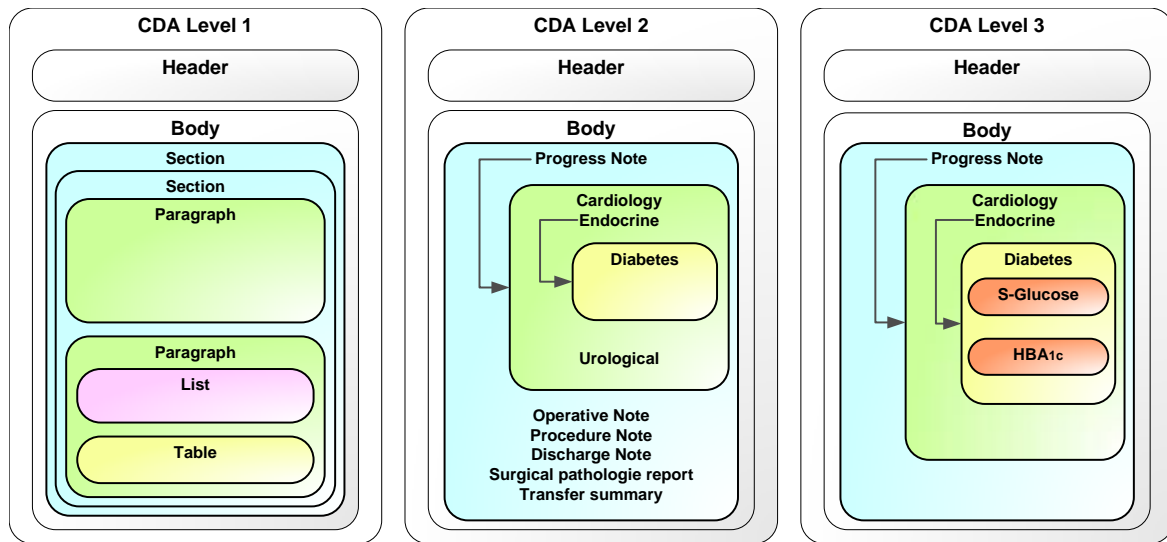


Fig. 1. HL7 v3 CDA structure outline for levels 1, 2, and 3 (cf. [16], [17])

In contrast, documents are coarse-grained, self-contained, and viable. A document carries its own context information and can exist independently from the system it stems from. Changes are not propagated by update information, but by creating an updated document that replaces its predecessor. The *document-oriented integration* focuses on available information, and the integration method affects the semantic scalability of document models, using standardized and minimal interfaces for hand-over. Redundancy in data distribution is not critical with documents because, due to the self-containedness, a synchronization in the classical sense is not required. Likewise are data consistency checks confined to the scope of the document.

HL7 v3 CDA provides semantic scalability for healthcare documents, both because this has been an inherent feature of the underlying RIM and because CDA is particularly structured in three levels of semantic abstraction: In Fig. 1 a basic outline of the three CDA levels is provided as illustration. CDA level 1 is the unconstrained CDA specification. CDA level 2 applies section-level templates. CDA level 3 applies entry-level templates. For example, CDA level 1 simply ensures the ability to display a document like a PDF file. Any CDA document can be accepted without immediate support for processing. Advanced semantic processing support of CDA level 2 or 3 can be added to the system, seamlessly enhancing the information value of already stored CDA documents.

The *deferred system design* principle of evolutionary systems [27] requires semantic decisions not to be frozen

in an interface schema because they are hard to revise. HL7 v3 CDA supports deferred system design by its semantic scalability. The proposed solution uses CDA for typing the *Document Messages* [28, p. 147] in its mediated publish-subscribe scenario. Event models have a major impact on the flexibility and usability [29]; applying a document-oriented approach improves the adaptability of the systems by deferring schema decisions from design-time to deploy- or run-time [30].

IV. PROPOSED SOLUTION

The proposed solution architecture, the *distributed electronic patient file update system* (DEUS), applies the document-oriented idea in form of DEUS *digital cards*⁴: digital cards are self-contained and viable containers of information. A digital card is authored by an information provider (first part of the digital card ID; e.g. a physician) and it concerns a person (second part of the digital card ID; e.g. a patient). Because an author can provide several digital cards about a concerned person, the author provides an additional discriminator as third part of its identifier. The digital card as information container is subject to PKI signatures, as it will be detailed by the DEUS scenario description. DEUS supports even closely meshed physician teams and fosters trans-sectional, life-long, and patient-centered healthcare documentation.

From a conceptual perspective, such exchanged information is a document as required by a document-oriented architecture approach. However, the intended

⁴Unfortunately, the preferred term *information card* has been patented in 1972 by Paul P. Castrucci, US patent 3702464.

granularity of a DEUS digital card is a more fine-grained one than the one that is experienced from paper-based working practice in healthcare. This improves the structure of the patient files, and provides higher selectivity in retrieval and display. One DEUS digital card contains, for example, a diagnostic finding, clinical evidence, a diagnosis, a therapeutic measure, an order, or a prescription. The digital card metaphor has been inspired by the Higgins project [31] with its I-Cards as foundation of an open source identity framework. The exemplary actor “Dr. Higgins”, as actual contributor of digital cards in the DEUS scenario in the next section, is a homage to this project.

A. Interaction Scenario

The basic DEUS scenario is outlined in Fig. 2. The patient (“Alice”) has recently visited a healthcare professional (“Dr. Higgins”) and the obtained information has to be shared to other involved parties (inter alia “Prof. Bob”). The local *healthcare information system* (HCIS) of Dr. Higgins, the author of the obtained information, bundles the information into a digital card. This digital card is electronically signed by its contributor and becomes the subject of information distribution. Subsequently, it is contributed into the node’s local DEUS system extension.

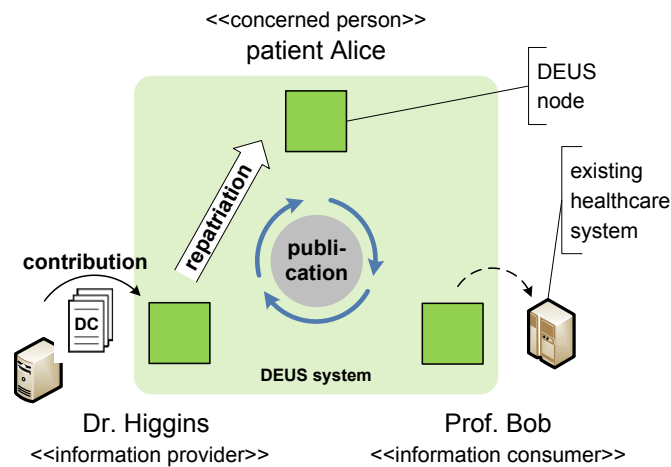


Fig. 2. The DEUS scenario as mediated publish-subscribe system

The exported digital card is subsequently transferred to the account of patient Alice who is the person being concerned by the medical information. The patient as sovereign of information distribution decides whether the information is accepted into the pool of digital cards that builds his or her *personal patient file*. The process of transferring a digital card from the contributing DEUS

system to the patient’s DEUS system together with the patient’s decision about the acceptance of the digital card is named *repatriation*. Subsequently, the digital card is published to any subscriber DEUS systems, like Prof. Bob. He will consume the information later, for example when Alice is visiting next time.

Each DEUS participant owns a DEUS account. A DEUS node is a healthcare information system with an installed DEUS extension. A DEUS node can host multiple DEUS accounts. For the mediated publish-subscribe interactions, it is transparent whether an account resides on the same or another DEUS node.

B. The Significance of Human Mediation

“Patient Empowerment” is a general issue in eHealth, for example being addressed by the EU in form of a ministerial declaration [32]. Empowering the patient specifically to control information distribution might still seem radical: Today, in working practice, the healthcare professionals can request patient information from each other. The cornerstones are the thorough trust of a patient in his or her physician and the trust of physicians among each other, due to the Hippocratic oath. Information exchange does not necessarily involve patients, aside from instrumenting the patient as a useful surrogate for postal delivery.

Life-long, trans-sectional, electronic storage and communication cannot be compared to such traditional approach: If a physician requests information in form of a paper-based, signed request it is checked and kept filed by the addressed institution. A plausibility check by a human actor is integral part of any such exchange, and any abuse in paper-based scenarios using faked requests risks instant detection; especially if more than a few patients are concerned. In contrast to paper-based requests, in electronic patient file scenarios, a high quantity of patients being affected by an attack will not automatically result in a timely detection, if the system-supported communication bandwidth and latency allows for fast assault automation.

