Intelligence Community and Department of Defense Content Discovery and Retrieval Integrated Project Team

IC/DoD Content Discovery and Retrieval Reference Architecture

v1.0

19 December 2009

REVISION/HISTORY

Doc Revision	Revised By	Revision Date	Revisions
0.1	CDR IPT	24 September 2009	1st draft for internal CDR IPT review
0.2	CDR IPT RA Subgroup	22 October 2009	Major updates across all sections
0.3	CDR IPT RA Subgroup	19 November 2009	Major updates/revisions to each section for clarification
0.4	CDR IPT RA Subgroup	16 December 2009	Updates based on 11/19 adjudicated comments. This is the pre-baseline draft.
0.5	CDR IPT	18 December 2009	Editing of document.
1.0	CDR IPT	19 December 2009	Internal Baseline

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

The Intelligence Community (IC) and Department of Defense (DoD) support an evolving and dynamic mission and business landscape, so their information systems must enable information sharing and interoperability by default. To achieve this objective, program managers, IT managers, and developers require guidance to plan, manage, model, and build interoperable information systems.

In July 2007, the DoD Chief Information Officer (CIO) and the Director of National Intelligence (DNI) CIO signed a memorandum committing the IC and DoD to a joint vision and shared oversight for realizing a common services-based environment. The key value proposition is to enable greater and more flexible information and capability sharing within and across the IC/DoD Enterprise. Under this vision, the Joint IC/DoD Content Discovery and Retrieval (CDR) Integrated Project Team (IPT) was established to develop and publish a set of architecturally driven standards and specifications to enable content discovery and retrieval of all[1] IC and DoD data assets.

This CDR Reference Architecture (RA) serves as the keystone CDR IPT guidance artifact and describes an overall architecture to support Enterprise-wide content discovery and retrieval. The goal of the CDR RA is to establish a flexible framework for the IC/DoD environment that is extensible, and scalable to support evolving mission/business requirements. Moreover, this reference architecture is intended to provide conceptual and contextual guidance to architects, engineers, and developers tasked to implement CDR specifications. The CDR RA describes key interaction patterns to guide different implementers in meeting their specific needs. In this sense, the CDR RA supplements each of the CDR services specifications to provide the foundation for understanding their interactions, external dependencies, common assumptions and constraints.

1 *The term "all" here implies that the CDR IPT will provide the technical guidance that any data provider can leverage (as identified during the creation of this document), but does not have the authority to enforce which data providers must implement the guidance.*

1 Introduction

The Intelligence Community (IC) and Department of Defense (DoD) support an evolving and dynamic mission and business landscape, so their information systems must enable information sharing and interoperability by default. To achieve this objective, program managers, IT managers, and developers require guidance to plan, manage, model, and build interoperable information systems.

In July 2007, the DoD Chief Information Officer (CIO) and the Director of National Intelligence (DNI) CIO signed a memorandum committing the IC and DoD to a joint vision and shared oversight for realizing a common services-based environment. The key value proposition is to enable greater and more flexible information and capability sharing within and across the IC/DoD Enterprise. Under this vision, the Joint IC/DoD Content Discovery and Retrieval (CDR) Integrated Project Team (IPT) was established to develop and publish a set of architecturally-driven standards and specifications to enable content discovery and retrieval of all[2] IC and DoD data assets.

1.1 Document Overview

The main sections of this Reference Architecture are organized as follows:

Section 1 – Introduction: Describes the introduction and document overview.

Section 2 – Purpose: Describes the purpose, scope, guiding principles, assumptions and constraints, and dependencies.

Section 3 – Key Priorities: Provides an overview of the requirements elicitation and analysis process as well as an overview of the key CDR features driving the definition of the needed capabilities.

Section 4 – CDR Capabilities and Components: Provides an overview of the capabilities and components needed to enable the key priorities for content discovery and retrieval. Furthermore, this section provides a capability development roadmap for CDR services.

Section 5 – Architecture Patterns: Illustrates how the CDR components can be leveraged within different business/mission contexts.

Appendix A – Acronyms and Abbreviations: Defines the abbreviations used in this document.

Appendix B – Reference Documents: Lists the documents used in preparation of this document.

Appendix C – **Search Patterns:** Describes in more detail the search patterns indentified in Section 2.

---

2 *The term "all" here implies that the CDR IPT will provide the technical guidance that any data provider can leverage (as identified during the creation of this document), but does not have the authority to enforce which data providers must implement the guidance.*

Appendix D – Content Collection Representation Realization Patterns: Provides guidance on how one might realize the CCR Component (defined in Section 3) under the auspices of each Core CDR component individually, or in concert with each other.

2 Purpose

This CDR Reference Architecture (RA) serves as the keystone CDR IPT guidance artifact and describes an overall architecture to support Enterprise-wide content discovery and retrieval. The goal of the CDR RA is to establish a flexible framework for the IC/DoD environment that is extensible, and scalable to support evolving mission/business requirements.

The following subsections describe this reference architecture's relationship to other CDR IPT architectural elements, its intended uses, and its intended audience.

2.1.1 Relationship to Other CDR IPT Architecture Elements

The CDR IPT prescribes an abstract-to-concrete model for the development of architecture elements and guidance for content discovery and retrieval. Each layer or tier of the model is intended to provide key aspects of the overall guidance to achieve the goals and objectives of the CDR IPT. The following graphic illustrates this model.

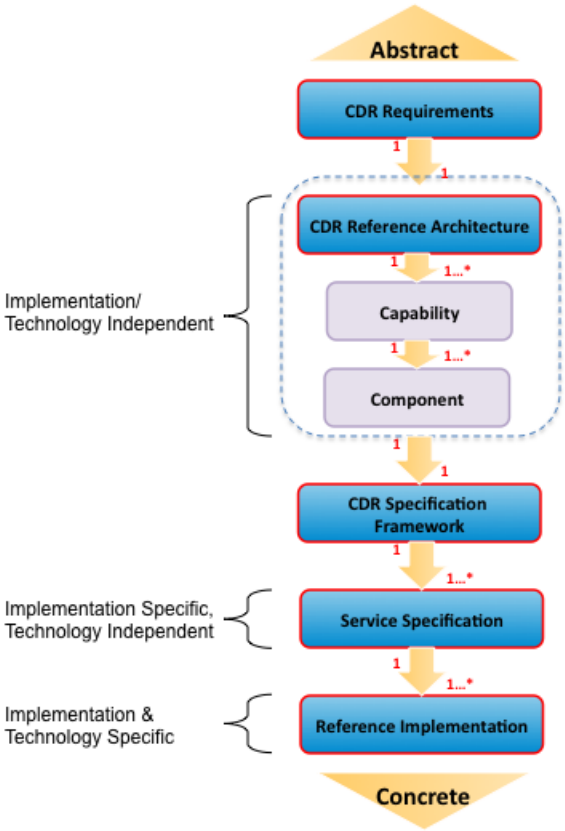

Figure 1 – CDR Architecture Model

As illustrated in Figure 1 above, the following descriptions are provided:

