



Implementing a semantically-aware SOA at NCI:

From Theory to Reality

Charlie Mead (meadch@mail.nih.gov)
*Senior Advisor to the Director,
NCI Center for Biomedical Informatics and
Information Technology (NCI CBIT)*

OMG: SOA in Healthcare
Sydney, AU
January 2011

sSOA at NCI: *From Theory to Reality*



The National Cancer Institute's (NCI) Cancer Bioinformatics Grid (caBIG®) project was launched approximately 7 years ago with the goal of providing an infrastructure that would use contemporary IT tools and technologies to link clinicians, researchers, and patients in the oncology community via a seamless infrastructure, a "World Wide Web for Oncology" that would integrate information and function across the Translational Medicine Continuum, i.e. "from bedside to bench and back...and beyond."

As experience was gained with a deployed code base into a distributed community, it became clear that a second-generation of caBIG® that leveraged maturing SOA knowledge, frameworks, capabilities, and technologies would substantially increase the value proposition of caBIG® for all of its stakeholders. The trajectory was dubbed as a "semantically-ware SOA (sSOA)."

As a consequence, the next-generation of caBIG® was launched in early 2010. This talk will briefly describe the NCI and its goals as the SOA development context for the next-generation of caBIG®, the theoretical reasons that led NCI to make the decision to pursue SOA as its underlying and overarching distributed architecture paradigm, and the realities and lessons learned over the past year. We expect that at least some of the latter will be of interest to an organization interested in developing software components that enable computable semantic interoperability from both an informational/static and behavioral/dynamic semantic perspective.

sSOA AT NCI: *THE CONTEXT*

sSOA at NCI: The Reality



- **sSOA (NCI) == ?? == SOA**
- **“sSOA” is an NCI moniker used to emphasize the importance of unambiguously defining (and focusing on) standards-derived semantics**

CBIIT Development: *Bringing Standards to the Community*



Standards
Development
Organizations



Center for Biomedical
Informatics &
Information
Technology



Software
Development
Companies and Open
Source Community



- CBIIT consumes and develops conformant standards-based specifications to resolve business problems
- CBIIT validates the applicability of specifications via reference implementations
- CBIIT informs the commercial and open-source vendor communities of the “state of the art” by deploying/handing off its reference implementations.



caBIG[®]:

a *Network* to Transform Biomedicine



caBIG[®] is a virtual network of interconnected organizations, individuals and data that collectively redefines how research is conducted, care is provided, and patients/participants interact with the biomedical research enterprise.



The Cancer Community Seeks to*



- Leverage biomedical data for research
- Connect workflows
- Streamline data collection
- Increase accuracy of processes
- Share data to enable collaborative research
- Perform complex analysis across multiple data sets (repurpose data)
- Identify best practices



* Based on 2004 Cancer Centers survey

caBIG[®] “three C’s” (circa 2010)

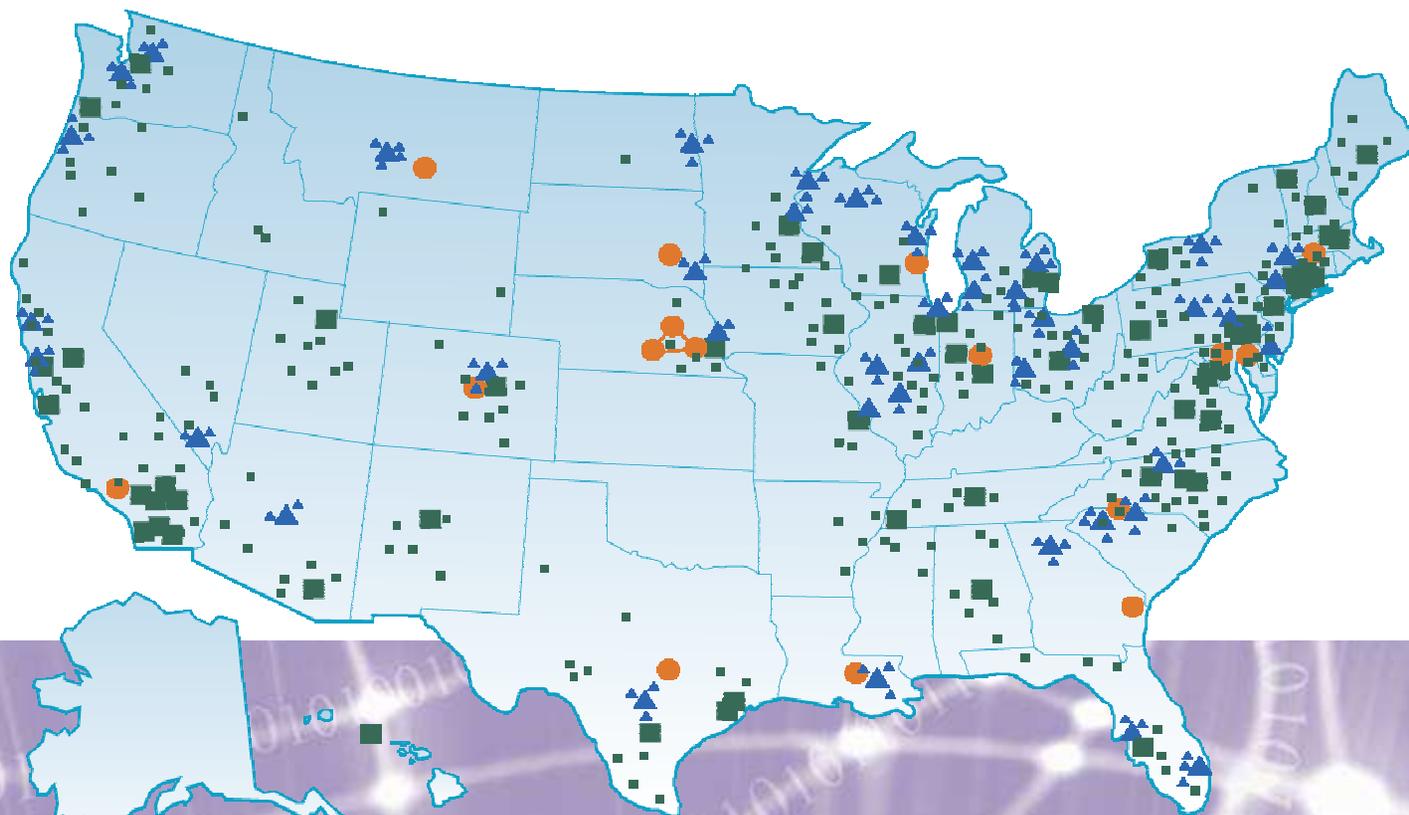


- Community
 - 2,300+ participants
 - 700+ organizations
 - 15+ countries
 - 19 licensed Support Service Providers
 - 1,100+ attendees registered for 2010 caBIG[®] Annual Meeting
- Connectivity
 - 78+ applications
 - 149 “nodes” connected to National Grid via caGrid
- Content
 - 2.17 million biospecimens available through caGrid
 - 4.76 million images stored in National Biomedical Imaging Archive
 - 39,952 microarray experiments available on caGrid

Linking The Cancer Community

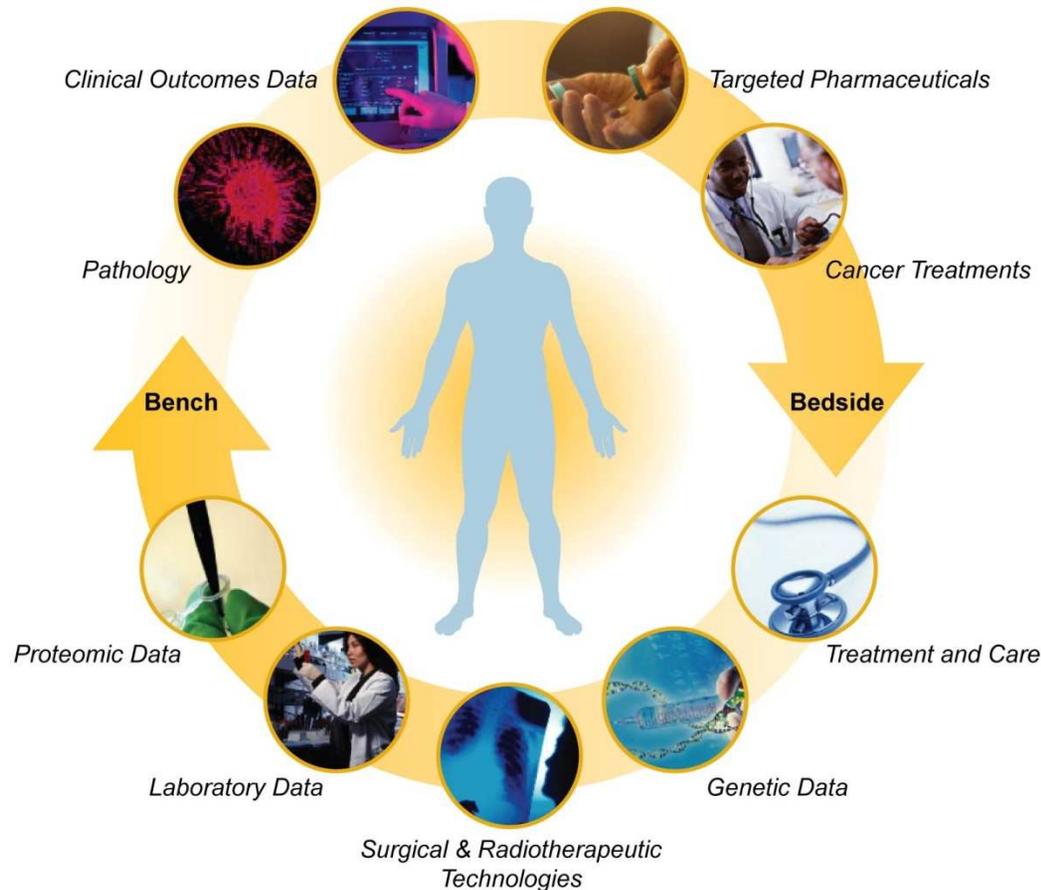


■ - NCI Designated Cancer Centers ● - NCI Community Cancer Center Program Sites ▲ - Community Oncology Practices