In conclusion, the life-long distribution of patient-related information requires a human authority over inter-institution-ally shared information, so that the induced transparency of patient information by a server-side electronic storage does not erode doctor-patient confidentiality. The only appropriate healthcare participant for publication control is the patient as the concerned person. If the patient lacks the abilities to sovereign his or her healthcare information, it is possible to delegate this role to a legitimate proxy person or institution, possibly

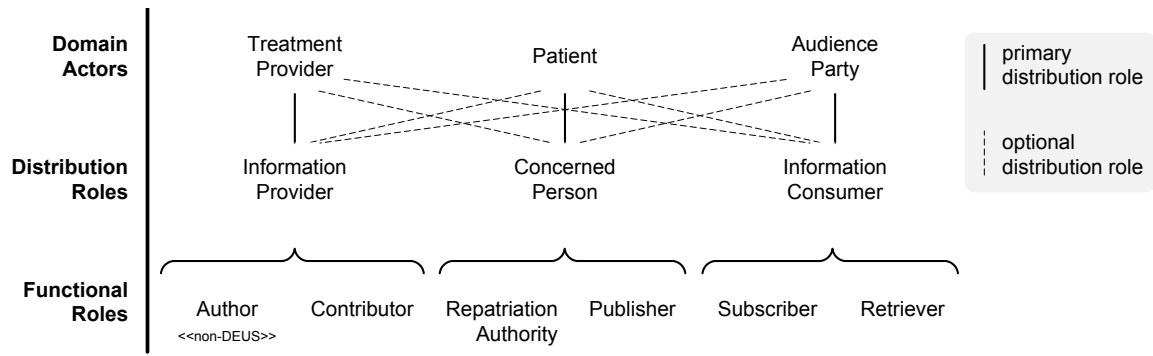


Fig. 3. The DEUS roles of healthcare participants

a general practitioner. The mediation of HCIS-to-HCIS communication by the patients' DEUS accounts reintroduces interception and control by a conscious human in an electronic environment for gradually automated information interchange.

C. The Significance of Repatriation

The necessity for a repatriation mechanism can further be clarified by describing its notable absence in a common distributed information environment – the PGP key server infrastructure: It is possible for a rogue signer to attack a PGP key that is published to the global directory by creating thousands of garbage keys, signing the attacked key with all of them, and posting the attacked key with a bloated signature extension to the key server network. A PGP key server does not check whether a key is submitted by the owner of the key, but allows anyone to post additional key extensions to the key server network. Henceforward, each key-updating owner and consumer is plagued by the bogus signatures, possibly leading to an unusable key and therefore successful denial of service. An additional phase has not been considered in PGP key server infrastructure – allowing only key extensions that are counter-signed by the owner for upload to the key server network would have solved the issue, implicating a kind of “repatriation”. Luckily, PGP users seldom attract attackers that would exploit this floodgate design flaw.

Distributed healthcare infrastructure for long-term information must preclude attackers from injecting bogus information by implementing a kind of repatriation. This allows the concerned person to check a digital card received from an information contributor. The concerned person will counter-sign an accepted digital card, because this enables the subscribers to verify the authenticity, finally. A public key infrastructure that allows the concerned person to securely white list author accounts

as valid information contributors optionally helps to automate the repatriation decision. The proposed *mediated publish-subscribe architecture* essentially decouples the roles of information author and information publisher into separate actors. A description in subsect. IV-G will portray how DEUS could be used in organizationally closed environments with special trimmed-down logical patient accounts for scenarios in which patient interaction is not feasible.

D. Actors and Roles

The three basic participating actors and their different roles in the DEUS information interchange are summarized in Fig. 3. Each actor has its dominant role in the distribution process, visualized as straight vertical lines between the layer of *domain actors* and the layer of *distribution roles*. However, each actor can assume, by its DEUS account, each distribution role as it is visualized by the clashed lines: For example, the patient can act as information provider about himself, providing information like allergies or legacy paper documents that he contributes in transcribed or scanned form. On the other hand, a treatment provider can use its DEUS account in the role of a concerned person to provide and publish business card information or consultation hours information. Even associations of statutory health insurance physicians could participate by contributing certificates like the ones required for accredited in-station breast cancer treatment centers, and health insurance funds could contribute the patient master data.

The distribution roles are related to the specific role that a DEUS account assumes in an overall distribution scenario. The *functional roles* are responsibilities that are deduced from a distribution role.

The information provider acts as author of the information, which takes place inside the HCIS and is not part of DEUS. The information provider acts as contributor by

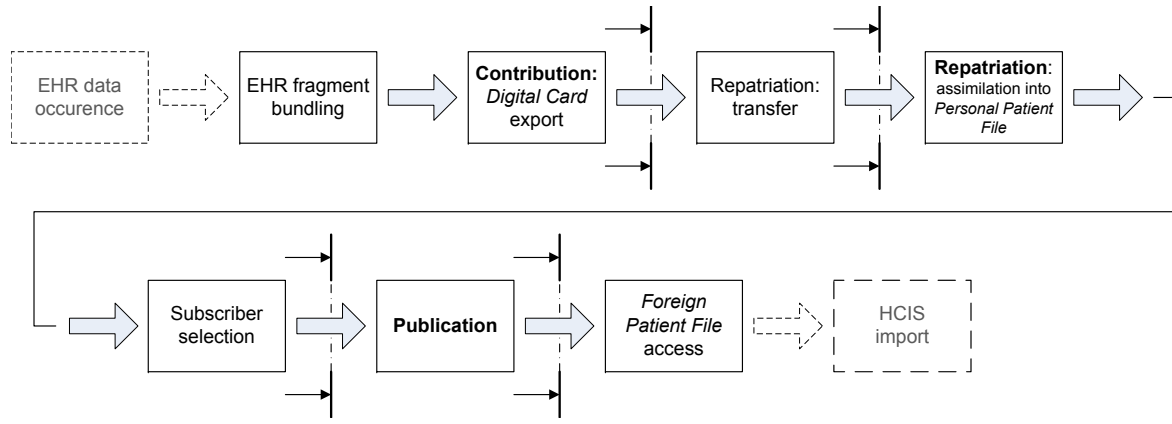


Fig. 4. The DEUS contribution-repatriation-publication chain

exporting a digital card and handing it over to the actual DEUS account. The concerned person acts as repatriation authority by deciding about the validity of a repatriated digital card. The concerned person acts as publisher by applying the selection of subscribers and performing the publication transfer. The information consumer acts as subscriber by establishing subscriptions to the account of the concerned person and by accepting published digital cards. Finally, the information consumer acts as retriever by accessing the information pool. Each functional role will further be detailed by the equivalent subsystems in the DEUS architecture description.

E. Phases of Information Interchange

The three basic phases in DEUS information interchange are *contribution*, *repatriation*, and *publication*. Fig. 4 describes an information chain from the perspective of a new patient-related information in an EHR of a selected HCIS. The vertical separators delimit system borders: The EHR with the bundling component and export mechanism for digital cards resides in the contributor environment. The repatriation transfer is done by point-to-point network infrastructure.

The repatriation acceptance and assimilation is handled in the environment of the concerned person that also comprises the subscriber selection: A group-based selection is considered that allows to group physicians into teams so that repatriated digital cards are only published to team members. Enabling subscribers to express predicates over classifying attributes or over the content of the digital cards, providing channel- or content-based subscriptions [33], is not supported due to tremendous unsolved medical and legal implications.

The published digital card becomes part of the *foreign patient file*, the read-only subset of the personal patient

file at subscriber-side. In distributed environments, a released information cannot be revoked completely: Any receiver has the possibility to subsequently process the received information in its local systems (last optional step of HCIS import in Fig. 4).

F. Data Architecture

The overview of the basic artifacts in DEUS is provided in Fig. 5. The healthcare *domain layer* artifacts are digital cards based on HL7 CDA. A digital card is the subject of interchange. The *personal patient file* is the central patient-oriented information hub. The *foreign patient file* is the read-only representation at subscriber-side, being a collection of digital cards. If the patient manages distinct physician teams, the foreign patient files are a team-specific subset of the original comprehensive personal patient file. The *distributed patient folder* is managed by DEUS and consists of all subscribed foreign patient files of a DEUS account.

The identity of a DEUS account is essential to the addressing scheme of the artifacts because the DEUS account IDs from different actors take part in several artifact composite primary keys. A DEUS account is identified by an OpenID [34]. The optional application of XRIs⁵ is in preparation. Both identifier types can be used with XRDS⁶ service discovery. XRDS can be used for transfer protocol negotiation as it will later be elaborated in the description of the DEUS transfer layer access.