- **CDR Requirements** – The critical driver of the CDR RA is the collection, analysis, and management of CDR requirements from disparate stakeholders across the IC/DoD Enterprise. Section 2 of this RA describes the requirements and the articulation of key priorities for the CDR RA. The consolidated CDR requirements are maintained in the CDR Requirements Master List[3].
- **CDR Reference Architecture –** The CDR RA is the keystone artifact for the overall set of guidance artifacts. The CDR RA defines a prioritized and extensible set of capabilities and components that are realized via service specifications. Furthermore, the CDR RA describes various architecture interaction patterns and the specific technical framework to enable robust, scalable, and repeatable implementations across varying business and mission boundaries.
- **Capability** – A capability represents a logical grouping of requirements or prioritized needs and is defined within the CDR RA. Each capability can be developed independently as long as they follow the same constructs identified in the CDR RA. CDR Capabilities are described within Section 3 of this RA.
- **Component** – A component represents an architectural element that is implementation and technology independent and is defined within the CDR RA. Components identified within the CDR RA include key components to realize the capabilities but also address external dependencies (e.g., security). CDR Components are described within Section 3 of this RA.
- **CDR Specification Framework –** Describes the behavioral model for CDR Service Specifications including the description of their key properties and a decomposition of key interaction patterns.
- **Service Specifications** – A service specification is a formal description of a mission/business function to be offered based on an abstract capability or component defined within the CDR RA. A service specification includes an interface definition, an information model, a quality model, and a behavior model. A service specification is analogous to Application Programming Interface (API) documentation for software in that it provides the semantics and interfaces for a service capability. A service specification provides stakeholders with a description of the behavior and structure of the service capability. It also provides service consumers with information to integrate with (i.e., consuming) a particular service, and service providers the information to implement (i.e., expose) the service. A service specification can be implemented by Commercial Off-the-Shelf (COTS) software, Government Off-the-Shelf (GOTS) software, custom software, or a combination of custom software wrappers interfacing to COTS/GOTS software, and represents the fundamental building block of services.
- **Reference Implementation** - To validate a service specification within the context of the reference architecture, a reference implementation is a concrete realization of a service specification. This reference implementation acts as a future test harness for consumers who wish to use the service or for providers who wish to expose the service.

3 *The CDR Requirements Master List is referenced in Appendix B, item 8.*

2.1.2 Intended Use and Audience

This reference architecture is intended to provide conceptual and contextual guidance to architects, engineers, and developers tasked to implement CDR specifications. The CDR RA describes key interaction patterns to guide different implementers in meeting their specific needs. In this sense, the CDR RA supplements each of the CDR services specifications to provide the foundation for understanding their interactions, external dependencies, common assumptions and constraints.

2.2 Scope

Achieving a ubiquitous content discovery and retrieval solution presents a substantial challenge in the IC/DoD Enterprise where content exists in a large variety of structures[4], are represented in diverse[5] semantics, and are exposed through many types of technical implementations. To meet this challenge, the CDR RA describes an architecturally driven approach for guiding the IC/DoD Enterprise towards enabling content discovery and retrieval.

2.2.1 Content Discovery and Retrieval within the IC/DoD Enterprise

Organizations in the IC and DoD implement content discovery and retrieval in a variety of ways based on their individual business and data requirements. A CDR capability must support a number of different query, content, and metadata formats, depending on the context. Rather than presenting one, monolithic schema for discovery and retrieval, the CDR RA introduces a number of concepts, components, and interactions in an abstract framework that supports flexibility

2.2.2 The Big Picture – The IC/DoD Services-Based Environment

As described in the Introduction section of this reference architecture, the IC and DoD are committed to achieving a common services-based environment as the basis for information sharing and business interactions. The following graphic illustrates a conceptual framework for this services-based environment.

4 *Content resources and/or collections include structural attributes that effect their discoverability and accessibility.*

5 *In this sense, diversity also refers to the large percentage of IC/DoD information resources that lack any mature data models or adherence to any IC/DoD data standards*

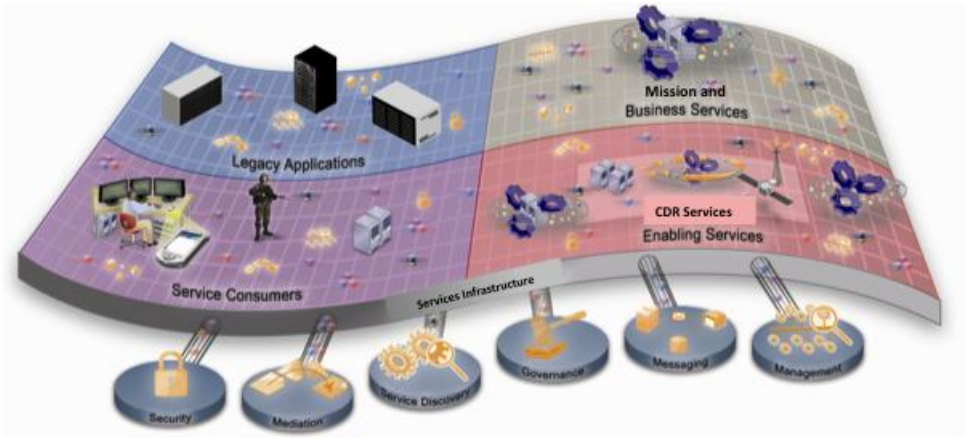

Figure 2 – The IC/DoD Services-based Environment

As illustrated in Figure 2 above, the vision within the IC/DoD services-based environment is that service-level infrastructure capabilities, such as Security and Service Discovery, are available for IC and DoD Enterprise consumption and enable network-transparent interactions across the environment. Consumers seamlessly interact with familiar domain applications to more efficiently and effectively perform particular mission and business functions. These applications interact with services from across the Enterprise, including local mission services, external mission services, and enabling services hosted by IC and DoD Enterprise.

The conceptual framework illustrates a strategic view of service consumers and exposed CDR services (as well as legacy applications exposed to the greater Enterprise) within a network-transparent environment. The underlying Services Infrastructure layer consists of core capabilities, including Security, Mediation, Service Discovery, Governance, Messaging, and Management. Both service consumers and service providers leverage the infrastructure layer in performing their business functions.

2.3 Guiding Principles

The core guiding principles of the CDR RA are as follows:

1. **Mission/business Driven:** Going forward, CDR requirements and statements of need will provide specific drivers or inputs to the CDR RA, which must be able to scale and adapt as needed. Improved information sharing requires a strategy that addresses discovery, access, context, trust, and interoperability. With a strategic eye on realizing these areas, and as constraints allow (including technology, standards, policy, culture, and resources), the CDR RA advances iteratively based on requirements collected.
2. **FEA Data Reference Model Compliant:** The Federal Enterprise Architecture (FEA) Data Reference Model (DRM) enables information sharing and reuse via the standard description and discovery of common data and the promotion of uniform data management practices. Iterations of the CDR RA primarily address the requirements of the FEA Data Sharing standardization area.
3. **Separation of Concerns:** The CDR RA prescribes technical guidance to enable an interoperable and repeatable CDR solution across an extremely heterogeneous

IC/DoD environment independent of policy and other technical/data standards. As such, this guidance is heavily reliant on Enterprise policy to ensure consistent implementation and enforcement, as well as Enterprise data standards and policies for ensuring common "views" of IC/DoD data resources[6].

4. **Open Standards:** The CDR RA prescribes an open standards-based approach. Open standards provide the basis for interoperability, extensibility, agility, vendor independence, and feasibility within and across IC and DoD mission/business domains.
5. **Iterative/Agile Realization:** The CDR RA will iteratively evolve and mature in increments to maintain a clear scope to facilitate timely and actionable guidance.
6. **Simplicity/Ease of Implementation:** Guidance prescribed within the CDR RA is intended to be both simple to use and non-intrusive. This accomplishes the overarching goal of lowering the barrier of entry for Enterprise members to interact with and to use the CDR RA services.
7. **Supplement Existing Information Discovery and Access Approaches:** The CDR RA defines a supplemental architecture to existing information discovery and access approaches, and NOT serve as a replacement. Organizations and programs may continue to develop and utilize Web portals and proprietary access mechanisms to meet their mission objectives. To facilitate information sharing for the IC/DoD Enterprise at large, the CDR RA defines the framework and constructs (e.g., components and service specifications) for achieving a common discovery and retrieval mechanism that spans all data assets within the Enterprise.

2.4 Dependencies

CDR services are loosely coupled and dependent on other IC/DoD Enterprise service efforts[7] including most notably Security, Service Discovery, and Messaging. The following bullets describe these focus areas.