The Goal: 21st Century Biomedicine

Personalized, Predictive, Preemptive, Participatory

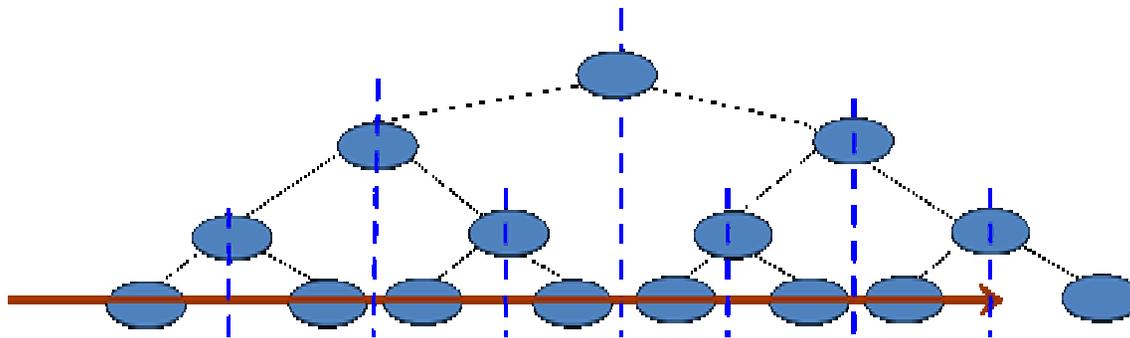


- Unifies discovery, clinical research, and clinical care (bench → bedside → bench) into a seamless continuum
- Accelerates the time from discovery to patient benefit
- Improves clinical outcomes
- Embraces the *global* nature of disease and the unique international insights in addressing its challenge
- Utilizes Health 2.0 technologies
- Engages all stakeholders
- Enables a *Rapid-Learning Health System*,

Describing Complex Systems

- More stakeholders →
- More agendas →
- More complexity →

Multiple vertical levels of organization and horizontal processes which – in order to produce the system’s products-of-value for its stakeholders/user – cross vertical (interface) boundaries. (adapted from Ivar Jacobson)



Experience in multiple domains has repeatedly shown that complexity is most effectively managed through layering and use of standards at the "interface boundaries."

sSOA AT NCI: *THE THEORY*

CBIIT "BIG 3" Priorities for 2010 - 2011



1. Enterprise-Level Interoperability

- Our applications, data, and analytic services need the capacity to connect and meaningfully exchange information and coordinate behavior.

CBIIT "BIG 3" Priorities for 2010 - 2011



2. Enterprise Security

- Develop and implement a clear, integrated and appropriately granulated enterprise strategy and associated operational tactics that unify and integrate our approach to security.

CBIIT "BIG 3" Priorities for 2010 - 2011



3. Deployment Support

- Comprehensively support the requirements of the 21st-century healthcare community, including
 - On-going stakeholder identification and management; and
 - Development of technology-independent specifications and reference implementations of software that provides clear benefit to that community.

Next-Generation caBIG®: *Overarching/Core Requirements*



- **Lower the barrier-to-entry for participation in caBIG®**
 - *caBIG® 1.x is too heavily front-loaded.*
- **Provide a “linear value proposition” to all stakeholders**
 - *Easy things should be easy to do.*
- **Support legacy data and functionality**
 - Next-generation caBIG® is evolution, not revolution
- **Leverage caBIG® 1.x Lessons Learned**
- **Leverage technology and semantic progress in the larger scientific and commercial communities**

Lower Barrier-to-Entry



- **Participation at multiple levels of organizational maturity**
 - Stakeholder evolution → Increasing involvement
- **Movement from *monolithic “all or nothing” single-purpose applications* to component-based, context neutral services**
 - Compatible with heterogeneous caBIG[®] community contexts
- **Enable participation via**
 - *Adoption* of NCI-developed code
 - *Adaption* of NCI-developed code
 - User Interface
 - Static semantics (local dictionaries, etc.)
 - Workflow composition (process vs capability)
 - *Development* based on NCI-defined specifications

Linear Value Proposition



- “Just enough security”
- “Just enough semantics”
- **Support for multiple programming models (developers)**
 - REST
 - SOAP
- **Layered governance**
 - Laptop
 - Enterprise
 - Community
- **Use cases remain essentially the same regardless of the governance layer**

Legacy Support (data and functionality)



- **Two basic options**
 - Facades
 - Migration and sun-setting
- **Facades**
 - Essentially invisible to the user
 - Can, however, prevent true evolution of an architecture
 - Decisions evaluated on a case-by-case basis
- **Migration and sunseting**
 - Often the “better” architecture solution
 - Data migration tools *must be provided* to prevent loss
 - Tools are often less-than-perfect
 - Data is often less-than-precise
 - Sun-setting dates must be realistic

SOA Strategic Goals

(CBIIT "semantics" frame these Goals)



- ***Intrinsic Interoperability***
 - Interoperability vs Integration
- ***Increased Federation***
 - Common endpoint and local governance
- ***Increased business/technology alignment***
 - Linear "degree of difficulty" for change
- ***Increased vendor neutrality options***
 - Specifications at a logical level (SAIF ECCF)

-- Thomas Erl, "Principles of Service Design"

sSOA AT NCI: *THE REALITY*

SOA Challenges **(SDP == SOA Design Principles)**



- 1. SDP → increased design complexity**
- 2. SDP → need for Design Standards (informational, behavioral)**
- 3. SDP → Change in software engineering process, e.g. "counter-agile" (contract-first) design/delivery**
- 4. SDP → new organizational roles, relationships, responsibilities**
- 5. SDP → governance requirements**