The DEUS storage and communication handles the artifacts in a general way, using the equivalent data arti-

⁵Extensible Resource Identifier,
<http://www.oasis-open.org/committees/xri>

⁶eXtensible Resource Descriptor Sequence,
<http://docs.oasis-open.org/xri/2.0/specs/xri-resolution-V2.0.html>

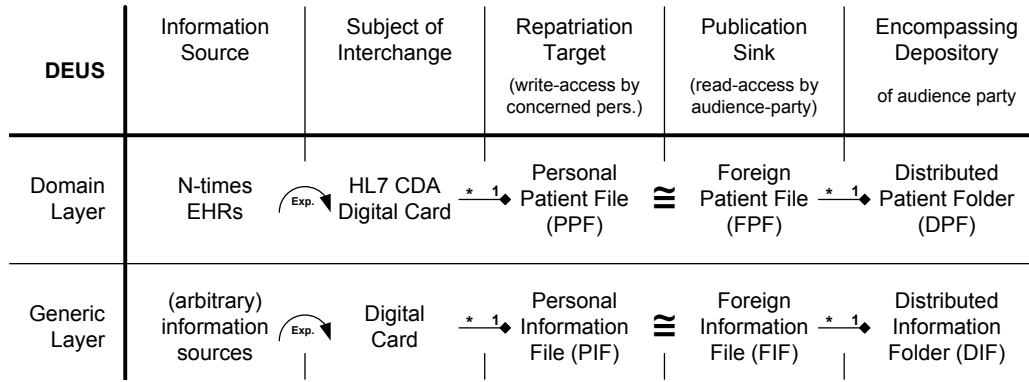


Fig. 5. The DEUS artifacts

facts of the *generic layer*. Any digital card is identified by a triple: The first component is a discriminator that is generated by the information provider. The second component is the account identity of the information provider itself. The third component is the account identity of the concerned person, about which the digital card provides information.

A *personal information file* (PIF) is identified only by the account identity of the concerned person it belongs to. The *foreign information file* (FIF) is identified by combining the concerned person's ID with the owning information consumer's ID. Finally, the *distributed information folder* (DIF) is owned and identified by the information consumer.

The actor-differentiated identification scheme is used by the persistence layer for each of the DEUS artifacts. It allows a DEUS node to host multiple DEUS accounts. This provides multitenancy, which enables a proxy institution to professionally and securely host DEUS accounts for multiple institutions. In conclusion, DEUS provides persistence that allows to scale discretionary from a central storage to a fully-distributed DEUS node topology.

G. From Cooperation to Collaboration

The basic DEUS scenario as described above is primarily a cooperative one and not a collaborative one. In cooperation, the participants are organizationally independent and the degree of interdependence remains low, although the information exchange is based on a mutual benefit. In contrast, collaborative scenarios require a team to achieve collective results that the participants would be incapable of accomplishing when working alone. Although collaborating partners remain organizationally independent, the degree of work-product interdependence is substantial.

Such a collaboration of healthcare professionals occurs for example in breast cancer treatment centers working with accredited partners. Here, an enforced patient interaction would be counterproductive to the collaborative workflow and is not feasible. Therefore, a specialized DEUS scenario can be found: Exchanging patient information without requiring a patient to decide about the distribution. A patient account that exists only for logical purposes could be specially trimmed-down, so that it would auto-accept any repatriation requests from certified contributors. The set of participants, comprising all accredited partners, could be templated and injected during application specific patient account creation as the list of contributors and subscribers. For cancer treatment the oncologists, radiologists, and the post-operative care actually require *all* available patient information unfiltered. Thus, the phase in which the patient account performs subscriber selection can be switched off by configuring the according subsystem to publish any digital card to all subscribers. Such a *virtual patient account* should be hosted by the breast cancer treatment center.

An organizationally closed environment is mandatory for such virtual patient accounts. Thus, this idea appears to be relevant for individual case environments, only, with moderate participant sets—not for life-long, large-scale patient files. In life-long scenarios the patient must be taken into account as mediator.

V. DEUS SYSTEM ARCHITECTURE

The DEUS system is decomposed vertically into tiers and sublayers. The horizontal decomposition into subsystems is deduced from the differentiation of DEUS functional roles in Fig. 3. Therefore, the modules have high functional cohesion and a distinct scope of work. As additional lateral decomposition, the Barker subsystem

provides cross-cutting functionality for system-to-user interactions. The resulting overall architecture is outlined in Fig. 6.

The neighbour systems in the presentation tier are not DEUS peers but institutional systems that access their DEUS system extension within an intranet by client/server access. DEUS peers communicate by arbitrary transfer protocols of the infrastructure tier.




Fig. 6. The DEUS architecture

A. Component Model

To enable clear interfaces and enforce module boundaries for side-effect free invocations, all pictured UML components are realized as distinct OSGi⁷ bundles. The OSGi framework provides a runtime system to support a defined lifecycle, including hot deployment, installation, dependency resolution and service registration. Component boundaries are controlled by explicitly specifying exported and imported interfaces and classes. Apache Maven⁸ was used as a build system to reflect the modularization in the project folder layout with Maven Multi-Modules and to support compile-time dependency resolution.

B. Basic Decomposition

The *presentation tier* offers a web portal for user interaction. Alternatively, neighbour systems inside the

institutional intranet can access the DEUS system by several remote invocation technologies that are provided by the *neighbour systems interface* (NSI) sublayer.

The *application tier* or *core layer* provides the NSI sublayer as well as the preparation sublayer, the Soul and both infrastructure access sublayers. The primarily disposed NSI remote invocation technology is a REST⁹ [35] interface, since the RESTful approach fits with the document-oriented perspective. Modules that provide NSI access by SOAP or RMI technology could be deployed additionally. All invocations to the subsystems that reside in the DEUS *Soul sublayer* are intercepted by the *preparation sublayer*. It checks for data validity and transforms the parameter types to domain data types, thereby shielding the Soul subsystems from calls with invalid data.

The *infrastructure tier* encompasses data stores and DEUS data schemas as well as transfer protocols for communication with remote DEUS nodes. The *infrastructure access sublayer* within the DEUS Core decouples from specific infrastructure technology. The *data access sublayer* provides an interface to the Soul that solely handles domain objects and is agnostic to any specific persistence technology. An implementation module provides a binding to a dedicated back-end data store. The *transfer access sublayer* provides the Soul with a facility to communicate with remote DEUS accounts. It is transparent to the Soul, whether the targeted account resides on another or on the same DEUS node.

The next three sections will detail the notable parts of DEUS architecture: the Soul subsystems and their cooperation, the transfer access sublayer with its plug-in architecture for binding arbitrary message-based communication middleware, and the interaction mechanism between the end-user, the subsystems, and the messaging.

C. Soul Subsystems

Fig. 7 depicts the DEUS Soul subsystem modules and their interfaces as well as their dependencies in cooperation. For lucidity, the delegations to the lateral Barker system, that provides system-to-user interactions, are omitted. The interplay between users, accounts, subsystems, and messages will be clarified later. The subsystems are grouped by the higher-level distribution roles. Each DEUS node implements all roles. The cooperation of the subsystem will be explained in the context of the contribution-repatriation-publication phases.

⁷Open Services Gateway initiative, <http://www.osgi.org>

⁸<http://maven.apache.org>

⁹Representational State Transfer

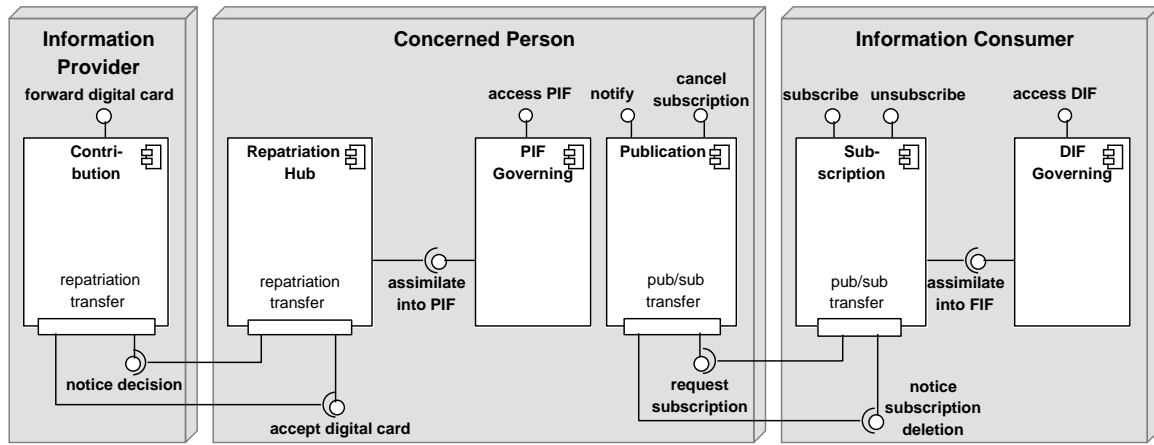


Fig. 7. DEUS Soul subsystems (without the lateral Barker)

The *Contribution* subsystem provides a facility to contribute digital cards that have been exported by an HCIS from its EHR into the DEUS system. For this purpose a `forwardToCP(DigitalCard)` method is exported by the NSI to institutional neighbour systems. The method takes a digital card, signs it, and forwards it to the concerned person's DEUS account using the DEUS peer transfer protocol infrastructure. The account identity of the concerned person is part of the composite identity of the digital card that has to be repatriated. The addressing is accomplished by using the account identity and resolving a transfer protocol ID by performing a handshake with the target account. The communication itself will be detailed in the next section. The *Repatriation Hub* of the receiving account accepts the digital card by an `accept(DigitalCard)` method. This subsystem presents the decision, whether to accept or decline the contribution, to the user by delegating the event to the Barker subsystem which appends it to an attention list presented to the user. The digital card is meanwhile persisted in a staging area. If the user confirms the decision about the digital card acceptance, the Barker subsystem triggers the *PIF-Governing* subsystem to migrate the digital card from the staging storage into the personal information file storage ("assimilate into PIF"). The PIF-Governing subsystem also exports its facility for personal information file access to the user-interface or to institutional neighbour systems.