- ***Security:*** The Security focus area provides a set of security-focused services to the IC and DoD for protecting access to services, data, and their interactions within the IC/DoD Enterprise. Integration of Security capabilities is advocated, both from the service discovery and the service access standpoint, to protect content providers and consumers from attack from any unknown entities. Security capabilities are responsible for authentication and authorization of consumers and consumer agents and also enable cross-domain search and retrieval. Furthermore, Security capabilities provide integrity, confidentiality, and audit services that CDR providers can leverage. CDR providers should reference the IC/DoD Security Reference Architecture[8] for guidance on utilizing the security capabilities.
- ***Service Discovery:*** The Service Discovery focus area provides the capabilities to publish and discover services (as well as other software entities) within the

---

6 *Policy/Standard conflict adjudication between the IC and DoD is not in scope for the CDR IPT*

7 *For more information on these specific Enterprise activities (as well as others), reference the IC/DoD Enterprise Services Engineering Review Board (ES ERB) collaboration workspaces on both the Defense Knowledge Online (DKO) and Intellipedia.*

8 *The Service Security Reference Architecture is referenced in Appendix B, item 4.*

IC/DoD Enterprise and make them visible to prospective consumers. Furthermore, the CDR RA technical framework is heavily dependent on leveraging Service Discovery to support context awareness of providers during federated discovery and retrieval requests. CDR providers should reference the IC/DoD Service Discovery Reference Architecture[9] for guidance on utilizing the service discovery capabilities.

- ***Messaging[10]:*** The Messaging focus area provides the capabilities for handling the addressing and re-routing functions necessary for asynchronous messaging. Furthermore, Messaging capabilities provide reliable messaging capabilities for potentially large sets of content.
- ***Service Description[11]:*** A service description provides the construct to describe the critical static and dynamic attributes about a service throughout its lifecycle, including linkages to relevant information (e.g., policies, technical artifacts, etc.). Utilization of service descriptions is recommended to enable static and dynamic discovery.

Apart from the specific dependencies on other infrastructure type services, CDR services also have some non-technical dependencies. These include:

- Policy and Governance Related Dependencies:
  - Service-level enforcement is dependent on a governance process to enforce compliance and ensure quality-of-service requirements.
  - Certification and Accreditation of CDR services must provide efficient processes for getting service implementations operational.
- COTS Tools and Vendor Dependencies
  - Across the full Service-Oriented Architecture (SOA) spectrum, the state of standards and COTS vendors vary in maturity.
- Network Dependencies
  - Performance and quality of service for CDR RA services are extremely reliant on the network infrastructure on which they are implemented.

2.5 Notational Convention

The key words "MUST," "MUST NOT," "REQUIRED," "SHALL," "SHALL NOT," "SHOULD," "SHOULD NOT," "RECOMMENDED," "MAY," and "OPTIONAL" in this specification are to be interpreted as described in the IETF RFC 2119. These keywords are thus capitalized when used to unambiguously specify requirements over protocol and application features and behavior that affect the interoperability and security of implementations. When these words are not capitalized, they are meant in their natural-language sense.

Throughout this RA, examples may be provided which cite specific IC, DoD, Joint IC/DoD, or industry standards. These examples provide a contextual understanding of a particular statement. To keep the CDR RA at an implementation and technology independent level and not subject to configuration management as standards evolve, the

9 *The Service Discovery Reference Architecture is referenced in Appendix B, item 5.*

10 *The Messaging Reference Architecture is referenced in Appendix B. item 6.*

11 *Reference Section 4.2.4*

standards cited do not imply implementation-level requirements. Detailed guidance may be included in implementation-level specifications derived from this RA.

3 Key Priorities

One of the primary objectives of the CDR RA is to codify an architecture framework that is scalable and flexible to support a variety of implementations suitable to the individual business and mission needs of explicit environments (i.e., COIs, domains, enclaves, etc.) within the IC/DoD Enterprise. As such, the critical driver of the CDR RA is the collection, analysis, and management of CDR requirements and needs from disparate stakeholders from across the IC/DoD Enterprise. This section of the CDR RA describes the overall requirements elicitation, analyses, and prioritization that drive the specific CDR services capabilities and framework.

3.1 Requirements Elicitation and Analysis

Organizations within the IC and DoD have captured and analyzed discovery and retrieval requirements that have been, or are to be used to build their own specific discovery and retrieval systems. To elicit specific requirements for the greater IC/DoD CDR solution, CDR IPT representatives from both the IC and DoD provided their organization's discovery and retrieval requirements as input to the development of the Master CDR IPT Discovery and Retrieval Requirements List[12].

The requirements provided by the Enterprise CDR IPT representatives addressed the full breadth of the discovery and retrieval topics. Although both discovery and retrieval have dependencies on processing, ingest, security and performance requirements, they were not included in the master requirements list, but instead recognized as critical dependencies. During the requirements analysis process, it was determined that a number of requirements were common across the greater IC/DoD Enterprise, while other requirements were determined to be unique to meet a particular business or mission function. Through analysis, some of the unique requirements were determined to be applicable to the Enterprise and were included in the master list. The list of requirements was normalized to provide a broader Enterprise context. From the Master List, key Content Discovery and Retrieval Features were extracted.

3.2 Content Discovery and Retrieval Key Priorities

Content Discovery and Retrieval were the two high level categories initially applied against the list requirements list. Requirements pertaining to metadata, query management, information sharing, ranking, and federated search were categorized as Content Discovery requirements. Retrieval requirements included delivery service, and mediation and delivery management requirements. For access to the specific requirements statements, reference the CDR IPT Master Requirements List. The following subsections describe the key priorities for Content Discovery and Retrieval.

3.2.1 Content Discovery

Across the IC and DoD, stakeholders have stated information systems must be searchable and discoverable. However, the ability to discover content is highly dependent on the content provider's ability to expose the content and the degree of allowing access to authorized consumers.

12 *The CDR Requirements Master List is referenced in Appendix B, item 8.*

The stakeholders desire flexibility when performing a search. The expectation is that the data will be searchable, available and accessible. The implied requirement is to ensure the content provider exposes the data in a manner applicable to its structure and context. Stakeholders stated multiple search filter patterns should be supported and the search parameters should be grouped and combined via logical operators[13].

Service consumers expressed the need for a capability to construct a query that will include instructions for when, where and how (e.g. using what formats and data types) the results are to be delivered. Consumers have also expressed the need to build, store, edit, and share their query with other consumers.

Federating a query across multiple search engines and accessing data stores with multiple data types ranks high with service consumers. Generally describing the requirement, federated search capability should inform the consumer of which data sources are available, which attributes may be searched, and the options the consumer may use to formulate the query. A key priority for federated search includes the ability for consumers to request the delivery of results as they are available (e.g., streaming) for immediate viewing and if desired, further processing. Several other key priorities include an extensible sorting/ranking mechanism upon delivery of results, as well as the ability for consumers to be able to distinguish between new results from those that have been previously reviewed. Additionally, the architecture must enable the consumer to re-sort the results upon receipt. Other desired features include: 1) results are to be de-coupled from presentation mechanisms; 2) results from multiple sources are to be aggregated upon delivery; 3) results are to be de-duplicated before delivery.

3.2.2 Retrieval

Many retrieval requirements contained within the CDR IPT Master Requirements List address stakeholders, ranging from those who are well connected to those who are bandwidth limited to those who are periodically disconnected. As such, the CDR architecture must account for this range of environmental conditions. The content provider should have available both synchronous and asynchronous delivery options, including support for such "use later" alternatives as FTP and caching results.

13 *Appendix C contains information regarding several key search patterns identified in the requirements.*

4 CDR Capabilities and Components

As described in Section 2, the current CDR RA defines two distinct, but inter-related capabilities:

- Content Discovery
- Retrieval

4.1 Capability Descriptions

In addition to the CDR component integration patterns illustrated in Section 4, the following sections provide an in-depth analysis of CDR capabilities.

4.1.1 Content Discovery

Content Discovery provides consumers with the ability to discover the existence of potentially relevant content resources and content collections from across the IC/ DoD Enterprise. Content Discovery also enables consumers to discover valuable metadata about content resources to support making them visible, accessible, understandable, and trusted.