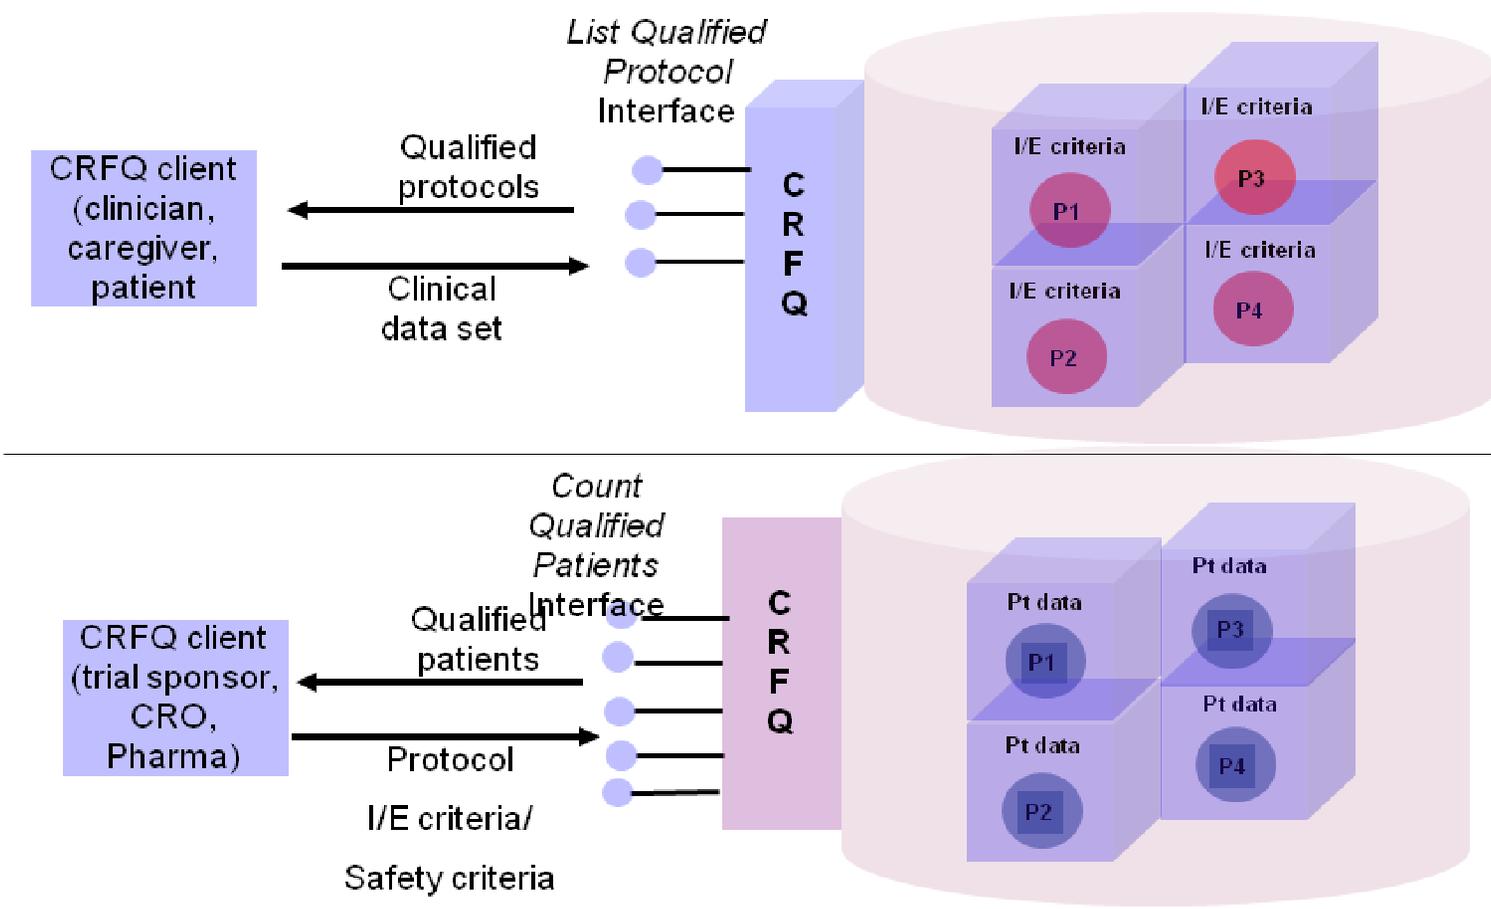


1. Increased Design Complexity

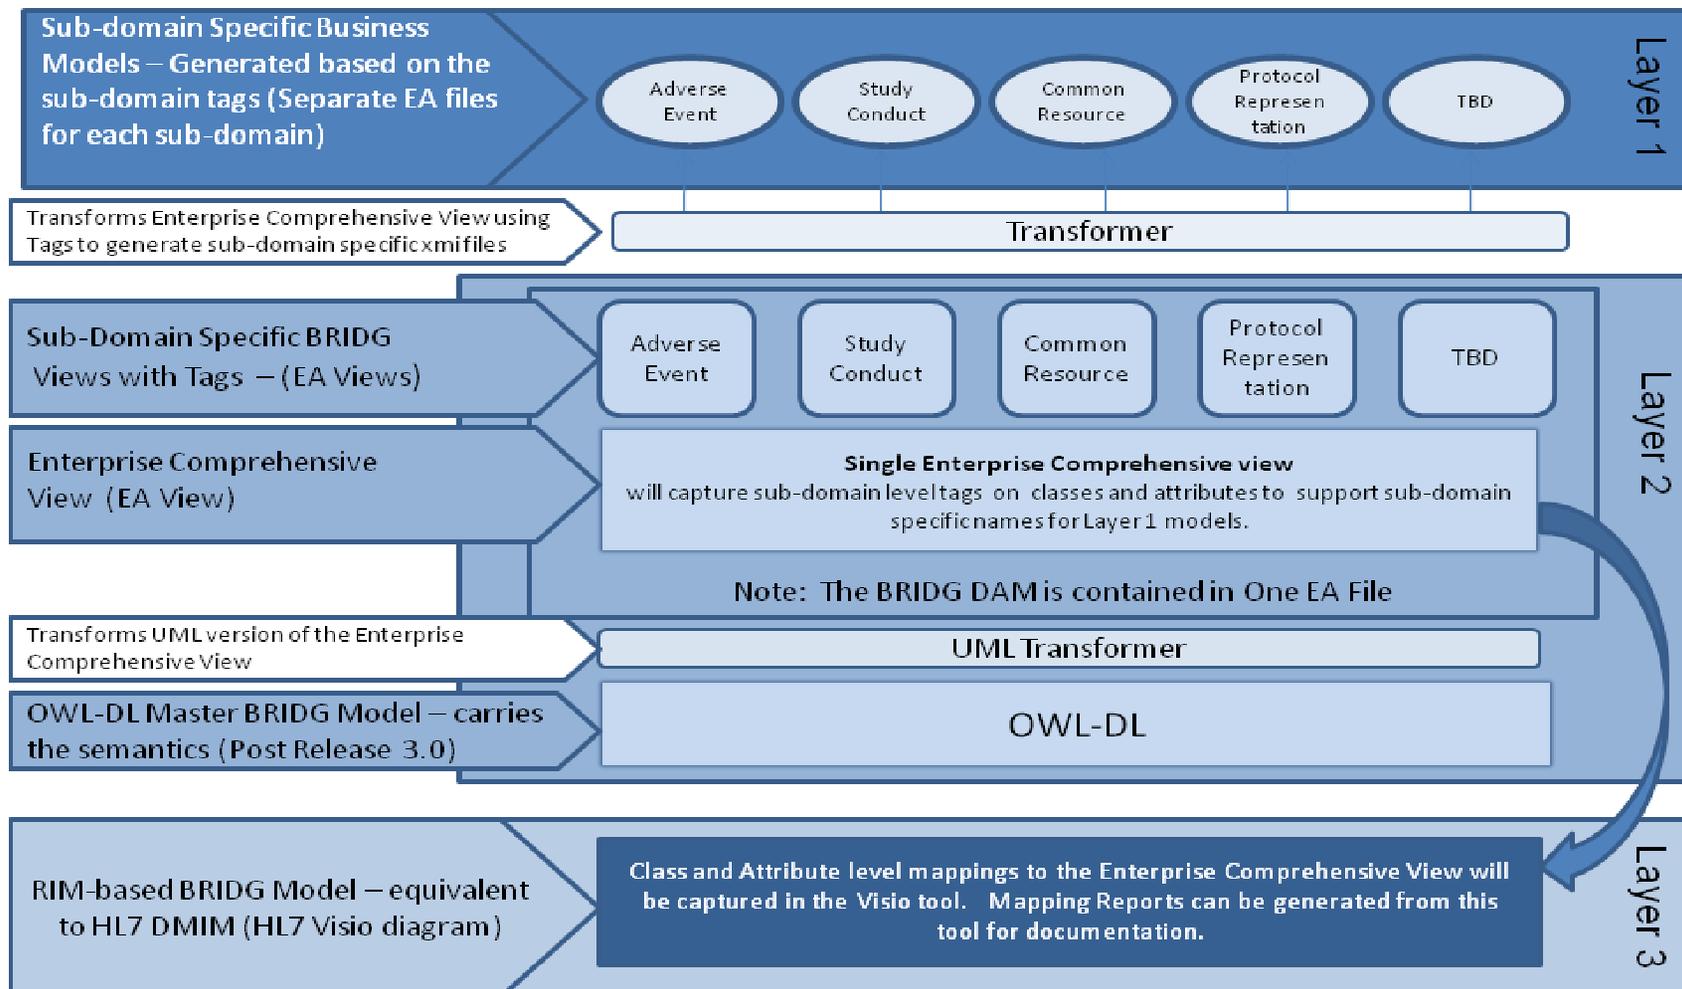


- **Clear separation of concerns**
 - Informational (static) vs Behavioral (dynamic) semantics
 - Agnostic vs non-agnostic “functional context”
- **Layering of complexity in a computable-semantic-interoperability context**
- **Coordination across teams**
- **Change Management and Education**
- **Governance!**

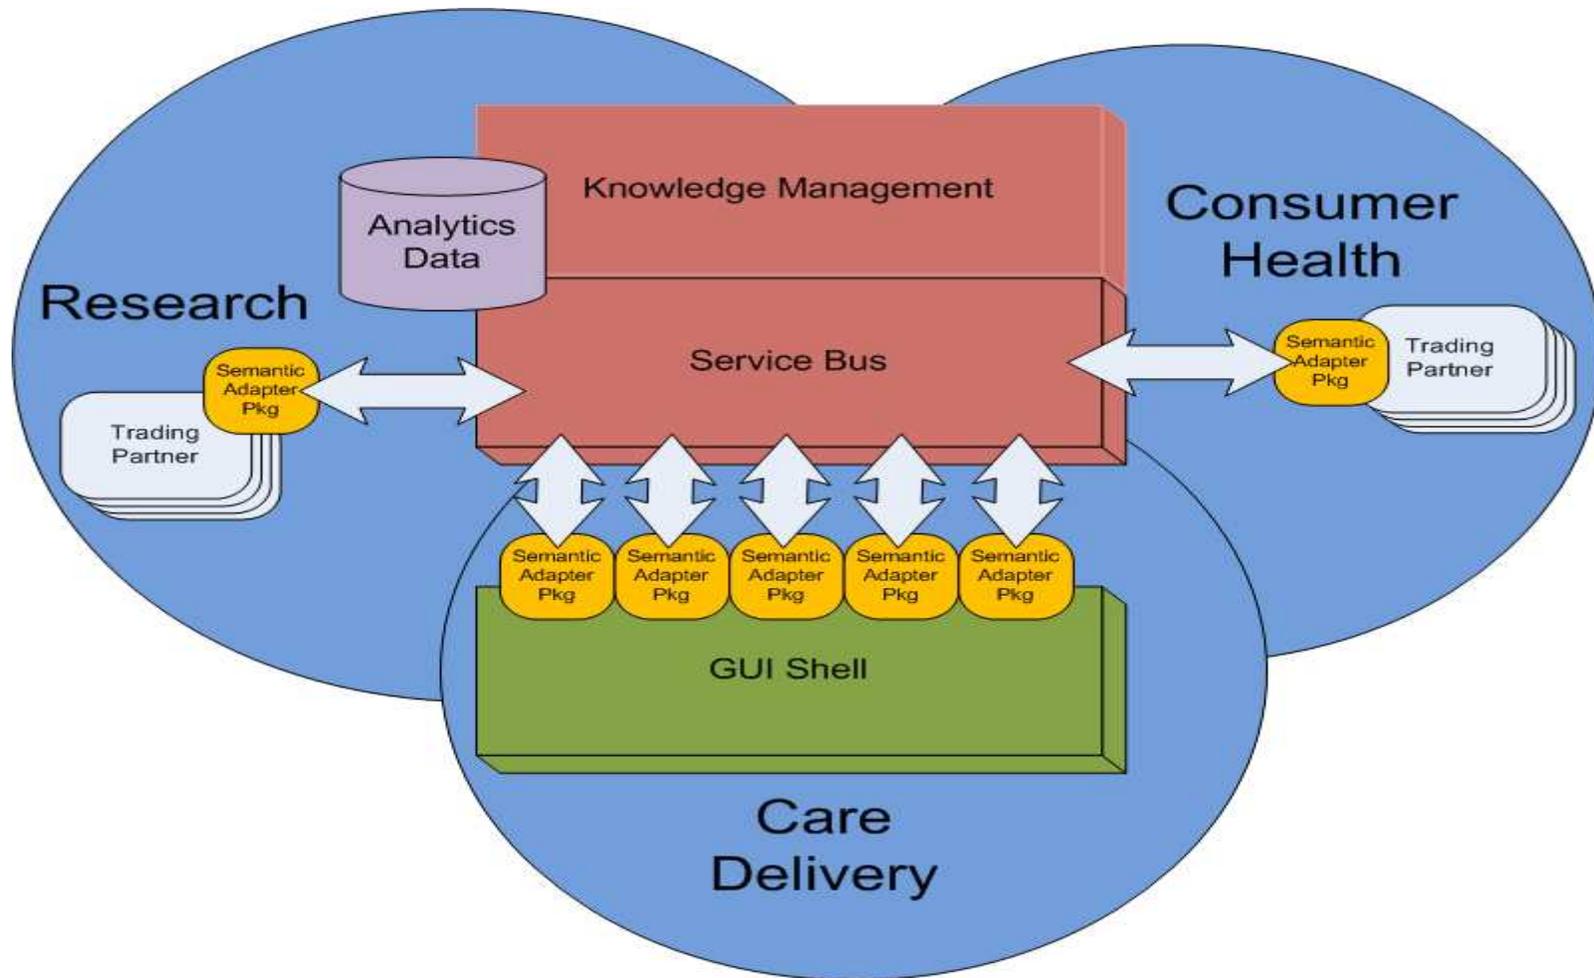
Collaboration Services: Clinical Research Filtered Query Service