The functionality around the traditional pub/sub capabilities of the DEUS system is comprised in the subsystems *Publication* and *Subscription* in Fig. 7: Establishing a publish-subscribe connection requires the subscriber to know the account ID of the publisher. The Subscription module implements a **subscribe** interface that is

exported by the NSI. It results in a signed **request subscription** to the addressed publication subsystem. The subscription request is delegated inside the Publication subsystem to the Barker, that presents the request to the user for decision. The necessary command messages are mapped to the resolved transfer protocol. The authentication relies on a pre-established PKI key exchange, which is discussed in the future work section. At the moment, an establishment is based on the account IDs. To complete the connection establishment at publication side, the publisher adds the subscriber ID to a *list of accepted subscribers*. The assignment of a subscriber into a specific publication group can be done during the acceptance of the request subscription. The decision is signed and sent back. To complete the connection establishment at subscription side, the subscriber adds the publisher ID to its *list of confirmed publishers*. An overview over the DEUS vocabulary for connection management is provided by fig. 8; the publisher initiated connection establishment is grey because it has not yet been implemented and is not represented by the current interfaces.

The last step in connection establishment concerns replication consistency: The concerned person account initially publishes historic digital cards to the new consumer. Different strategies can be applied: First option would be the selection of digital cards that are always initially published to any accepted subscribers, containing for example the patient's master data and basic medical information like allergies or medication. Second option would be the manual selection of digital cards by the concerned person, specifically for the new consumer. Third option would be to remember a list of digital cards that have been published to a publication group,

	establish connection		terminate connection	
	publisher initiated	subscriber initiated	publisher initiated	subscriber initiated
use case name	invite subscriber	subscribe to publisher	cancel subscription	unsubscribe from publisher
message name	subscription offer	subscription request	subscription cancel	unsubscribe
decision on opposite side	confirm/repel	grant/deny	—	—

Fig. 8. The DEUS vocabulary overview for pub/sub connection management

and to publish these digital cards to the new subscriber being added to the group. Another most trivial option would be to initially publish nothing. The first option is use-case globally, group-independant, and instance-independant; the second is instance-specific; the third is group-specific; the forth is do-nothing. More options, like sophisticated rule-based strategies, could be applied, but have not been evaluated at the moment.

Deleting an established publish-subscribe connection between two accounts can be initiated by the concerned person with a **cancel subscription** that is provided by the Publication subsystem. It can also be initiated by the information consumer, with an **unsubscribe** that is provided by the Subscription subsystem. The DEUS system will delete the respective partner from the list of confirmed publishers or accepted subscribers. The deletion of the foreign information file from the consumer actor's distributed information folder will not be enforced; such decisions will be bound to judicial boundary conditions. The concerned person can demand such deletion, either as part of the **cancel subscription** invocation or as part of the notification about the **unsubscribe**. The Subscription subsystem, on the consumer side, will use its Barker to delegate the decision about the information deletion to its user.

The subsystem *DIF-Governing* provides a facility to absorb a received digital card into the foreign information file. Both, the DIF-Governing and PIF-Governing subsystems provide access to their governed data for the web portal or for institutional neighbour systems.

D. Transfer Access

For being independent of a specific communication protocol, bindings to arbitrary protocols can be implemented to provide remote transfer. This is reflected

in the design of the transfer access subsystem which implements a plug-in architecture. To decouple the Soul sublayer from a specific protocol binding the *Transfer Core sublayer* is introduced. This sublayer handles registration of transfer protocol bindings, negotiation of transfer protocols between two DEUS nodes, and resolution of User IDs to transfer protocol IDs. It is outlined in Fig. 9, together with an exemplary binding at the bottom.

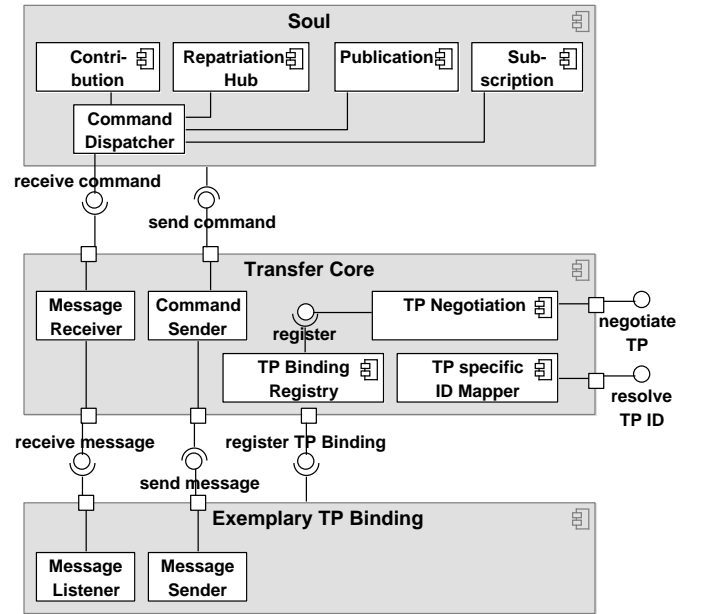


Fig. 9. Inner architecture of the transfer access sublayer

E. TP access layer mechanisms

Available transfer protocol bindings are registered at the TP Binding Registry on startup and added to the list of available transfer protocols available for protocol negotiation. If a subsystem from the Soul sublayer

needs to transfer a message to another DEUS account, it instruments the interface `send` command provided by Transfer Core. The component **Command Sender** first negotiates the transfer protocol that should be used: The list of supported protocols is remotely retrieved from the receiver of the command. These protocols are subsequently matched with the ones supported by the local DEUS node and a common transfer protocol is chosen. Both communication participants have priorities attached to their supported transfer protocols, which are taken into account during protocol negotiation.

Subsequently, the user ID of the remote DEUS account is resolved to the transfer protocol ID of the negotiated transfer protocol using the remote offered interface `resolve TP ID`. Following the *Command Message* pattern in [28, p. 145], the command to send and its parameters are marshalled and included in a message object. The `send message` interface of the chosen transfer protocol binding is thereupon used to deliver the message.

Registered protocol bindings are listening for remote messages being sent by other DEUS accounts. If a message is received by a binding, binding makes a callback to the `receive message` interface of Transfer Core. The message is unmarshalled, the contained command is extracted and passed to the Soul sublayer using its callback interface `receive command`.

The protocol bindings are realized as OSGi bundles, so that communication protocol bindings are registered at the startup of the OSGi bundle. Using this approach, transfer protocols can be plugged in at runtime and thus are instantaneously available for selection during transfer protocol negotiation.

F. Transfer Protocol Bindings

The transfer layer for DEUS account-to-account communication must support a *Point-to-Point Channel* [28, p. 103] for repatriation and a *Publish-Subscribe Channel* [28, p. 106]. The most basic transfer protocol binding is a local loopback protocol, for communication between user accounts that reside on the same DEUS node. The next binding module is based on a REST delivery that implements a simple application layer multicast. The REST-based multicast requires a minimal `POST` interface that accepts messages. Obviously, the simple REST approach does not provide *Guaranteed Delivery* [28, p. 122]. Assuming that DEUS nodes may fail and be offline for a certain period of time would require a sophisticated implementation of persistent queues. In-

stead, another protocol binding is based on XMPP¹⁰: The transfer access layer plug-in is based on an XMPP client¹¹ that acts as *Messaging Mapper* [28, p. 477] and instruments a local XMPP server¹² as persistent queue. The XMPP server takes any measures that delay and guarantee the delivery of messages directed to these nodes during their downtime.