To address the full range of requirements for Content Discovery, the CDR RA describes several critical assumptions to ensure the appropriate CDR components are defined. These assumptions include:

- The IC/DoD Enterprise employs multiple, independently developed content collections[14] that manage a large variety of content resource types. Across these content collections, common characteristics may have different implementation solutions and include:
  - the interface or interfaces for accepting queries
  - the supported query syntax types
  - the content collection type (e.g., database, document repository, etc.)
  - the metadata format describing the content collection
  - the types of content resources managed by the content collection (e.g., database records, imagery files, etc.)
  - the content format/structure for the content resources
  - the subject of the content
- Content consumers may submit a general search request to ascertain what information may be available from across the Enterprise based on a given set of search criteria, or content consumers may perform much more specific/refined searches based on specific attributes of content collections and/or content resources.
- Queries are expected to be save-able, discoverable, and reusable both manually and automatically through subscription.

14 *Content collections and content resources are provider components described in Section 3.2.2*

4.1.2 Retrieval

Retrieval provides the ability for consumers to retrieve and route specific content resources. To do so, content providers are to ensure their content is accessible to the Enterprise.

Additionally a consumer may submit a query from one location but have the content delivered at a specified time and/or to a different location. For this reason, the Retrieval capability should support both synchronous and asynchronous delivery of content.

4.2 Component Descriptions

Components within the CDR RA constitute implementation and technology-independent concepts. In this sense, components may be used to define a logical set of properties, functions, or behaviors. Some components may lead to the realization of a specific service implementation, while other components may describe key characteristics or inform specific operations that are realized across many service implementations.

The CDR RA identifies four main component types including:

- Consumer Components
- Provider Components
- Core CDR Components
- Key CDR Dependency Components

The following diagram illustrates the CDR RA components.

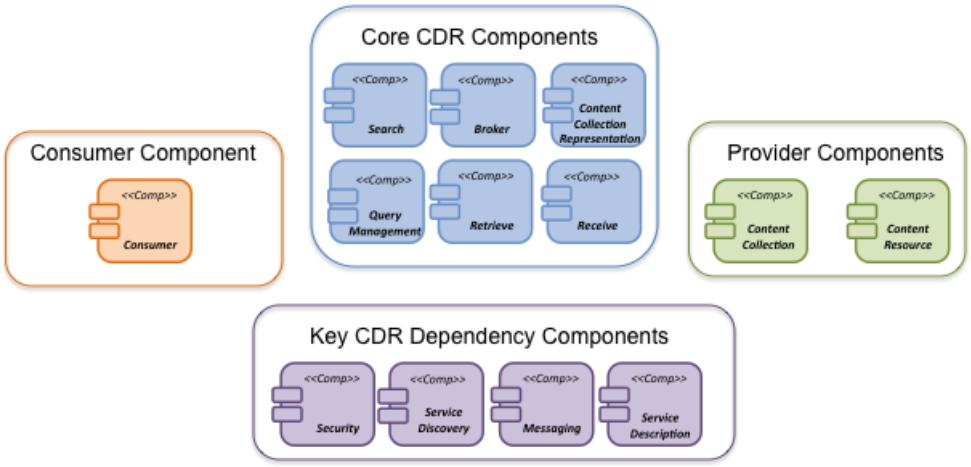

Figure 3 – CDR RA Components

The following subsections define each of these component types. Section 4 of this RA illustrates some key component interaction patterns.

4.2.1 Consumer Component

The Consumer Component represents any entity that initiates a content discovery or content retrieval interaction with any of the Core CDR Components.

It is important to note that Consumer Components may be *anticipated* consumers (i.e., those resident on IC/DoD networks and infrastructures) or *unanticipated* consumers (e.g., U.S. and coalition partners). The Security Component (one of the Key Dependency Components described in Section 3.2.4.1) provides the necessary mechanism(s) to ensure only authorized consumers are able to perform the requested function(s).

4.2.2 Provider Components

Provider Components describe the applicable constructs related to the providers of content to support content discovery and retrieval use cases. This component grouping includes a Content Collection Component and a Content Resource Component. Each is described below.

4.2.2.1 Content Collection Component

The Content Collection Component represents a data source that holds one or more content resources. In this sense, a Content Collection Component may be a database, a document repository, an imagery repository, an authoring system, etc.

Content Collection Components are generally characterized by exposing certain information (i.e., metadata), including, but not limited to, the name of the content collection, the organization or element that serves as steward for the managed content, and a characterization of the type of content resources that are stored in the collection. For example, if a Content Collection Component was a database, the characterization may include a logical model for the database, in which the terms are expressed in COI-coordinated vocabulary and common representations.

4.2.2.2 Content Resource Component

The Content Resource Component represents an individual data resource that is being made visible and accessible. Examples of Content Resource Components include a PDF file, a database record, a satellite image, a web page, etc. Information (i.e., metadata) about these components may be used by Consumer Components to further inform the Consumer whether the resource should be accessed and what transforms may be needed to optimize its utility. For a specific instance, this can include the size of the resource, when it was created or last edited, its geospatial coverage scope, its subject coverage scope, its releaseability, and any alternate formats (e.g., MIME types) that might be available. To address conformity and broad adoption of these types of metadata-related concerns, the IC and DoD are engaged in joint activities to develop, promulgate, and enforce metadata standards[15].

4.2.3 Core CDR Components

The Core CDR Components directly address specific requirements that enable Content Discovery and Retrieval capabilities. This component grouping includes Search, Broker, Content Collection Representation, Query Management, Retrieve, and Receive Components. These components are described in the following subsections.

---

15 *For example, in this specific case relating to resource metadata, the IC and DoD have adopted the DoD Discovery Metadata Specification (DDMS) as the standard for resource information format.*

Additionally, the CDR Specification Framework document describes the CDR framework around these components.

4.2.3.1 Search Component

The Search Component serves as the primary content discovery mechanism content providers implement to expose their content collections for discovery and accessibility. This component provides a common interface and behavioral model for IC and DoD content collections, enabling content consumers the ability to discover relevant content resources from disparate collections across the IC/DoD Enterprise.

In a general sense, the Search Component provides the underlying content collection with a well-defined discovery request. In turn, it will interpret this request to discover content resources. It is important to note that a Search Component does not return the actual content resources, but rather metadata about the content resources.

4.2.3.2 Broker Component

The Broker Component serves as the primary mechanism to 1) facilitate the distribution of queries to applicable/relevant content collections (exposed as Search Components) and, 2) aggregate the returned results. To satisfy the basic "federated search" use case, a Consumer component would submit a query request to the Broker Component that would then distribute the query to the applicable Search Components. The Search Components might interpret the search request as best they can, compile a list of search results[16] and send them back to the Broker Component. Depending on the consumer preference, the Broker Component may then:

- aggregate the results from all the Search Components into a consolidated set and then deliver the set to the originating consumer, or
- stream the search results directly back to a consumer as they become available from each Search Component.

Besides the high-level functions of the Broker Component described above, the Broker Component may also perform other more detailed functions in support of federated search-related processes. This may include, for example, the ability for the consumer to ascertain details regarding which content collections were queried, including any sources that were not available (if any).

4.2.3.3 Content Collection Representation Component

The Content Collection Representation (CCR) Component serves as the primary mechanism for content collectors (i.e., providers) to expose information to describe the context, access constraints, and current inventory status of its underlying content resources to support static and dynamic discovery and accessibility of a content collection. This component is leveraged by the Broker and Consumer Components to discriminate whether the content collection contains content resources that are relevant to the consumer's query. To support a wide array of use cases, the CCR Component should

16 *Results should be encoded in documented format, e.g. the IC/DoD supported format of Atom augmented with DDMS.*

reflect both the static[17] and dynamic[18] information about the underlying content collection.

It is important to note that although the Provider components may provide the input and management of some CCR information, the CCR as a component is recognized as core to the CDR RA. With that said, the CCR Component is not meant to imply that CCR information is maintained and managed as a single, central instance.