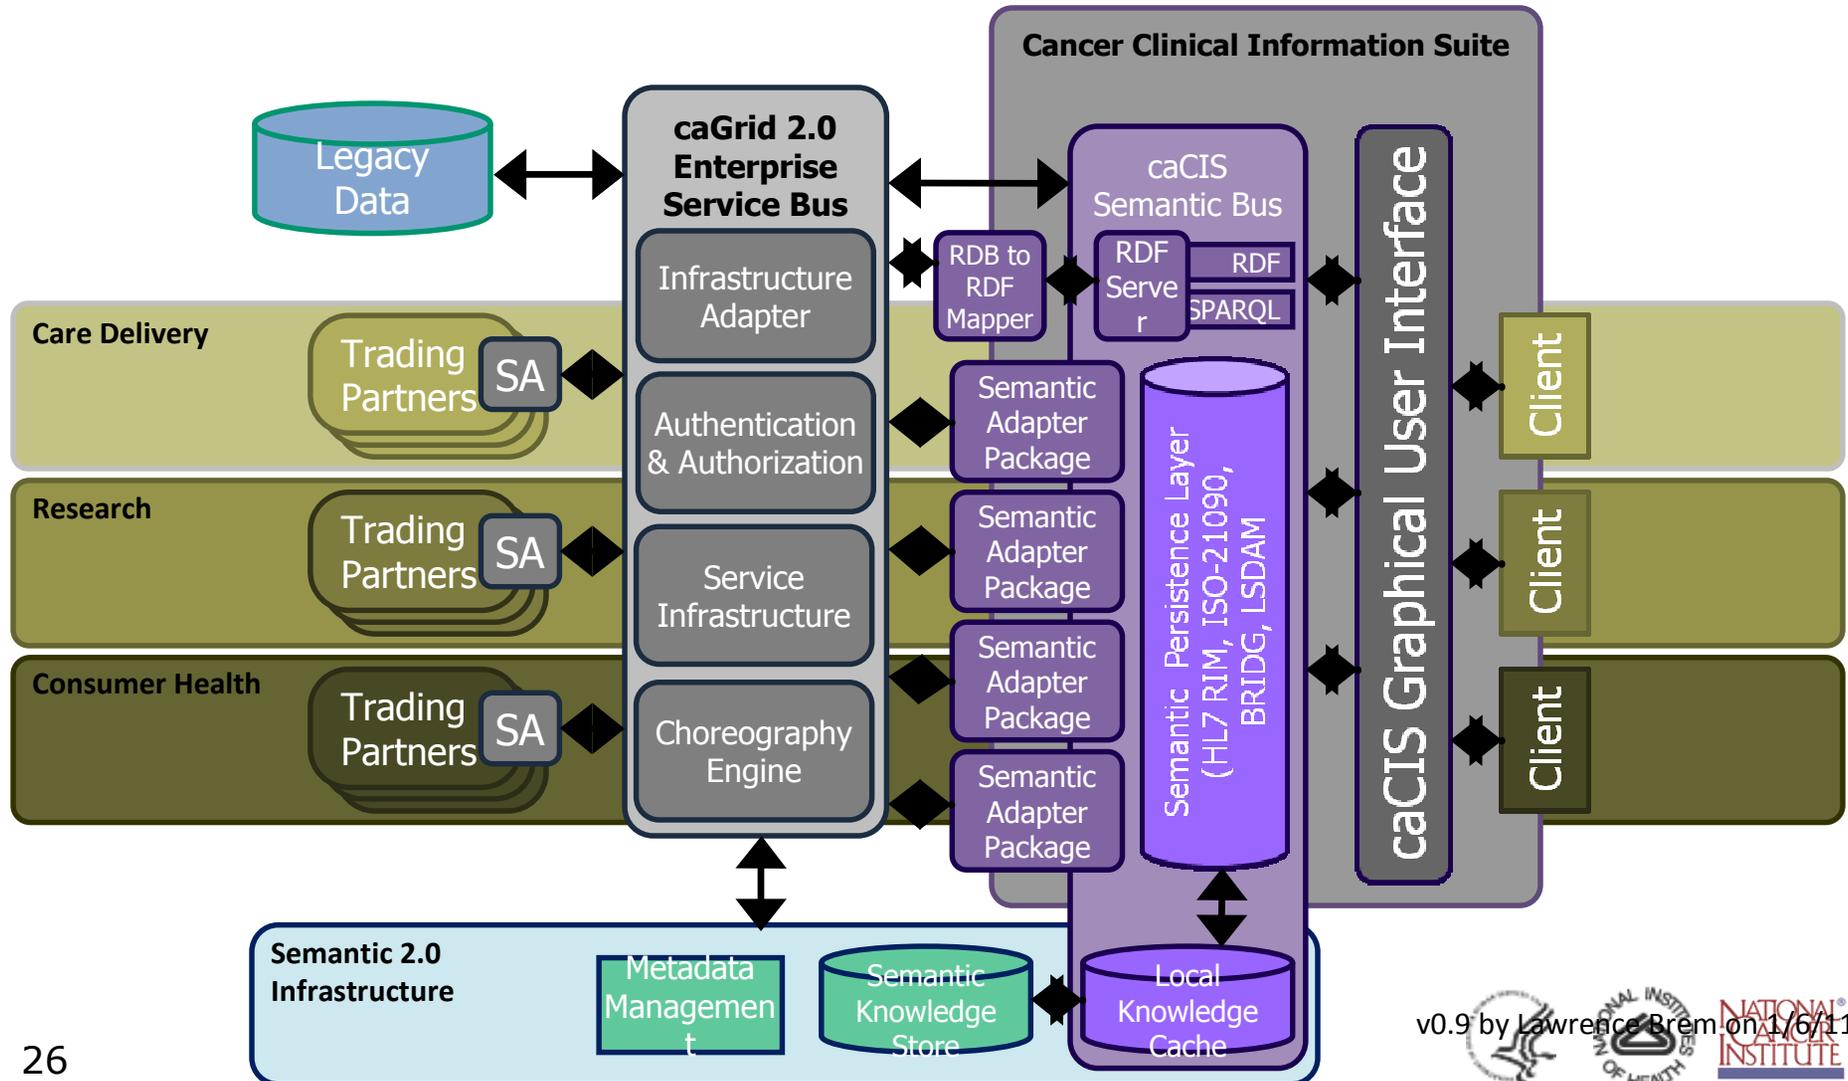
The BRIDG Model, V3.0



caCIS: Next-generation integration of EHR functionality in the ambulatory oncology context