Further transfer protocol bindings are in consideration: JMS¹³ could be implemented with a JMS client instrumenting configurable JMS providers; messaging middleware like MQSeries/WebSphere MQ or the Tibco service bus could augment the portfolio. This is not supported by REST and thus would need to be implemented using message queues on the publisher side. Yet, the abstraction level of the protocol to be bound should ease the mapping of DEUS high-level concepts: XMPP provides the abstraction of users and their accounts, being uniquely identified by an XMPP ID, as well as concepts of subscription management, message routing, and delivery. Support for a transfer protocol inherent multicast is also provided – by an XMPP extension (XEP-0060) that extends the original store-and-forward into a basic publish/subscribe message delivery. Furthermore, XMPP provides many official XMPP Extension Protocols including support for service discovery, in-band account registration, search, message receipts, reliable data transfer, time-sensitive delivery, expiration of transient messages and more.

G. User-Barker Interaction and Message Exchange

The message protocol comprises command messages for establishment and termination of repatriation relationships, document messages for the repatriation as well as command messages for subscription management. The document message for publication is the same as for the repatriation, the difference lies within the transferred digital card which is signed twice. There are three types of command messages: request, decision, and notice messages.

Fig. 10 outlines the interplay between the user actors, the Subscription/Publication subsystems, the messaging, and the Barker subsystem. Prof. Bob subscribes to Alice, invoking the `subscribe(userID)` with Alice's account ID. A `request subscription` message is sent to Alice's account and is dispatched to the Publication

¹⁰Extensible Messaging and Presence Protocol – the core protocol of the Jabber instant messaging technology

¹¹DEUS is using the Smack library

¹²DEUS is using OpenFire

¹³Java Message Service, JSR-914

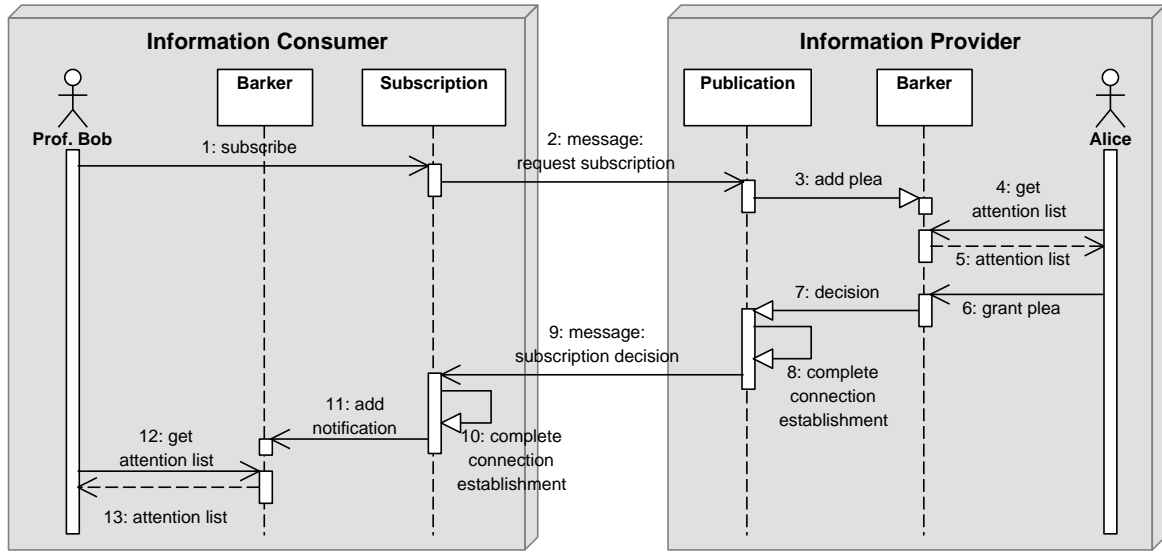


Fig. 10. Request subscription sequence

subsystem. The transition from a command (\rightarrow messaging) to an attention element (\rightarrow human-interface) is done by the Publication subsystem that adds a *plea* object to the Barker. The Barker integrates the plea as *attention element* in the *attention list*. Alice, upon next login, displays the attention list and can grant or deny any plea. The decision is delegated by the Barker to the Publication subsystem. The example in Fig. 10 describes the “grant” case, therefore the connection establishment is completed (in the “deny” case, any persistently staged objects for a request are cleaned up). The decision is transferred to the requesting account, which performs its own request completion and adds the decision as *notification* to the Barker. Again, the Barker absorbs the notification in the attention list. Upon reading a notification, the user can mark it, so that it will not be displayed any more in the attention list. The Barker stores a complete history of any notification and decision elements.

VI. DISTRIBUTION OVER CENTRALIZATION

A risk in instrumenting a central content storage, like German D2D¹⁴ or Google Health¹⁵, is an information leak that potentially involves all patients. This is not comparable to any possible abuse scenario in today’s paper-based infrastructure: No current healthcare institution hoards information about so many patients as will do any centralized solution for inter-institutional scenarios.

¹⁴Doctor to Doctor, <http://www.d2d.de>, based on PaDok cryptographic infrastructure [36]

¹⁵<http://www.google.com/health>

The distributed approach mirrors the current state in paper-based working practice: The patient information is available only to the directly involved healthcare systems. As a result, the consequences of a security breach are limited to a fraction.

Yet, even the smallest security breach still remains fatal due to the criticality of the involved information. Hospital infrastructure commonly hosts electronic patient information and already applies profound security measures, e.g. [37], but primary care participants might not be accustomed to the required security standards. Therefore, a distributed approach has to allow proxy institutions to professionally and securely host accounts, i.e. for primary care participants or for patient participants. For example, the health insurance funds or the associations of statutory health insurance physicians could provide account hosting. In conclusion, participating peer systems in the globally distributed healthcare environment are required to locally adhere to a multi-tenant data architecture as it is provided by DEUS.

In medical care, the availability of information at the right time and at the right location (the “point-of-care”) is crucial [38]. In contrast to the DEUS push-based information distribution, a pull-based approach could allow a HCIS to query a patient account ad-hoc when information is needed. Since a local replicate of the electronic patient file is absent, its advantages are lost: These benefits include a reduced response time, because a remote call is avoided, which elevates end-user acceptance. Furthermore, the absence of a local copy requires the continuous availability of the patient DEUS node

for information provision, which can not be guaranteed. Nevertheless, even a pull-based approach requires a trust relationship to be built between the concerned person and the information consumer in advance, analogous to the DEUS architecture. In conclusion, the DEUS push-based approach is more efficient than pull-based solutions.

VII. RELATED WORK

Besides data integration, the *functional integration* has to provide operative interoperability between information systems. Functional integration is addressed on a syntactic technical level and a semantic domain-specific level. On the technical level, mature solutions exist, i.e. instrumenting component models like EJB [39] and .Net or lightweight ones like Spring and OSGi as well as remote procedure calls like IIOP, RMI, or XML-based protocols like SOAP. Semantic functional integration requires interface and protocol standards for medical services.

Standards for electronic information exchange between the *practice management systems* of the primary care and the hospitals and institutions of the secondary care are rare. No universal exchange protocol and format exists for interchange of referral vouchers and discharge letters. In Germany, the governmental project “Elektronische Gesundheitskarte” (cf. sect. VIII-E) has not provided a solution for the issue since the project’s outset in 2003. Effective platforms like D2D require a central server for document handover.

Existing protocol standards for information exchange in distributed healthcare scenarios mainly focus on hospitals of the secondary care. These standards encompass the earliest efforts in this area due to the complexity of a major hospital and its need for inter-sectional information exchange. Available standards in healthcare, like DICOM¹⁶ [40] or the *cross-enterprise document sharing* (XDS) [41] standard from IHE¹⁷ [42], focus on the information exchange between hospital information systems (HIS) cooperating with ancillary systems like radiology information system (RIS), cardiology, and pathology systems, or laboratory information management system (LIMS). Furthermore, there exist tailor-made regional integration efforts, but they are based on a central database system (DBS) with distributed transaction systems and diverse communication middleware. Even wide-area RHIN architectures like HYGEIANet [43] on the island Crete require a federated database schema

[44] and are therefore tightly-coupled by their complex infrastructure being inadequate for transregional scaling.

Solving the information exchange in healthcare in a document-oriented fashion seems to be targeted by IHE XDS which allows for distributed document repositories and access delegation. For gaining experience with such standards, we implemented *XdsRig*, an IHE XDS test stand environment; it is published in [45] and available as open source. Yet, in order to find documents in such a repository, a single central document registry is specified, reusing ebXML registry methodology to provide a centralized method of indexing documents. The central registry is a global system node that allows queries and that delegates the access to referenced documents to the original document repositories. Such architecture targets complex hospitals with associated ancillary systems and is even applicable to regional integration efforts, but fails for nationwide application due to its centralized approach.

In RHINs, several hospitals and ancillary institutions normally develop a shared set of IT services for information exchange. Neutral organizations like IHE try to establish interaction standards in format and protocol, little by little, based on standards like DICOM or HL7. Concurring standards currently in use can be placed into a classifying matrix of integration [3]. In conclusion, a “semantic gap” is revealed, that is not covered by standards concerning the semantic integration.