Appendix D of this RA provides more detail for the CCR Component including several key realization patterns in support of other Core CDR Components.

4.2.3.4 Query Management Component

The Query Management Component is the primary mechanism that enables content consumers to build queries, store them, access the results, and allow other consumers the ability to subscribe to them. In this sense, this component supports asynchronous alerting of new content to subscribing consumers as well as asynchronous "pushing" of content resources.

4.2.3.5 Retrieve Component

The Retrieve Component enables synchronous access to discovered content and is the primary mechanism for content consumers to access one or more specific content resources from a specific content collection. The Retrieve Component then locates content resources related to the incoming request and directly returns that resource to the requesting consumer.

4.2.3.6 Receive Component

The Receive Component enables asynchronous access to discovered content and is the primary mechanism for content consumers to reroute discovered results to a different consumer. Generally in these types of interactions, a consumer provides a Retrieve Component with not only a request for content resources but also routing information as to where those content resources should be ultimately delivered. In this case, the provider may again locate the desired content resource as in the synchronous model, but instead of returning it immediately to the requesting consumer, the provider may rely on a messaging component (refer to Section 3.2.4.3) to route the content to the location. To support this scenario, the requested location must implement a Receive Component.

4.2.4 Key CDR Dependency Components

These components represent the key external dependencies needed to realize the CDR architecture. Specific guidance around the realization and implementation of these component types are out of scope for the CDR RA, but they are described here to

---

17 *Static information includes, but is not limited to, features of the exchange schemas (both discovery and content) that have been engineered into the collection. For example, the Provider may declare that it has implemented the DDMS 2.0 schema for discovery metadata, and that all metacards will reference "Unclassified" or "Secret" data resources.*

18 *Dynamic information includes, but is not limited to, features of the exchange schemas (both discovery and content) that are present in the current collection. Issuing queries against the content collection produces this information. For example, the Provider may declare that for the DDMS 2.0 schema that has been implemented, and there are a total of 1200 metacards (instances of the DDMS schema).*

illustrate the high-level dependencies between Core CDR Components and other applicable Enterprise service activities.

4.2.4.1 Security Component

The Security Component impacts each of the components in the CDR RA as depicted in Figure 3 above. Generally, the Security Component provides the authentication, authorization, integrity, confidentiality, and non-repudiation functions. Interfaces between security components such as policy enforcement, decision points and the Core CDR Components will be necessary to ensure that only authorized consumers gain access to content and metadata. When a consumer searches for content collections, the Search Component utilizes the Security Component to build and exchange consumer authentication and authorization data to determine if the consumer can discover specific metadata about a given content collection. Furthermore, this function is also used to determine if the searching consumer can have access to specific content resources.

The Core CDR Components utilize the Security Component to protect the confidentiality and integrity of the content exchanged between the consumer and the provider. Also, interactions between the Core CDR Components utilize the Security Component to pass and maintain consumer credential and identity data with each request while ensuring that this data is kept confidential. Furthermore, the Core CDR Components use the Security Component to enable cross-domain content discovery and retrieval, to protect content transfer between security domains, and to prevent unauthorized access from one security domain to any of the others.

To allow for both anticipated and unanticipated consumers (e.g. U.S. Coalition consumers) to discover providers and perform queries, the Core CDR Components use the Security Component to obtain access and policy information as well as to determine whether to permit or deny access to the Core CDR Components based on the defined policy rules. The Core CDR Components also utilize the Security Component to allow for certified identity and attribute data of authorized and unanticipated consumers to be associated with search requests so that authorization decisions can be made in a distributed manner.

4.2.4.2 Service Discovery Component

The Core CDR Components rely on the Service Discovery Component for mechanisms to discover distributed content across the IC/DoD Enterprise. When the consumer searches for content providers, the Service Discovery Component provides the support necessary to discover any applicable/relevant content collections that support the needs of the consumer's request.

4.2.4.3 Messaging Component

The Messaging Component provides the support for search requests and query results to be delivered synchronously or asynchronously. It provides a framework to allow consumers to subscribe to content providers so they are notified when content is updated or published.

4.2.4.4 Service Description Component

The Service Description Component provides the information about a service and the service provider to assist consumers in understanding the capabilities of the given service. It provides descriptive information about the service (e.g., service name), details about how to access the service, along with any policy and constraint information associated with the service. This component works in conjunction with the Service Discovery Component so consumers can become aware of content providers and query them.

5 Architecture Patterns

The following architecture patterns describe how the components in the CDR RA may be used in various contexts. Using this information, a Consumer or Content Collection provider may determine which patterns work best to meet their needs. These patterns are not an exhaustive list of all possible integrations; however they provide several example usages that build upon the basic interaction patterns.

For this release, the CDR RA discusses seven core interactions that are used as a baseline for advanced architecture patterns. For each pattern, the interactions among the components are shown as a sequence of steps, with each step containing corresponding key architectural information indicative of the information that the receiving aspect of the interaction must have to proceed. For example, a Search Component receiving a query must understand the query type being sent and the semantics of the query parameters being used. While the target values assigned to the query parameters convey the current needs of the requester, these values are part of the usage of the architecture and not an architectural concern.

The specific details of what happens within the bounds of Core CDR Dependency components are not discussed in this section, though references to these patterns are provided. For instance, security concerns crosscut nearly every aspect of the following interaction patterns. However, these concerns are excluded from the diagrams in the interest of clarity and to focus on the core interactions between components introduced in this document or those external components that are not crosscutting in nature. This does not imply that security is not pertinent to these interactions in any way[19].

19 *For a full discussion on security patterns, implications, and details, please refer to the Service Security Reference Architecture, referenced in Appendix B item #4*

5.1 Interaction Pattern 1: Direct Search

CDR Interaction Pattern 1: Direct Search

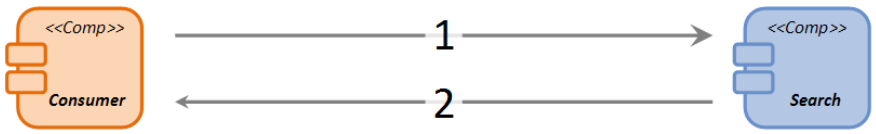

Figure 4 – CDR Interaction Pattern 1: Direct Search

Step	Description	Key Architectural Information
1	A Consumer Component sends a search request to a Search Component.	Query information, including well-defined query type and query parameters
2	The Search Component performs the search and return metadata describing the data resources that fit the criteria outlined by the search request to the Consumer Component.	Format used in returning metadata results

5.2 Interaction Pattern 2: Direct Asynchronous Search

CDR Interaction Pattern 2: Direct Asynchronous Search

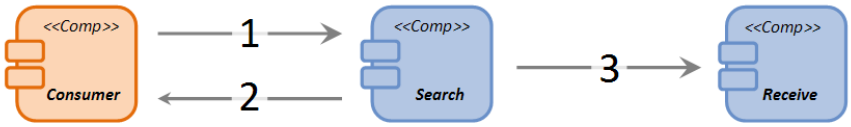

Figure 5 – CDR Interaction Pattern 2: Asynchronous Direct Search

Step	Description	Key Architectural Information
1	A Consumer Component sends a search request and a Receive Component reference to a Search Component implementation.	Query information, including well-defined query type and query parameters; location of the Receive Component, where results will be forwarded
2	The Search Component sends a response back to the Consumer Component containing a unique identifier for the search request that may be used to interrogate the Search Component for the current status of processing the search request, or to request that the Search Component cancel the processing of the search request.	Unique Search identifier
3	The Search Component performs the search and asynchronously sends the metadata describing the data resources that fit the criteria outlined by the search request to the specified Receive Component.	Format used in returning metadata results