caCIS (conceptual) Architecture



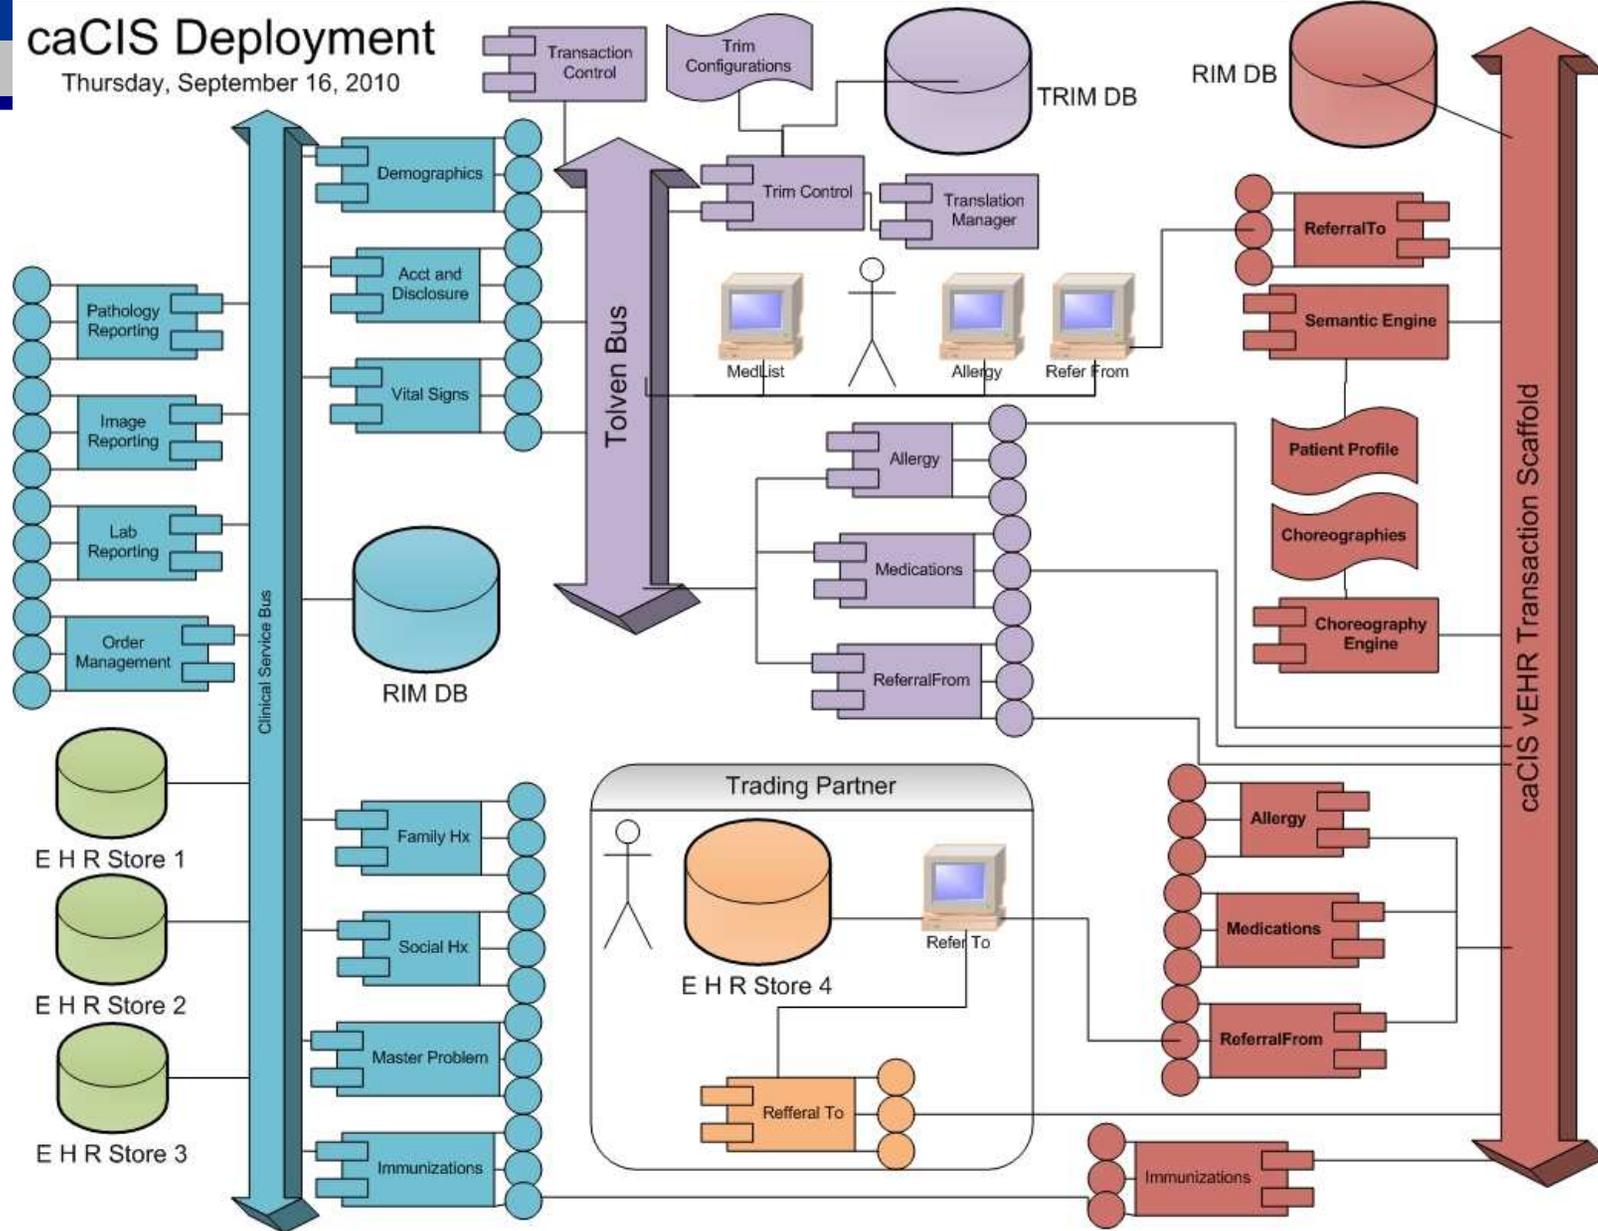
caCIS Deployment: Phase 1



Informatics
ology

caCIS Deployment

Thursday, September 16, 2010



2. Need for Design Standards/Patterns/Best Practices



- **Reference Information Model**
 - CDA
 - 21090 Data Types

- **CTS2 Terminology Services**

- **SAIF Behavioral Framework**

- **Numerous SOA Design Patterns**
 - *Canonical* Patterns (e.g. schema, versioning)
 - *Centralization* Patterns (e.g. schema, resources, rules, policies)
 - Legacy Wrapper, Façade, etc.

- **Governance!**

NES Taxonomy



Reference:		http://www.soamethodology.com/
Service Type	Service Sub-Type	Definition
Task Context-Sensitive (Non-Agnostic) Business Centric		A business service with a functional boundary directly associated with a specific parent business task or process is based on the task service model. This type of service tends to have less reuse potential and is generally positioned as the controller of a composition responsible for composing other, more process-agnostic services. Each capability of a task service essentially encapsulates business process logic in that it carries out a sequence of steps to complete a specific task.
	Orchestrated Task	Services with a functional context defined by a parent business process or task can be developed as standalone Web services or components – or – they may represent a business process definition hosted within an orchestration platform. In the latter case, the design characteristics of the service are somewhat distinct due to the specific nature of the underlying technology. In this case, it may be preferable to qualify the service model label accordingly. This type of service is referred to as the orchestrated task service .
Entity Context-Insensitive (Agnostic) Business Centric		The entity service model represents a business-centric service that bases its functional boundary and context on one or more related business entities. It is considered a highly reusable service because it is agnostic to most parent business processes. As a result, a single entity service can be leveraged to automate multiple parent business processes. Several of its capabilities are reminiscent of traditional CRUD (create, read, update, delete) methods. An entity service can be composed of other entity services.
	Enterprise Entity	NCI Specific sub-categorization. This is done primarily for governance and management purpose. Enterprise entities are entities that are used by one or more domains. Hence, any modifications to such services requires an enterprise wide agreement and approval.
	Domain Entity	NCI Specific sub-categorization. This is done primarily for governance and management purpose. Domain entities are entities that are used within a specific domain (eg. CS, LS, caCIS etc.) only. Hence, any modifications to such services requires an agreement and approval within that domain. As scope and usage of these services increases beyond one domain, domain entity services become enterprise entity services.
Utility Context-Insensitive (Agnostic) Non-Business Centric		The utility service provides reusable, cross-cutting utility functionality, such as event logging, notification, and exception handling. It is ideally application agnostic in that it can consist of a series of capabilities that draw from multiple enterprise systems and resources, while making this functionality available within a very specific processing context. Utility services are also known as application services, infrastructure services, or technology services.

NESIB: Task Services



Id	Service Name	Short Description	Service Type	Current Status	Product Team
T01	Safety Reporting	The purpose of the Safety Reporting Service is to provide a standard set of interfaces to create, manage, and submit safety reports. This service will enable the searching, analysis, sharing, and data mining of these reports and also provide the ability to evaluate harmful or potentially harmful events against criteria to determine appropriate reporting requirements.	Orchestrated Task	Draft Scope	caAERS
T02	Subject Registration	The Subject Registration service provides capabilities to manage the registration of a subject on a trial, including multi-site trials. It also enables querying of registration data and will provide authoritative information on a subject's registration.	Orchestrated Task	Draft PIM	C3PR
T03	Consult Note	Provides access to a patient's History from another Organization.	Orchestrated Task	Assigned	caCIS
T04	Study Outcomes	Study Outcomes establishes a set of reporting patterns for collecting data in CDISC SDTM format and validating it in order to report on a particular study's outcomes.	Orchestrated Task	Not Assigned	
T05	Protocol Life Cycle Management	This service will capture protocol state modifications for a study protocol by collecting information about various activities during the lifecycle of a study protocol and will support key reporting requirements. Examples of these activities are accrual, compliance of subjects to the schedule, adverse events etc.	Orchestrated Task	Not Assigned	
T06	Recruitment	This service will enable investigators to identify potential subjects by applying high-level demographic and medical history inclusion/exclusion criteria against a repository. It will also enable patients and healthcare providers to identify trials for which they may be candidates.	Orchestrated Task	Not Assigned	
T07	IRB Management	This service will support and track the interactions between participating sites, their Institutional Review Boards and the Study Coordinating Center. This service have capabilities to track and record the status of various artifacts such as study protocol, LOI, amendments etc. as they go through IRB approvals.	Orchestrated Task	Not Assigned	
T08	Image Management	Image Management service includes the ability to manage image, including intersections with order reporting.	Orchestrated Task	Assigned	caCIS
T09	Image Study Management	Imaging Study Management service includes the ability to order, track, and manage imaging orders, including intersections with order reporting.	Orchestrated Task	Assigned	caCIS
T10	Trial Registration	The Trial Registration service provides investigators with the ability to submit descriptive information about a clinical trial. It also provides the ability to search registered trials based on various criteria.	Task	Implemented	CTRP

NESIB: Task Services (2)



Id	Service Name	Short Description	Service Type	Current Status	Product Team
T11	Randomization	The Randomization service provides the capability to randomize a subject to a trial arm according to either one of a set of predefined standard randomization algorithms, or an external customized randomization algorithm.	Task	Assigned	C3PR
T12	Accrual	This service will support publishing/reporting requirements for both single-site and multi-site studies. This service will hold the data that is needed for reporting to various agencies such as CTEP/CDUS. Cancer Centers will submit the necessary data which will be used to generate the accrual report. This service can also be consumed by Protocol Lifecycle Tracking.	Task	Assigned	CTRP
T13	Order Request Management	Order Request Management provides capability to manage requests for clinical services (such as Labs and Images). Part of the Request / Fulfillment Pattern.	Task	Draft CIM	caCIS
T14	Order Fulfillment Management	Order Fulfillment Management provides access to fulfillment of clinical services (such as Labs and Images). Part of the Request / Fulfillment Pattern.	Task	Draft CIM	caCIS
T15	Patient Immunization	Immunization service establishes a central service of record for the immunizations tied to patients.	Task	Assigned	caCIS
T16	Decision Support	This service provides an algorithmically-neutral way to manage decision support services for patients.	Task	Assigned	caCIS
T17	Eligibility	Eligibility service provides the capabilities associated with establishing eligibility for a patient to register on a Study, based on inclusion/exclusion criteria specified in the Study Protocol.	Task	Assigned / On-Hold	ESST
T18	Laboratory Management	Lab Management service includes the ability to order, track, and manage laboratory orders from a clinical perspective, including intersections with specimen management and pathology reporting.	Task	Assigned	caCIS
T19	Pharmacy	Pharmacy Management service includes the ability to order, track, and manage prescriptions from a clinical perspective.	Task	Assigned	caCIS
T20	Scheduling	The Scheduling service provides the capabilities associated with scheduling a particular appointment for a particular encounter.	Task	Assigned	caCIS

NESIB: Task Services (3)



Id	Service Name	Short Description	Service Type	Current Status	Product Team
T21	Lab Result Grading Service	The Lab Result Grading service will provide the capability to compare lab values against values specified in the study protocol and assign a grade to each value.	Task	Not Assigned	
T22	Trial Participating Site Management	The Trial Participating Site Management Service will be used to update trial records in the Clinical Trials Reporting Program (CTRP) system with participating sites information for any given trial. It will support creating and updating participating sites data for multi-site trials that are registered with CTRP.	Task	Assigned	CTRP
T23	Experiment Design	The Experiment Design service will provide capabilities which support the creation of an experimental study, such as establishing the description, design type and independent variables. This service will likely have a dependency with the Standard Operating Procedure service.	Task	Not Assigned	