Another term that describes DEUS in regard to its patient-centered approach is *Personal Health Record* (PHR), which is an old term that appeared first in 1978 [46]. The term has gained new momentum. However, solutions that are currently marketed as a PHR are Web-based and centralized; Sittig provides an overview in [47]. Well-known examples are systems like Dossia¹⁸, PatientsLikeMe, Microsoft HealthVault¹⁹, and the discontinued Google Health. Whereas the original PHR concept naturally emphasizes on a paper perspective (like we do in DEUS) the emerging Web-based PHR systems, however, are intrinsically database-oriented with semantically rich data models. Rich data models imply integration efforts in a proprietary way and in a fragmented fashion. In conclusion, they become “yet another data silo” or even a “data tomb”: When Google closed their Google Health service, the users got a database dump in form of an XML export. Granted, such an export was better than expected (and other PHR systems do

¹⁶Digital Imaging and Communications in Medicine

¹⁷Integrating the Healthcare Enterprise, www.ihe.net

¹⁸Dossia is limited to employees of the few signee companies.

¹⁹Microsoft HealthVault is not available in Germany.

not provide an export), but ultimately it is futile because other PHR vendors have incompatible data models and an import is not provided. The IHE started to attend issues of PHR integration by adopting it into their Patient Care Coordination (PCC) efforts [48]. DEUS can be considered as a distributed PHR; the document-orientation guarantees that it is possible, with the collection of electronic documents that constitute an DEUS patient file, to export it into the file system at the user desktop at any time.

In the context of integrated ambient systems in healthcare there exist research projects for large-scale distributed middleware [49]. They demonstrate that XMPP is even applicable to real-time healthcare scenarios with event-based body sensor networks. Yet, the focus is completely different; they do not consider the three party scenario of DEUS and the authoritative integration of the patient in information distribution.

VIII. FUTURE WORK

There are several open issues on different conceptual levels and technological layers.

A. First Contact

In DEUS, a sophisticated solution for the first contact problem is required: The OpenID identification of each DEUS account could be integrated in a distributed master patient index (MPI) for the simplification of account connection establishment. MPI systems like IHE PIX²⁰ [41] or OMG²¹ PIDS²² [50] instrument hierarchical federation with central system nodes and are not applicable in distributed environments. Therefore, a loosely-coupled distributed patient identification service for inter-institutional purpose is required.

B. Content Filtering

An altogether open issue in DEUS is digital card filtration because it requires a deep knowledge of medical taxonomies and ontologies. An approach that enables subscribers to express predicates over the content of the digital cards could provide content-based subscriptions [33]. Subscriber-defined predicates over well-known classification attributes of published digital cards would provide channel-based subscriptions. The same filter mechanisms could even be applied as further empowerment of the concerned person: Publisher-defined filters as a *share set* for each individual subscriber could enable the concerned person with sophisticated

filter mechanisms for extended control over published data. Yet, any signed digital card can only be filtered completely or not at all; filtering information content of the repatriated digital card can only be applied for unsigned information with limited trust, for example contact information like telephone numbers. However, diverse kinds of filtering digital card information would technically be possible, but would have tremendous unsolved medical and legal implications. Singh et al. [51] are considering similar concepts for the English National Health Service (NHS).

Another consideration about the content is the medical consistency of its information. With the OXDBS project [52], we extended a native XML database system with validation by consistency checking of OWL-DL ontologies. A possible integration of such a facility into our exchange platform seems promising.

C. Group Management

The subscriber selection requires some efforts from the concerned person, in form of physician group management. Yet, there exists another approach to the subscriber selection: The contributor could provide reliable hints to the concerned person about which subscribers are appropriate in relation to a digital card. Formalizing such hinting could automate the subscriber selection effectively; the PKI trust between the concerned person's and information provider's account has to be established for repatriation anyway. In healthcare, the publication hints would require a classification system for the medical roles that a participant can take in treatment supply chains.

D. Process Model

The document-oriented content distribution platform can be considered as a foundation for inter-institutional process support. A distributed process model that adheres to the diagnostic-therapeutic cycle as coarse-grained institutional reference from working practice is not yet available. With a project named DMPS [53], we worked on a platform that provides a shared process ID and maintains process history in order-entry and result-reporting scenarios. The process history also provides information about the pretreatment or mutual treatment providers. Additional characteristics for workflow support needs yet to be integrated into DEUS like a formal model for pan-diagnostic and pan-therapeutic processes.

E. Alignment to German eGK

The German governmental project "Elektronische Gesundheitskarte" (eGK) was initiated after the German par-

²⁰Patient Identifier Cross-Referencing

²¹Object Management Group

²²Patient Identification Service

liament passed a bill, the “GKV-Modernisierungsgesetz”, to modernize health insurance cards in 2003. The objective of the modernization plans was a telematics infrastructure for interconnecting healthcare facilities [54]. With some new electronic health smartcards as its cryptographic foundation.

The primary eGK function is a public-key infrastructure (PKI) with smartcards and card readers as well as connectors for online communication. On top of this basic communication platform, the *mandatory eGK applications* comprise health insurance master data (eGK application VSDD: “Versichertenstammdatendienst”) and electronic prescriptions (eGK application VODD: “Verordnungsdatendienst”). In addition, so called *optional eGK applications* had been described in the official documentation [55]. It comprised a requirement analysis for the following applications: a patient’s pharmaceuticals history (application AMDD: “Arzneimitteldokumentationsdienst”), the patient’s emergency health data (application NFDD: “Notfalldatendienst”), an infrastructure for physicians’ result reporting (application ABD: “Arztbriefdienst”), an infrastructure for physician-provided electronic patient files (application EPAD: “Elektronischer Patientenaktendienst”), and an infrastructure for patient-provided information to his/her electronic file (application PDD: “Patientendatendienst”). No system specification actually exists to any of these optional eGK applications. As mentioned in sect. VII, the eGK has not provided any solution, neither for its mandatory nor for its optional applications, since the project’s outset in 2003.

In eGK terms, DEUS represents a combined solution to EPAD & PDD. (Our co-project DMPS represents a solution to the ABD.) Yet, an alignment with the specific eGK concepts and technological infrastructure remains an open issue.

F. Abstraction from Public-Key Infrastructures

The research does not focus on PKI, but relies on existing PKIs, like eGK, in order to build authentication and authorization, during the establishment of contribute/repatriate or pub/sub relationships. Yet, there exists competing PKI infrastructures in healthcare, like PaDok [36] or IHE ATNA [41]. In addition, an integration of general messaging middleware with role-based access control [56] seems promising²³, too. To sign digital cards with arbitrary PKI technology, a generalized DEUS PKI-

component integration is required that provides independence of the various PKI infrastructures.

G. Case Files

Maintaining patient files requires a system extension at each participating site. For ad-hoc establishment of data exchange without preceding system integration we have done considerable work on distributed case files. The umbrella project is α -Flow, where the central artifact is the α -Doc that is a distributed case file. The “ α ” stands for “active” as in *active documents* [59]. The general idea is to extend electronic documents with process status information to achieve workflow support. Further information about the various aspects of the α -Flow project can be found in [60]–[67]. A seamless transition from α -Flow case files to DEUS patient files is prepared for by the conceptual similarity in their data model. It would be required to include the considerations for case-related adaption of DEUS (from sect. IV-G) into an integrated patient-&-case-file approach. Hence, it remains an open issue to integrate both systems.

IX. CONCLUSION

The DEUS architecture achieves a document-oriented process support between strict autonomous institutions following the paper-based work practice as reference model. For this purpose, DEUS applies document-orientation based on a publish-subscribe design. The essential argument for document-oriented integration over interface- or message-oriented integration lies in its capacity to support deferred system design. Deferred system design is necessary for healthcare information systems due to their adaptive-evolutionary character.

The DEUS architecture appoints the patient as integral participant of the information supply chain. The warranty of data protection, imperative in healthcare, requires PKI integration on technological level but additionally requires a repatriation phase with end-user interception capability, resulting in a mediated publish-subscribe architecture. The DEUS architecture supports local autonomy, multi-tenant data architecture, platform independence, and loose coupling.

The current DEUS platform has several unsolved aspects. It requires concepts for the first contact problem, content filtering, and a connection to PKI infrastructure. It would also benefit from a technical alignment with the German national eGK project. DEUS is a platform for patient files in which patients are empowered to control data distribution. Support for individual cases requires that the patient file artefact model is refined

²³For example, OASIS [57] is a role-based access control, with HERMES [58] as overarching project that integrates it.

one level, in the future, in order to model case files as subsets of a patient file. Support for the planning and monitoring of case-centric treatment progress in form of work lists requires a document-oriented process model. Furthermore, DEUS is a light-weight system extension at each participating site and requires an installation setup. In contrast, α -Flow implements the active documents metaphor such that it can be deployed without prior software installation. Currently, the α -Flow approach to distributed case files, excluding the patient from the data flow, is better suited for true ad-hoc establishment of data exchange between healthcare professionals. DEUS is better suited for life-long patient files. In conclusion, the current DEUS implementation is a necessary first step towards a distributed platform for trans-sectional, long-term, patient-centered healthcare documentation.