5.3 Interaction Pattern 3: Brokered Search

CDR Interaction Pattern 3: Brokered Search

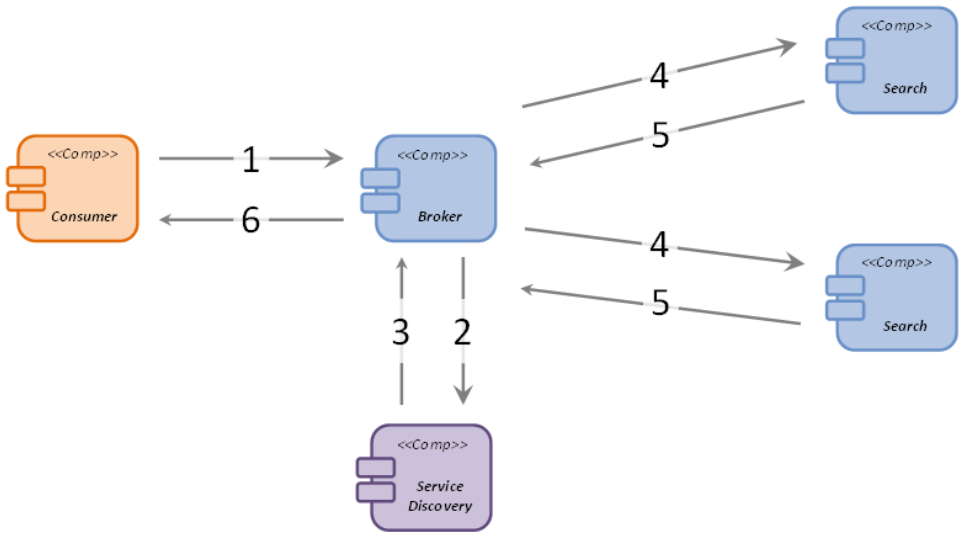

Figure 6 – CDR Interaction Pattern 3: Brokered Search

Step	Description	Key Architectural Information
1	A Consumer Component sends a search request to a Broker Component.	Query parameters
2	The Broker Component queries the Service Discovery Component for "relevant" Search Components, based upon the context of the search request, and the supported capabilities of the available Search Components.	Metadata describing supported capabilities of available Search Components; message exchange standards between core architectural components
3	The Service Discovery Component returns the "relevant" Search Components.	Message exchange standards between core architectural components
4	The Broker Component distributes the search request to the identified Search Components.	Query information, including well-defined query type and query parameters
5	The Search Component performs the search, and returns metadata describing the data resources that fit the criteria outlined by the search request to the Broker Component.	Format used in returning metadata results
6	The Broker Component aggregates the distributed results from the invoked Search Component into a single result set and returns it to the Consumer Component.	Format used in returning metadata results, including identifier of associated Search Component

5.4 Interaction Pattern 4: Content Retrieval

CDR Interaction Pattern 4: Content Retrieval

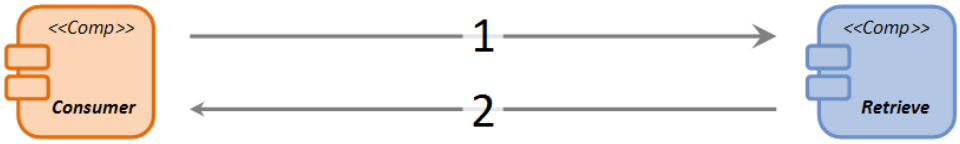

Figure 7 – CDR Interaction Pattern 4: Content Retrieval

Step	Description	Key Architectural Information
1	A Consumer Component sends a retrieve request to a Retrieve Component.	Query information, including well-defined query type and query parameters
2	The Retrieve Component executes the retrieve request and returns the Content Resources that fit the criteria outlined by the retrieve request to the Consumer Component.	Format used in returning content results

5.5 Interaction Pattern 5: Asynchronous Content Retrieval

CDR Interaction Pattern 5: Asynchronous Content Retrieval

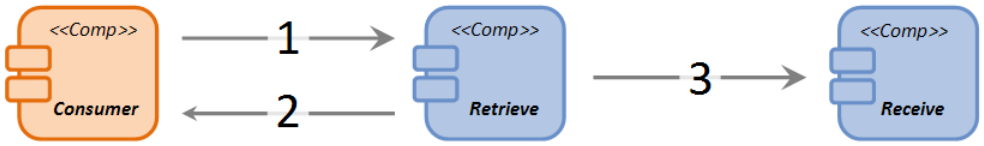

Figure 8 – CDR Interaction Pattern 5: Asynchronous Content Retrieval

Step	Description	Key Architectural Information
1	A Consumer sends a retrieve request and a Receive Component reference to a Retrieve Component.	Query information, including well-defined query type and query parameters
2	The Retrieve Component sends a response back to the Consumer Component. This response contains a unique identifier for the retrieve request that may be used to interrogate the Retrieve Component for the current status of processing the retrieve request. Or, this response can be used to request that the Retrieve Component cancel the processing of the retrieve request.	Unique Search identifier
3	The Retrieve Component executes the retrieve request and asynchronously returns the Content Resources that fit the criteria outlined by the retrieve request to the specified Receive Component.	Format used in returning content results

5.6 Interaction Pattern 6a: Query Management – Save Query

CDR Interaction Pattern 6a: Query Management – Save Query

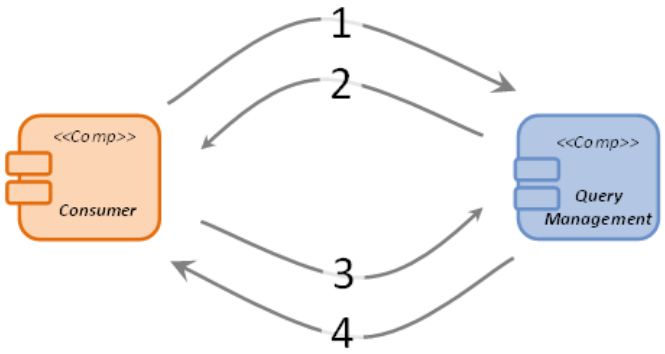

Figure 9 – CDR Interaction Pattern 6a: Query Management – Save Query

Step	Description	Key Architectural Information
1	A Consumer Component sends a search request to the Query Management Component to persist.	Message exchange standards between core architectural components[20]
2	The Query Management Component saves the search request and returns a reference to the saved search request to the Consumer Component.	Format used in returning metadata results
3	At some point later, a Consumer Component requests to invoke the search request by providing a reference to the saved search request to the Query Management Component.	Message exchange standards between core architectural components
4	The Query Management Component retrieves the saved search request and executes the query, returning the results to the Consumer Component.	Format used in returning metadata results

20 *Query Management Component needs to accept query package that will be stored but does not itself need to parse query or related architectural information to process query.*

5.7 Interaction Pattern 6b: Query Management – Alerting

CDR Interaction Pattern 6b: Query Management – Alerting

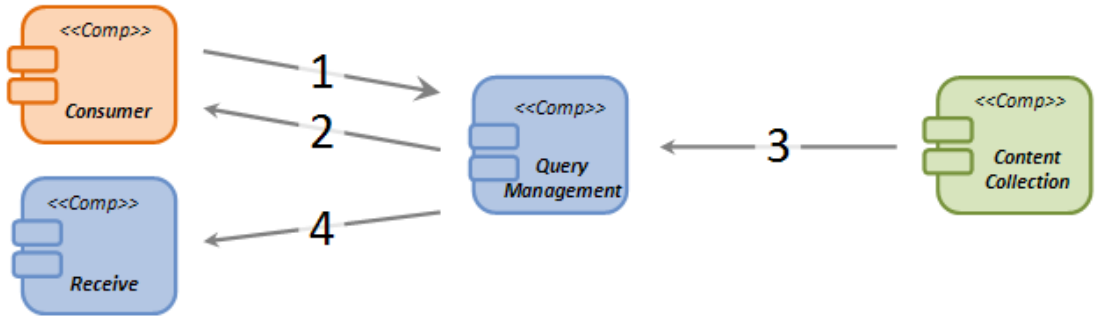

Figure 10 – CDR Interaction Pattern 6b: Query Management – Alerting

Step	Description	Key Architectural Information
1	A Consumer Component registers a subscription update from a Query Management Component for a search request or content resource change.	Subscription Query information, including identity of what is being monitored, what changes are being monitored; location of the Receive Component, where results will be forwarded
2	The Query Management Component persists and activates the subscription, returning a success indicator to the Consumer Component.	Format used in subscription success notification
3	A Content Collection that is subject to the Consumer Component's subscription notifies the Query Management Component of a change.	Format used in sending change information of Content Resource
4	The Query Management Component determines that a change relevant to one or more subscriptions has occurred and notifies the Consumer Component-specified Receive Component of the change by providing Search Component-like results.	Format used in returning metadata describing the changes Content Resource

5.8 Interaction Pattern 7: Content Collection Representation

The following interaction pattern is an example of how the Content Collection Representation could be used in conjunction with a Search Service, but this interaction can be applied across CDR RA Components.