NESIB: Entity Services



Id	Service Name	Short Description	Service Type	Current Status	Product Team
E01	Patient Outcomes	The Patient Outcomes service provides clinicians or administrators at clinicians' offices with the ability to submit outcome data for the cancer patients. It also provides them with the ability to query the outcomes data they submitted.	Domain Entity	Implemented	caCIS
E02	Discharge Summary	Provides access to a patient's Discharge Notes for a given encounter.	Domain Entity	Assigned	caCIS
E03	Referral	Referral is intended to provide core capabilities needed by health care organizations to facilitate the business process of patient referrals and consultations among providers.	Domain Entity	Assigned	caCIS
E04	Procedure History	This service provides access to a patient's History of Procedures from another Organization.	Domain Entity	Assigned	caCIS
E05	Treatment Plan Management	The Treatment Plan service provides the ability to create and update Treatment Plan templates, generate and view prospective calendars of patient activities, track activities as they occur, and manage patient calendars as they change during a plan.	Domain Entity	Assigned	caCIS
E06	Clinical Statements	Clinical Statements provides an semi- or non-structured note capturing resource to coordinate patient care.	Domain Entity	Assigned	caCIS
E07	Specimen ID Management	This service will manage relationships between specimens and their related entities across multiple sites/applications to facilitate interoperability.	Domain Entity	Assigned	caTissue
E08	Clinical Data Management	The Clinical Data Management service enables the collection of data using electronic CRFs in to Clinical Data Management Systems.	Domain Entity	Not Assigned	
E09	Adverse Event Management	The Adverse Event Management service facilitates the creation and management of an adverse event. Capabilities include associating the event with other data sources such as laboratory results and enabling review by study personnel.	Domain Entity	Draft Scope	caAERS
E10	Credentialing	The Credentialing Service will provide capabilities to submit and manage healthcare professional and/or organizational credential information and documentation.	Domain Entity	Assigned	FIREBIRD

NESIB: Entity Services (2)



Id	Service Name	Short Description	Service Type	Current Status	Product Team
E11	Molecular Annotations	There are a variety of molecular annotations that provide information on the biological functions and classifications of the gene, genetic variants associated with the gene, the physical location of the gene, diseases and agents associated with a gene, and other annotations of interest. The Molecular Annotation Service provides a set of interfaces for retrieving such molecular annotation information.	Domain Entity	Conformance Testing	caBIO
E12	Protocol Abstraction	The Protocol Abstraction Service establishes a central service of record for non-annotated, "flat" information about Protocols and their relation to trials.	Domain Entity	Implemented	CTRP
E13	CRF Template Management	The CRF Template Management service manages the life cycle of CRF templates by providing the capability to create, aggregate, publish and query CRF templates.	Domain Entity	Draft PIM	ESSRT
E14	Assay Data Management	Many of the attributes and workflows related to the description and management of biological assay data are common across a wide range of assays in cancer research. This service will provide capabilities to manage data from diverse assay types.	Domain Entity	Not Assigned	
E15	Equipment	The Equipment service will act as a Master Inventory of equipment that will store equipment information along with settings/calibration information. This service can store information about a variety of equipment types such as lab and imaging equipment.	Domain Entity	Not Assigned	
E16	SDTM	The SDTM service enables the creation and management of tabular data for a clinical trials submission.	Domain Entity	Not Assigned	
E17	History and Physical	The History and Physical service provides access to patient's History and Physical document. This service will facilitate requests for and provision documents that capture the clinical history of the patient during a specific encounter.	Enterprise Entity	Draft CIM	caCIS
E18	Protocol Management	The Protocol Management service will support creation and management of a Study Protocol by providing capabilities to create, amend and query a computable part of the protocol.	Enterprise Entity	Assigned	CTRP
E19	Subject Management	The Subject Management service will support creation, updating, querying and deactivation of a potential study subject.	Enterprise Entity	Draft PIM	C3PR
E20	Specimen Management	Specimen Management Service will provide the core specification for specimen management from both the clinical and research perspectives regardless of the nature of specimen transactions that occur or the type of specimens involved in the transaction.	Enterprise Entity	Assigned	caTissue

NESIB: Entity Services (3)



Id	Service Name	Short Description	Service Type	Current Status	Product Team
E21	Planned Calendar	The Planned Calendar service enables the creation, amendment, state management and querying of calendar templates created for a particular study.	Enterprise Entity	Draft PIM	PSC
E22	Scheduled Calendar	The Scheduled Calendar service enables the creation and management of subject specific Scheduled Calendars. Combining a template with the Subject's start date enables the creation of a personalized calendar.	Enterprise Entity	Draft PIM	PSC
E23	Medication	Medications service establishes a central service of record for the Medications tied to patients.	Enterprise Entity	Assigned	caCIS
E24	Prescription	The Prescription Management service provides the capability to order, track and manage prescriptions for a Clinical Trial.	Enterprise Entity	Not Assigned	
E25	Lab Results Service	The Lab Results service provides a standard set of interfaces to manage the laboratory results of clinical trial participants and also enables the queries of this laboratory results data.	Enterprise Entity	Draft Scope	Lab Viewer
E26	Correlation	The Correlation service enables the creation and management of relationships (structural roles) between Persons and Organizations.	Enterprise Entity	Implemented	CTRP
E27	Organization	The Organization service enables the creation and management of entities of type "Organization". This service will be the central store for Organizations across the enterprise thereby avoiding duplication.	Enterprise Entity	Implemented	CTRP
E28	Person	The Person service enables the creation and management of entities of type "Person". This service will be the central store for demographic information for persons across the enterprise thereby avoiding duplication.	Enterprise Entity	Implemented	CTRP
E29	Master Problem List	Master Problem service establishes a central service of record for the Problems tied to patients.	Enterprise Entity	Assigned	caCIS
E30	Allergy	Allergy service establishes a central service of record for the Allergies tied to patients.	Enterprise Entity	Assigned	caCIS

NESIB: Entity Services (4)



Id	Service Name	Short Description	Service Type	Current Status	Product Team
E31	Drug	Drug service establishes a central service of record for the Drugs that may be tied to patients.	Enterprise Entity	Assigned	caCIS
E32	Agent	The Agent service establishes a central service of record for information about various agents used in trials. It includes capabilities for creating, updating, querying and deletion.	Enterprise Entity	Not Assigned	
E33	Disease	The Disease service establishes a central service of record for information about various diseases used in trials. It includes capabilities for creating, updating, querying and deletion.	Enterprise Entity	Not Assigned	
E34	Activity	The Defined Activity service facilitates the creation, updating, deletion and querying of reusable activities . For. e.g. Planned Calendar services can reuse activities to create quicker and consistent calendars.	Enterprise Entity	Draft Scope	ESSRT
E35	Standard Operating Procedure	This service will store detailed step-by-step Standard Operating Procedure templates for various experiments. This service can be used by applications such as SPOREs, TCGA, caLIMS and possibly BRD as well. The proposed service could be a generic service supporting other workspaces that need such a capability. This service should be explored to be able to support Clinical Science's Study Protocol Management and Life Science's Experiment Design Service. A standard structure needs to be determined for storage of such SOPs.	Enterprise Entity	Not Assigned	
E36	Annotation and Markup	The Annotation and Markup service provides annotation support for both Radiology and Pathology images. The service will enable the creation, management and querying of annotation markup for both individual images and series of images.	Enterprise Entity	Not Assigned	

NESIB: Utility Services



Id	Service Name	Short Description	Service Type	Current Status	Product Team
U01	Notification Management	Notification Management service will provide a centralized capability to create, manage and send event-based notifications to other applications and services.	Utility	Draft CIM	caBIG® Integration Hub
U02	Messaging	Messaging service will provide capability to exchange messages with other services / applications.	Utility	Draft PSM	caBIG® Integration Hub
U03	Global Unique Identifier	This service provides capability to generate unique identifier for all enterprise level objects (e.g. Persons, Organizations, Studies) as well as their sources.	Utility	Draft Scope	Enterprise Security
U04	Audit Management	The Audit Management service provides a set of capabilities to log, de-identify, query, analyze events and messages across the enterprise, thus facilitating the auditing of the access to sensitive data.	Utility	Draft Scope	caBIG® Integration Hub, caCORE
U05	Identity Management	Identity Management service provides means of provisioning and managing user identities as well as testing an identity against a set of credentials. It also provides ability to plug into existing identity providers to authenticate user accounts thereby allowing users to use their local identities to access Enterprise Services.	Utility	Draft Scope	Enterprise Security
U06	Credential Management	This service provides users a free and open source mechanism to obtain X.509 certificates. It provides credentials at different Levels of Assurance as directed by OMB's E-Authentication Guidelines.	Utility	Draft Scope	Enterprise Security
U07	Trust Management	Trust Management service provides a set of interfaces to maintain a list of trusted credential providers at different Levels of Assurance as directed by OMB's E-Authentication Guidelines.	Utility	Draft Scope	Enterprise Security
U08	Credential Delegation	This service facilitates delegation of user's authority by allowing the delgator to define a delgation policy dictating which applications/services can use the delegated user identity.	Utility	Draft Scope	Enterprise Security
U09	Role Management	This service allows provisioning and management of business roles. It provides capabilities to request, provision, approve and retrieve role assignments.	Utility	Draft Scope	Enterprise Security