REFERENCES

- [1] L. Leape, D. Bates, D. Cullen, J. Cooper, H. Demonaco, T. Gallivan, R. Hallisey, J. Ives, N. Laird, and G. Laffel, "Systems analysis of adverse drug events. ADE Prevention Study Group," *JAMA*, vol. 274, no. 1, pp. 35–43, 1995.
- [2] R. Lenz and K. Kuhn, "Towards a continuous evolution and adaptation of information systems in healthcare," *Int J Med Inf*, vol. 73, no. 1, pp. 75–89, July 2003.
- [3] R. Lenz, M. Beyer, and K. Kuhn, "Semantic integration in healthcare networks," *Int J Med Inf*, vol. 76, no. 2-3, pp. 201–207, 2006.
- [4] C. P. Neumann, F. Rampp, M. Daum, and R. Lenz, "A Mediated Publish-Subscribe System for Inter-Institutional Process Support in Healthcare," in *Proc of the 3rd ACM Int'l Conf on Distributed Event-Based Systems (DEBS 2009)*, Nashville, TN, USA, Jul. 2009.
- [5] C. van Walraven, M. Taljaard, C. Bell, E. Etchells, K. Zarnke, I. Stiell, and A. Forster, "Information exchange among physicians caring for the same patient in the community," *Canadian Medical Association Journal*, vol. 179, no. 10, p. 1013, 2008.
- [6] K. Lorig, D. Sobel, D. Laurent, and V. Gonzalez, *Living a Healthy Life With Chronic Conditions: Self-management of Heart Disease, Arthritis, Diabetes, Asthma, Bronchitis, Emphysema & Others*. Bull Publishing Company, 2000.
- [7] M. Tattersall, P. Butow, J. Brown, and J. Thompson, "Improving doctors' letters," *Med J Aust*, vol. 177, no. 9, pp. 516–20, 2002.
- [8] J. Samers, "Report on Integrated Care in Advanced Cancer Project," Inner and Eastern Melbourne BreastCare Consortium, Tech. Rep., mar 2004. [Online]. Available: <http://www.health.vic.gov.au/breastcare/downloads/integratedcare.pdf>
- [9] R. Lenz, M. Beyer, C. Meiler, S. Jablonski, and K. Kuhn, "Informationsintegration in Gesundheitsversorgungsnetzen," *Informatik-Spektrum*, vol. 28, no. 2, pp. 105–119, 2005.
- [10] R. Fletcher, M. O'Malley, S. Fletcher, J. Earp, and J. Alexander, "Measuring the continuity and coordination of medical care in a system involving multiple providers," *Med Care*, vol. 22, pp. 403–411, May 1984.
- [11] M. Müller, F. Ückert, T. Bürkle, and H. Prokosch, "Cross-institutional data exchange using the clinical document architecture (CDA)," *International journal of medical informatics*, vol. 74, no. 2-4, pp. 245–256, 2005.
- [12] M. Williams, G. Venters, and D. Marwick, "Developing a regional healthcare information network," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 5, no. 2, pp. 177–180, 2001.
- [13] S. Rothschild and S. Lapidus, "Virtual Integrated Practice: Integrating Teams and Technology to Manage Chronic Disease in Primary Care," *Journal of Medical Systems*, vol. 27, no. 1, pp. 85–93, 2003.
- [14] S. Brucker, U. Krainick, M. Bamberg, B. Aydeniz, U. Wagner, A. DuBois, C. Claussen, R. Kreienberg, and D. Wallwiener, "Rationale, funktionelles Konzept, Definition und Zertifizierung," *Der Gynäkologe*, vol. 10, p. 862, 2003.
- [15] J. Powell and I. Buchan, "Electronic Health Records Should Support Clinical Research," *Journal of Medical Internet Research*, vol. 7, no. 1, 2005.
- [16] L. Alschuler and K. U. Heitmann, "CDA Introductory Tutorial," in *HL7 International CDA Conference*, Oct. 2002. [Online]. Available: <http://www.hl7.de/cda2002/absbiopres/cdaintro.pdf>
- [17] B. Sippel, "Evaluation und Integration von Standards zum Datenaustausch im medizinischen Umfeld," Diplomarbeit, Friedrich-Alexander-Universität Erlangen-Nürnberg, Mar. 2005.
- [18] C. McDonald, J. Marc Overhage, P. Dexter, B. Takesue, and J. Suico, "What is done, what is needed and what is realistic to expect from medical informatics standards," *International Journal of Medical Informatics*, vol. 48, no. 1-3, pp. 5–12, 1998.
- [19] A. Zaiß, *Überleitungstabelle zwischen ICD-9 und ICD-10*. Deutscher Ärzte-Verlag, Köln, 1996.
- [20] M. Stearns, C. Price, K. Spackman, and A. Wang, "SNOMED clinical terms: overview of the development process and project status," in *Proc AMIA Symp*, vol. 662, 2001, p. 6.
- [21] A. Forrey, C. McDonald, G. DeMoor, S. Huff, D. Leavelle, D. Leland, T. Fiers, L. Charles, B. Griffin, F. Stalling *et al.*, "Logical observation identifier names and codes (LOINC) database: a public use set of codes and names for electronic reporting of clinical laboratory test results," *Clinical Chemistry*, vol. 42, no. 1, pp. 81–90, 1996.
- [22] HL7v2, "ANSI/HL7 V2.6-2007," http://www.hl7.org/Library/standards_non1.htm.
- [23] R. Dolin, L. Alschuler, C. Beebe, P. Biron, S. Boyer, D. Essin, E. Kimber, T. Lincoln, and J. Mattison, "The HL7 Clinical Document Architecture," *Journal of the American Medical Informatics Association*, vol. 8, no. 6, pp. 552–569, 2001.
- [24] K. Heitmann, R. Schweiger, and J. Dudeck, "Discharge and referral data exchange using global standards – the SCIPHOX project in Germany," *International Journal of Medical Informatics*, vol. 70, no. 2-3, pp. 195–203, 2003.
- [25] eArztbrief – D2D Telematik-Plattform der Kassenärztlichen Vereinigungen, "Elektronischer Arztbrief im D2D-System," <http://www.d2d.de/index.php?id=17>, 2005.
- [26] R. Lenz and K. Kuhn, "A strategic approach for business-IT alignment in health information systems," *Lecture notes in computer science*, pp. 178–195, 2003.
- [27] N. Patel, *Adaptive Evolutionary Information Systems*. Idea Group Inc, 2002.
- [28] G. Hohpe and B. Woolf, *Enterprise integration patterns: Designing, building, and deploying messaging solutions*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 2003.
- [29] S. Rozsnyai, J. Schiefer, and A. Schatten, "Concepts and Models for Typing Events for Event-Based Systems," in *International conference on Distributed event-based systems (DEBS'07)*, 2007. [Online]. Available: <http://cococon.ifs.tuwien.ac.at/pub/debs/debs2007.pdf>