CDR Interaction Pattern 7: Content Collection Representation

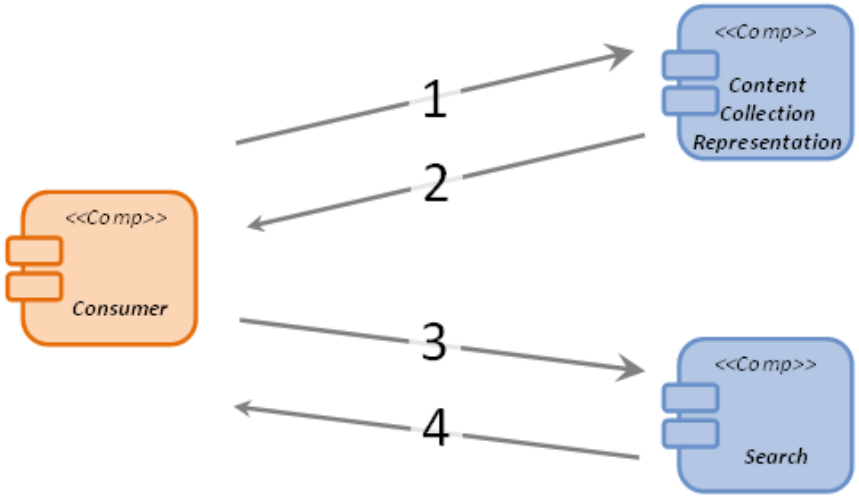

Figure 11 – CDR Interaction Pattern 7: Content Collection Representation

Step	Description	Key Architectural Information
1	A Consumer Component requests information describing a Content Collection Component's context, access constraints, and/or current inventory from the Content Collection Representation Component.	Message exchange standards between core architectural components, including identifier of component of interest
2	The Content Collection Representation Component returns the requested descriptions to the Consumer Component.	Format used in returning metadata describing the Content Collection
3	A Consumer Component sends a search request to a Search Component.	Query information, including well-defined query type and query parameters
4	The Search Component performs the search and returns metadata describing the data resources that fit the criteria outlined by the search request to the Consumer Component.	Format used in returning metadata results

Appendix A – Acronyms and Abbreviations

Table 1 – CDR RA Acronyms and Abbreviations

Acronym	Definition
API	Application Programming Interface
CCR	Content Collection Representation
CDR	Content Discovery and Retrieval
CIO	Chief Information Officer
COI	Community of Interest
COTS	Commercial-off-the-Shelf
DDMS	DoD Discovery Metadata Specification
DISR	DoD IT Standards Registry
DKO	Defense Knowledge Online
DNI	Director of National Intelligence
DoD	Department of Defense
DRM	Data Reference Model
DTIC	Defense Technical Information Center
ES ERB	Enterprise Services Engineering Review Board
FEA	Federal Enterprise Architecture
FIPS	Federal Information Processing Standards
FTP	File Transfer Protocol
GIS	Geospatial Information Service
GML	Geography Markup Language
GOTS	Government-off-the-Shelf
HTML	Hypertext Markup Language
IC	Intelligence Community

Acronym	Definition
ICSR	Intelligence Community Standards Registry
IETF	Internet Engineering Task Force
IPT	Integrated Project Team
ISO	International Organization for Standardization
IT	Information Technology
KIP	Key Interface Profile
MIME	Multipurpose Internet Mail Extensions
NGA	National Geospatial-Intelligence Agency
OASIS	Organization for the Advancement of Structured Information Standards
PDF	Portable Document Format
RA	Reference Architecture
RFC	Request for Comment
RSS	Real Simple Syndication
SOA	Service-Oriented Architecture
U.S.	United States
UTC	Coordinated Universal Time
W3C	World Wide Web Consortium
XML	Extensible Markup Language

Appendix B – Reference Documents

The following references were used in the development of the CDR Reference Architecture:

1. IC Services Strategy, v3.2, February 2008
2. DoD Net-centric Services Strategy, v1.0, February 2007
3. OASIS SOA Reference Model, v1.0, October 2006
4. Joint IC/DoD Security Reference Architecture, v1.2, July 2008 (v2.0 set for release in early 2010)
5. IC Service Discovery Reference Architecture, v1.2, September 2007 (Joint IC/DoD v1.0 set for release in early 2010)
6. IC SOA Messaging Reference Architecture, v1.1, December 2006
7. Joint IC/DoD Memorandum of Agreement for Services-Based Environment, July 2007
8. CDR Requirements Master List, October 2009
9. IC/DoD Unified Authorization and Attribute Service, Concept of Operations, v1.11, December 2008

Appendix C – Search Patterns

The Master CDR IPT Discovery and Retrieval Requirements List indicates consumers should be able to conduct searches using a combination of facets, including but not limited to keyword, geospatial, and temporal facets. Anecdotal evidence from participants collected during the DoD's Empire Challenge 09 noted a high priority for searching on these three facets in particular. This appendix will discuss each of these facets in more detail (see figure below). Search patterns may include additional concepts not explicitly discussed in this appendix, such as ontology-based mechanisms.

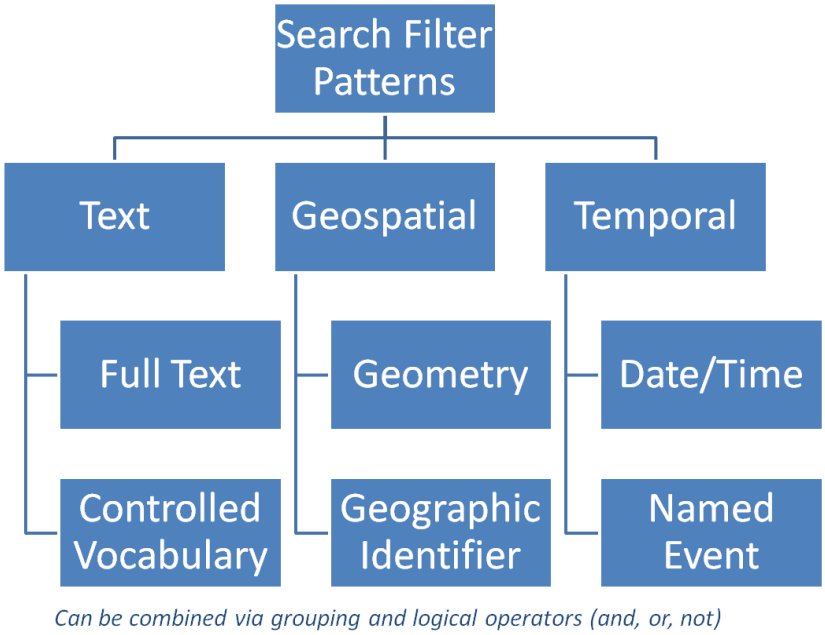

Figure 12 – Search Patterns

Text Search Patterns

Full text Search

Most popular search products today, either on the web or internal to the Enterprise, focus on full text search. Best illustrated by the major web search engines, full text search approaches commonly target unstructured and semi-structured data such as HTML, RSS/Atom, PDF, and MS Word. Many tools are also able to integrate structured sources into their platforms with special adapters that flatten the structured data into forms that make more sense with a syntax-based search. Also, for many potential content collection providers, full text search is the lowest barrier to entry and will enable more content collection providers to implement the specifications and to participate in the collaboration.