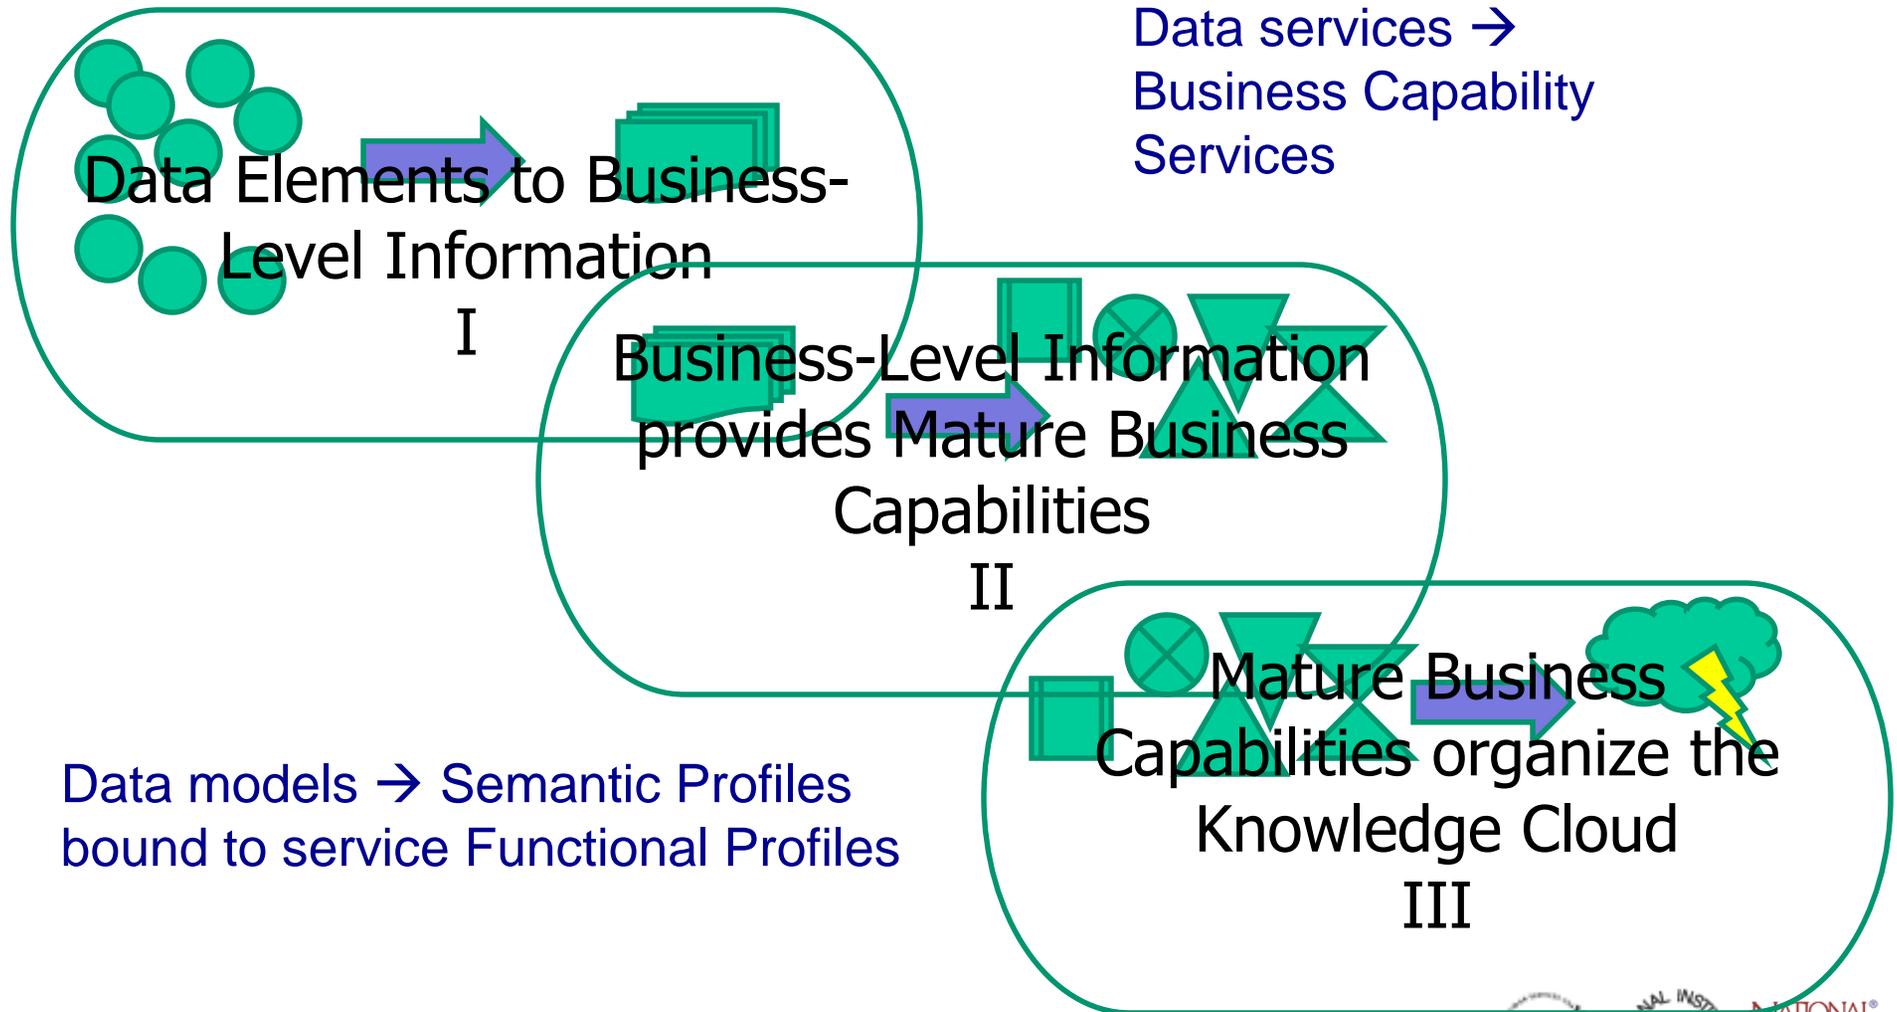
NESIB: Utility Services (2)

Id	Service Name	Short Description	Service Type	Current Status	Product Team
U10	Authorization Management	This service allows provisioning, publishing and accessing authorization policies at user level and role level. It supports different levels of granularity of the resources to be protected.	Utility	Draft Scope	Enterprise Security
U11	Data Sharing Policy	A Data Sharing policy typically lists all the guidelines, regulations and policies which the consumer of the data needs to adhere to. The Data Sharing Policy service allows the dataset owners (applications/services) to publish their data sharing policy and the users to review and accept it.	Utility	Draft Scope	Enterprise Security
U12	De-identification	De-identification service provides a way to identify those data attributes that have identifiable information or are otherwise sensitive. It also has capability to mask such data attributes.	Utility	Draft Scope	Enterprise Security
U13	Validation	The Validation service verifies structural and semantic consistency across messages used in interoperability scenarios.	Utility	Not Assigned	
U14	Transformation	The Transformation service provides a functional end point to manage and enact mappings between syntactically disparate information types.	Utility	Not Assigned	
U15	Knowledge Management	Knowledge Management service represents a series of capabilities around the storage, versioning, and expression of the semantics supporting key capabilities. Its capabilities include repository of ECCF specifications and supporting artifacts, artifact metadata management, Support transforms and the linking of business rules to enable viewing information in multiple formats, contract management, artifact and ECCF conformance validation.	Utility	Not Assigned	
U16	Governance	Governance Service includes predefined templates, workflows, and governance policies for governing the service lifecycle as well as an approval and review process for service specifications and the ability to promote services through the stages of the service lifecycle.	Utility	Not Assigned	
U17	Enterprise Vocabulary	Enterprise Vocabulary services support the management, storage, and mapping of terminologies and value sets.	Utility	Draft PIM	LexEVS
U18	Semantic Decision Support	The Semantic Decision Support Service provides capabilities to: <ul style="list-style-type: none"> • Support both design-time and run-time contract discovery • Construct and execute semantically rich queries • Develop custom analysis reports to support decision making • Run reasoners (like OWL reasoners) and view results 	Utility	Not Assigned	
U19	High-throughput Data	Provides capabilities to facilitate transfer of large sized files as well as large number of files in an efficient manner from the provider to the consumer. This service should be capable of providing this facility in synchronous and asynchronous mode allowing the clients to be able to download the files during off-peak hours.	Utility	Not Assigned	

NESIB: Utility Services (3)

Id	Service Name	Short Description	Service Type	Current Status	Product Team
U20	Rules Management	Provides capabilities to define and manage business rules, policies etc. across the enterprise. It also provides an engine which can execute these rules on input data and provide results of the execution.	Utility	Not Assigned	
U21	Service Instance Registry	Service instance registry is a central store which would store and publish information regarding all the services that are part of NCI Enterprise. This information can contain details such as the service address, the service owner, point of contact, various service level agreement parameters for performance, uptime, availability etc.	Utility	Not Assigned	
U22	Service Orchestration	This service will provide capability to define and invoke services in a harmony to perform a business task or process. This service will allow users to define and publish the business process in one of the standard workflow execution languages such as WS-CDL, BPEL etc. It will also provide an execution engine which can execute these	Utility	Not Assigned	
U23	Assertion	To be defined by caGrid 2.0 Roadmap	Utility	Not Assigned	
U24	Non-repudiation	To be defined by caGrid 2.0 Roadmap	Utility	Not Assigned	
U25	Privacy	To be defined by caGrid 2.0 Roadmap	Utility	Not Assigned	
U26	Computation	This common computational service is envisioned to provide a registry of algorithms. Users can upload data to this service and select the algorithms to be applied to that data. Alternatively, users can download any of the published algorithms to be applied to their data locally as well.	Utility	Not Assigned	
U27	Image Visualization	Image Visualization Service is a service which provides a google map like visualization capabilities. Various types of Images are rendered into JPEG/PNG/GIF Tiles which are sent across the wire to the local visualization client. This service can either provide the JPEGs/PNG/GIF (that are already stored) or render them from other formats such as DICOM. This service will also provide image zooming and panning features sending the next set of image tiles back to the visualization client based on user actions.	Utility	Not Assigned	
U28	Mobile Services	This is a service or a group of services enabling the use of mobile infrastructure for clinical trial management.	Utility	Not Defined	

Evolution – not Revolution



3. Changing Software Engineering Process



- **Emphasis on more “up-front” analysis**
 - **Closer coordination between analysis and architecture**
 - **Organizational choice between**
 - Top-down
 - Bottom-up
 - Middle-out
 - **Each strategy has positives and negatives in terms of**
 - Time-to-production
 - Down-stream governance (including refactoring costs, resources, etc.)