- [30] R. Lenz, "Information Systems in Healthcare – state and steps towards sustainability," in *IMIA Yearbook of Medical Informatics*, A. Geissbuhler and C. Kulikowski, Eds. Stuttgart: Schattauer, 2009, accepted for publication.
- [31] Eclipse project, "Higgins Open Source Identity Framework," <http://www.eclipse.org/higgins/>.
- [32] eHealth, "Ministerial Declaration, Brussels," http://europa.eu.int/information_society/eeurope/ehealth/conference/2003/doc/min_dec_22_may_03.pdf, May 2003.
- [33] A. Carzaniga and A. Wolf, "Content-Based Networking: A New Communication Infrastructure," *Lecture Notes in Computer Science*, vol. 2538, pp. 59–68, 2002.
- [34] D. Recordon and D. Reed, "OpenID 2.0: a platform for user-centric identity management," in *Proceedings of the second ACM workshop on Digital identity management*. ACM Press New York, NY, USA, 2006, pp. 11–16.
- [35] R. Fielding, "Architectural Styles and the Design of Network-based Software Architectures," Ph.D. dissertation, University of California, 2000.
- [36] Fraunhofer Institut für Biomedizinische Technik (IBMT), "PaDok – Patientenbegleitende Dokumentation," http://www.ibmt.fraunhofer.de/fhg/Images/MT_padoknetzkonzept_de_tcm266-68980.pdf, 2000.
- [37] J. Vazquez-Naya, J. Loureiro, J. Calle, J. Vidal, and A. Sierra, "Necessary security mechanisms in a PACS DICOM access system with Web technology," *Journal of Digital Imaging*, vol. 15, pp. 107–111, 2002.
- [38] J. Anderson, "Clearing the way for physicians' use of clinical information systems," *Communications of the ACM*, vol. 40, no. 8, pp. 83–90, 1997.
- [39] M. Meyerhöfer and C. Neumann, "TestEJB – a measurement framework for EJBs," in *Proc of the 7th Int'l Symposium on Component-Based Software Engineering (CBSE'04) in conjunction with the 26th Int'l Conf on Software Engineering (ICSE'04)*, ser. Lecture Notes in Computer Science, I. Crnkovic, Ed., vol. 3054. Edinburgh, UK: Springer, Berlin, DE, May 2004, pp. 294–301.
- [40] W. Bidgood, S. Horii, F. Prior, and D. Van Syckle, "Understanding and Using DICOM, the Data Interchange Standard for Biomedical Imaging," *Journal of the American Medical Informatics Association*, vol. 4, no. 3, pp. 199–212, 1997.
- [41] ACC, HIMSS, and RSNA, *IHE IT Infrastructure Technical Framework, vol. 1 (ITI TF-1): Integration Profiles, Rev 4.0*, IHE Std., Rev. 4.0, Aug. 2007. [Online]. Available: http://www.ihe.net/Technical_Framework/upload/IHE_ITI_TF_4_0_Vol1_FT_2007_08_22.pdf
- [42] E. Siegel and D. Channin, "Integrating the Healthcare Enterprise: A Primer 1 Part 1. Introduction," *RadioGraphics*, vol. 21, no. 5, pp. 1339–1341, 2001.
- [43] M. Tsiknakis, D. Katehakis, and S. Orphanoudakis, "An open, component-based information infrastructure for integrated health information networks," *International Journal of Medical Informatics*, vol. 68, no. 1-3, pp. 3–26, 2002.
- [44] D. Katehakis, M. Tsiknakis, and S. Orphanoudakis, "Enabling Components of HYGIEAnet," in *Proc. of TEPR*, 2001, pp. 146–153.
- [45] C. P. Neumann, F. Wagner, and R. Lenz, "XdsRig – Eine Open-Source IHE XDS Testumgebung," in *Tagungsband der 54. GMDS-Jahrestagung*. Essen, DE: Deutsche Gesellschaft für Medizinische Informatik, Biometrie und Epidemiologie (GMDS), Sep. 2009.
- [46] G. Britain, "Computerisation of personal health records," *Health Visit*, vol. 51, p. 227, 1978.
- [47] D. F. Sittig, "Personal health records on the internet: a snapshot of the pioneers at the end of the 20th century," *International Journal of Medical Informatics*, vol. 65, no. 1, pp. 1–6, 2002.
- [48] IHE: Integrating the Healthcare Enterprise, "IHE Patient Care Coordination (PCC), Technical Framework, Volume 1, Revision 6.0," www.ihe.net/Technical_framework/upload/IHE_PCC_TF_Rev6-0_Vol_1_2010-08-30.pdf, Aug. 2010.
- [49] W. Labidi, J. Susini, P. Paradinas, and M. Setton, "XMPP based Health Care Integrated Ambient Systems Middleware," in *Developing Ambient Intelligence: Proceedings of the Second International Conference on Ambient Intelligence Developments (Ami. D'07)*. Springer, 2008, p. 92.
- [50] Object Management Group, Inc. (OMG), "Person Identification Service (PIDS) Specification," <http://www.omg.org/technology/documents/formal/index.htm>, April 2001.
- [51] J. Singh, L. Vargas, J. Bacon, and K. Moody, "Policy-Based Information Sharing in Publish/Subscribe Middleware," in *IEEE Workshop on Policies for Distributed Systems and Networks (POLICY)*, 2008, pp. 137–144.
- [52] C. P. Neumann, T. Fischer, and R. Lenz, "OXDBS – Extension of a native XML Database System with Validation by Consistency Checking of OWL-DL Ontologies," in *Proc of the 14th International Database Engineering & Applications Symposium (IDEAS'10)*, Montreal, QC, CA, Aug. 2010.
- [53] C. P. Neumann and R. Lenz, "A Light-Weight System Extension Supporting Document-based Processes in Healthcare," in *Proc of the 3rd Int'l Workshop on Process-oriented Information Systems in Healthcare (ProHealth'09) in conjunction with the 7th Int'l Conf on Business Process Management (BPM'09)*, Ulm, DE, Sep. 2009.
- [54] T. Weichert, "Die elektronische Gesundheitskarte," *Datenschutz und Datensicherheit*, vol. 28, no. 7, pp. 391–403, 2004.
- [55] Fraunhofer Institut, "Fachlogische Modellierung und spezifische Anwendungsdienste der elektronischen Gesundheitskarte," Mar. 2005. [Online]. Available: <http://www.ehealthopen.com/FuE/PDF/eGK-Fachmodlanw-v1.0.pdf>
- [56] A. Belokosztolszki, D. Eysers, P. Pietzuch, J. Bacon, and K. Moody, "Role-Based Access Control for Publish/Subscribe Middleware Architectures," in *2nd International Workshop on Distributed Event-Based Systems (DEBS'03)*, 2003, pp. 1–8.
- [57] J. Bacon, K. Moody, and W. Yao, "A model of OASIS role-based access control and its support for active security," *ACM Transactions on Information and System Security (TISSEC)*, vol. 5, no. 4, pp. 492–540, 2002.
- [58] P. Pietzuch and J. Bacon, "Hermes: A Distributed Event-Based Middleware Architecture," in *1st International Workshop on Distributed Event-Based Systems (DEBS)*, 2002.
- [59] A. LaMarca, W. Edwards, P. Dourish, J. Lamping, I. Smith, and J. Thornton, "Taking the work out of workflow: mechanisms for document-centered collaboration," in *Proceedings of the sixth conference on European Conference on Computer Supported Cooperative Work*. Kluwer Academic Publishers Norwell, MA, USA, 1999, pp. 1–20.
- [60] C. P. Neumann and R. Lenz, "alpha-Flow: A Document-based Approach to Inter-Institutional Process Support in Healthcare," in *Proc of the 3rd Int'l Workshop on Process-oriented Information Systems in Healthcare (ProHealth'09) in conjunction with the 7th Int'l Conf on Business Process Management (BPM'09)*, Ulm, DE, Sep. 2009.
- [61] —, "The alpha-Flow Use-Case of Breast Cancer Treatment – Modeling Inter-Institutional Healthcare Workflows by Active Documents," in *Proc of the 8th Int'l Workshop on Agent-based Computing for Enterprise Collaboration (ACEC) at the 19th*

Int'l Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE 2010), Larissa, GR, Jun. 2010.

- [62] A. Todorova and C. P. Neumann, "alpha-Props: A Rule-Based Approach to 'Active Properties' for Document-Oriented Process Support in Inter-Institutional Environments," in *Lecture Notes in Informatics (LNI) Seminars 10 / Informatiktage 2011*, L. Porada, Ed. Gesellschaft für Informatik e.V. (GI), Mar. 2011.
- [63] C. P. Neumann, P. K. Schwab, A. M. Wahl, and R. Lenz, "alpha-Adaptive: Evolutionary Workflow Metadata in Distributed Document-Oriented Process Management," in *Proc of the 4th Int'l Workshop on Process-oriented Information Systems in Healthcare (ProHealth'11) in conjunction with the 9th Int'l Conf on Business Process Management (BPM'11)*, Clermont-Ferrand, FR, Aug. 2011.
- [64] A. M. Wahl and C. P. Neumann, "alpha-OffSync: An Offline-Capable Synchronization Approach for Distributed Document-Oriented Process Management in Healthcare," in *Lecture Notes in Informatics (LNI) Seminars 11 / Informatiktage 2012*, L. Porada, Ed. Gesellschaft für Informatik e.V. (GI), Mar. 2012.
- [65] C. P. Neumann, A. M. Wahl, and R. Lenz, "Adaptive Version Clocks and the OffSync Protocol," in *Proc of the 10th IEEE Int'l Symposium on Parallel and Distributed Processing with Applications (ISPA-12)*, Madrid, Spain, Jul. 2012, accepted for publication.
- [66] C. P. Neumann, S. A. Hady, and R. Lenz, "Hydra Version Control System," in *Proc of the 10th IEEE Int'l Symposium on Parallel and Distributed Processing with Applications (ISPA-12)*, Madrid, Spain, Jul. 2012, accepted for publication.
- [67] C. P. Neumann and R. Lenz, "The alpha-Flow Approach to Inter-Institutional Process Support in Healthcare," *International Journal of Knowledge-Based Organizations (IJKBO)*, vol. 2, no. 3, 2012, accepted for publication.