- **Adoption of a “counter-agile” (contract-first) design/delivery**

- **Governance!**



4. New Organizational Roles, Relationships, Responsibilities



- **Service Analyst**
- **Service Architect**
- **Service Developer**
- **Service Custodian**
- **Service Administrator**
- **Schema Custodian**
- **Policy Custodian**
- **Service Registry Custodian**
- **Technical Communications Specialist**
- **Enterprise Architect**
- **Enterprise Design Standards Custodian (and Auditor)**
- **SOA Quality Assurance Specialist**
- **SOA Governance Specialist**
- **SOA Security Specialist**



4. New Organizational Roles, Relationships, Responsibilities (2)



■ At NCI:

- not all disciplines/skill sets were currently part of the organization →
- map governance to SDLC risk mitigation →
- provides support for recruiting/hiring actions and budgets

5. Governance Requirements



- **Adoption of Thomas Erl's Governance Framework.**
 - **"SOA Governance," Prentice Hall 2011**
- **NCI has mapped "Precepts/People/Processes/Metrics" to high-risk points on the NCI's service development life cycle (SDLC).**
 - **Risk mitigation through governance**
- **SOA will not succeed without governance.**

Governance Framework: *Key Components*

(Thomas Erl et al "SOA Governance" 2011 Prentice Hall)



- **Governance provides a systematic way for organizations to make decisions. A governance system is implemented using:**
 - **Precepts**
 - » Define the rules that govern decision-making
 - » Codified using principles, policies, standards and guidelines
 - **People**
 - » Make decisions in accordance to and within the constraints stipulated by Precepts
 - **Processes**
 - » Coordinate decision-making activities
 - » Provide the means and opportunities to control decisions, enforce policies and take corrective action
 - **Metrics**
 - » Measure compliance with Precepts
 - » Provide visibility into the progress and effectiveness of the governance system

Governance Framework (1)

(Thomas Erl et al "SOA Governance" 2011 Prentice Hall)



- **Precepts** -- *an authoritative rule of action including the roles responsible for making the decision, constraints associated with decisions, and the consequences of non-compliance with the decision.*
 - **Objectives** -- *the underlying motivation and overarching goals for establishing the Precept. In addition, the Objectives define the overarching responsibility, authority, and communications associated with a given Precept.*
 - **Policies** -- *statements which define specific rules associated with a given Precept ,and, in addition, establish decision-making constraints and consequences of non-compliance.*
 - **Standards** -- *statements which specify any mandatory content, formats, technologies, processes, actions, and/or metrics that people in roles assigned to produce artifacts associated with a given Precept must use and/or execute in order to implement one or more of the Policies associated with the Precept.*
 - **Guidelines** -- *non-mandatory recommendations , patterns ^{inspired} best practices that can be adopted in the context of the Precept.*

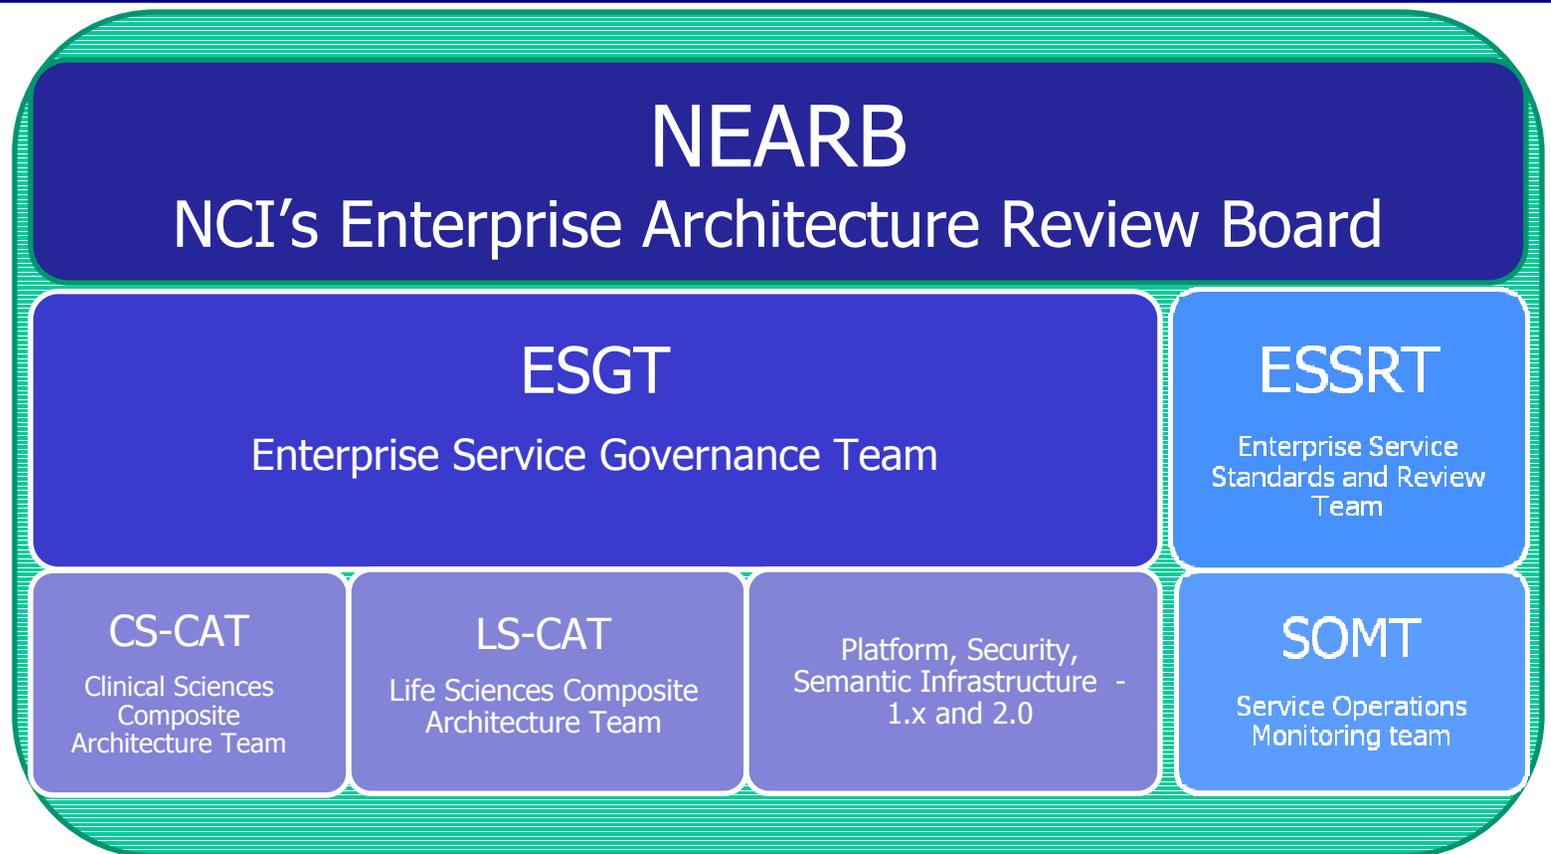
Governance Framework (2)

(Thomas Erl et al "SOA Governance" 2011 Prentice Hall)



- **Roles** -- *the various capabilities/capacities/competencies associated with people or organizations responsible for meeting the goals of a given Precept. In the CBIIT SAIF IG, roles – whether they are people- or organization/group-based are specified in a RACI-style chart associated with each Precept .*
- **Processes** -- *an organized representation of a series of activities that provide a means by which to control decisions, enforce policies, and take corrective action in support of the governance system. Although a number of management processes may be required to satisfy a particular governance Precept, these processes are not considered part of the governance system per se.*
- **Metrics** -- *quantitative, measurable/assessable criteria which a governing body can use to assess compliance/conformance with a given Precept.*

NCI sSOA: Enterprise Service Governance



Governance: Core Processes

(reference: *The Open Group*)



– **Compliance / Conformance**

- Performed at various points during the SOA Governance lifecycle
- Provide a mechanism for review/approval against the established criteria

– **Dispensation (Appeals and Exceptions)**

- Allows to appeal noncompliance to established processes, standards, policies as defined within the governance framework
- May result in a granted exception. Rejection of appeal may be followed by an escalation.

– **Communication**

- Educates, communicates and supports SOA Governance regimen
- Ensure Governance is understood and acknowledged within governed processes

– **Vitality**

- Maintains the applicability of the governance model
- Requires the governance model to stay current by refining processes to ensure ongoing usage and relevance



sSOA AT NCI: *THE SUMMARY*

sSOA at NCI: *The Summary*



- **SOA is a mature architecture paradigm for distributed computing**
 - The core design principles and patterns that have been collected over the past 5+ years provide a comprehensive, practical, and tractable framework for solving the multi-dimensional, complex problem of enabling Working Interoperability across the Translational Medicine Continuum: *"From bedside to bench and back...and beyond."*
 - Organization adoption has been "piecewise" across enterprise
- **SOA can be implemented using a number of technologies**
- **In additional to architecture and technology expertise, an enterprise commitment to SOA requires...**
 - Executive commitment and understanding of both benefits and challenges of SOA
 - Adequate funding and resources to support Change Management
 - Education
 - Governance

sSOA at NCI: *The Summary (2)*



- **NCI (and others) are involved in building systems for multiple communities**
 - healthcare delivery is part of a larger ecosystem involving research, development, public health, personal involvement in care and prevention
 - a *governed* SOA approach provides a set of patterns for the ecosystem to more organically develop and deploy capabilities that build on, add value to, and do not disrupt (evolution, not revolution) previous investments.
- **Based on the experience of the past 5+ years in developing a (relatively) small integrated healthcare community – in the broadest sense of that term – we believe that SOA is the only paradigm that will enable IT to succeed in...**
 - support *today* for the healthcare ecosystem
 - support *today* for the healthcare enterprise
 - support *tomorrow* for both as evolution based on increasing knowledge and innovation moves ever forward

sSOA at NCI: *Relevant URLs*



- **NCI CBIIT Enterprise Service documentation**
 - **NESI Blueprint:** <https://wiki.nci.nih.gov/x/Q4KtAQ>
 - **NESI Processes and Standards:** <https://wiki.nci.nih.gov/x/EIA1AQ>
- **_NCI CBIIT SAIF Implementation Guide:** <https://wiki.nci.nih.gov/x/nTdyAQ>
- **NCI CBIIT Semantic Infrastructure 2.0 Roadmap:** <https://wiki.nci.nih.gov/x/IRnDAQ>
- **NCI CBIIT caGrid 2.0 Roadmap:** <https://wiki.nci.nih.gov/x/4gTDAQ>

Q & A