While full text searches are intuitive for human users, there are some shortcomings with this approach. First, while the software appears to "understand" what you are asking for,

in reality the tools leverage advanced mathematical or statistics-based algorithms to determine relevance. In general, these algorithms do not concern themselves with the semantics of the query, rather just the words themselves. Unless tuned properly, these algorithms could potentially lead to a significant number of false-positives. Furthermore, many of these tools, such as Google, rely on the pre-indexing of data to achieve acceptable levels of performance. The downside of this is the data may not always be fresh depending upon the last index point of the corpus of data resources.

Controlled Vocabulary Search

Since full text search approaches support the lowest level of semantic disambiguation, additional information is needed to enable a query with more business context. The DoD Net-Centric Data Strategy advocates the development and use of COI vocabularies that are instrumental in differentiating the use of terms having different meanings in different communities. Candidates of controlled vocabularies include the Intelink Topic Directory, the UCore default taxonomy, the DTIC Thesaurus, and the DoD Core Taxonomy. The Intelink Topic Directory is a high-level taxonomy that breaks down the different "topics" of interest within the IC. Some of the topics might include "Air and Air Defense" or "Geospatial/Imagery." As well as using these controlled vocabularies to characterize content resources, these vocabularies are also useful to characterize services in service registries.

XML-Based Search

XML-based search models, unlike the text search models described above, allow consumers to discover content resources based on desired values for specific fields within a particular community XML data model. Combining a particular XML data model with a W3C standard filtering syntax for XML (such as XPath or XQuery) provides a very high fidelity means to query and discover metadata or a content collection. As well-defined community discovery metadata specifications such as the DoD Discovery Metadata Specification (DDMS) and data specifications such as UCore, XML Pubs, and COI-coordinated exchange specifications are registered in Enterprise-level metadata registries and gain traction within the IC and DoD, community members can retrieve the specifications and other amplifying documentation to develop a shared understanding of the business meaning. In turn, they will be able to exchange information with other members of the community that implement and support those data specifications.

Geospatial Search Patterns

Much progress in harmonizing geospatial representations has been made in the international standards organizations, such as the ISO, both in terms of geometry and named locations, such as a geographic identifier like country code. The Open GIS Consortium has made significant progress towards developing common physical representations of geometric geospatial standards via the Geography Markup Language (GML), thereby promoting better interoperability. In the case of named locations, such as country codes, the discovery metadata and the search should conform to federal announcements such as the rescinding of FIPS 10-4. NGA is working with several DoD

and IC agencies to develop a migration plan for developing common geospatial standards.

Temporal Search Patterns

In addition to using common text strings and geospatial coincidence, a popular search filter pattern is to search based on temporal coincidence, whether it is based on a discrete date/time or a named event. The common representation for date/time format is ISO 8601, UTC format. The advantage of this 'CCYY-MM-DDThh:mm:ssZ' formatting is that values can be compared via simple string comparisons, and it is in widespread use in both DoD and IC. The common use of Greenwich-based date/time representations further simplifies filtering by date/time. Searching by named events, however, is problematic due to the lack of standardization for named events.

Appendix D – Content Collection Representation Realization Patterns

As introduced in Section 3 of the Content Discovery and Retrieval Reference Architecture (CDR RA), the Content Collection Representation (CCR) Component "serves as the primary mechanism for content collectors (i.e., providers) to expose information to describe the context, access constraints, and current inventory status of its underlying content resources to support static and dynamic discovery and accessibility". Given the heterogeneous nature of the Enterprise and the need to ensure that the CDR RA is resilient and interoperable with the constantly evolving state of technology and implementation, abstract components are at the heart of the CDR RA interaction patterns. The CCR Component provides the connectivity that matches generalized functionality to specific implementations, allowing new implementation options to be introduced to the community on demand while continuing to provide support for prior realizations of the CDR RA. While the CCR Component's capabilities are a key enabler to the CDR RA functionality, it may not be realized into custom specification that directly realizes CCR functionality. Rather, this notional component will likely be realized through one or more CDR RA Components.

This CDR RA Appendix provides guidance on how one might realize the CCR Component under the auspices of the Service Discovery, Search, and Query Management Components, individually, or in concert with each other. It is intended as "**Possible Realization Guidance"**, meaning that programs within the IC/DoD should review it and consider adopting the architecture and utilize the patterns that it describes.

"**Possible Realization Guidance**" **IS NOT** any of the following:

- A DoD Key Interface Profile (KIP)
- A DISR Entry
- A mandatory IC Standards Repository (ICSR) entry
- Mandated by any DoD or IC policy or memorandum

Following the guidance is **optional**, but it is recommended to facilitate interoperability of IC/DoD content. After an initial period of use, this guidance may be considered for endorsement as more formal, mandated guidance (mandated by an IC/DoD policy/memorandum or defined as a KIP or KIP-equivalent).

CCR Component Realization Pattern 1: Via Service Discovery

Many organizations with the IC/DoD Enterprise have been leveraging the concept of Service Discovery for some time, often registering their services and associated metadata artifacts to expose functionality across enclaves. Given the existing level of adoption, Service Discovery provides a well-understood path towards realizing the CCR Component.

Registering Static Information

Service Discovery Registries are well established for registering services with associated metadata. A realized service is typically registered with a series of descriptions that can be used to visually or programmatically determine the usefulness of that service to potential consumers. In the CDR RA, this concept is extended to register information based on the type of CDR RA Component. For instance, a Content Collection Provider may register a service that realizes the CDR RA Search Component within Service Discovery. In doing so, registered information may include the name and description of the component and the types of queries it supports.

Often, Service Discovery providers require that service information is submitted to administrators who manually enter this information into a Registry, though this is not a universal requirement. Regardless of the entry mechanism, this level of CDR RA Component realization information is unlikely to change at a frequent rate, making it a solid fit for traditional Service Discovery Registries. However, this slow-changing information only provides one aspect of the metadata needed to fully determine a CDR Component's potential usefulness for a consumer. Dynamic information must also be considered to have a fully capable realization of the CCR Component.

Registering Dynamic Information

For content collections, CDR Component realizations need to describe requisite static informational concepts such as requests, responses, and implementations. Additional information about content collections may change on a frequent basis, making it a poor fit for manual registration within Service Discovery. Consumers need to be able to not only discover those services that expose data via the type request and response they desire, but also those that have characteristics in terms of data, performance, or other fungible classifications. For instance, a consumer may only be interested in content collections that contain data less than 24 hours old. By allowing dynamic data to be registered in Service Discovery, a content collection could advertise the "freshness" of their data, allowing consumers to select only those CDR Search Component realizations that advertise that they contain applicable data. Similar arguments can be made for a wide variety of information, ranging from geospatial coverage of content collections to Quality of Service metrics/measures.

While there are many use cases for registering and using dynamic information in conjunction with Service Discovery, the realization options for this functionality is less straightforward, particularly given the commonality of Service Discovery mechanisms that require an administrator to update registries. Several possible mechanisms exist to help fulfill the promise of dynamic information registration within a Service Discovery realization, ranging from information pulling mechanisms, such as a CDR Search Component, to pushing mechanisms, such as a CDR Query Management Component, which could be incorporated into Service Discovery to aid in the registration and maintenance of static and dynamic information.

CCR Realization Pattern 2: Via Search

Another option for realizing the CCR Component is through a CDR RA Search Component. Such a solution would allow providers to expose their CCR Component data through the same infrastructure they would leverage to expose other CDR RA Search Components (if applicable), allowing this data to be requested (or pulled) by consumers on demand. This mechanism may be particularly effective for exposing dynamic information, since the underlying CDR Search Component realization can always return real-time results. Additionally, the CDR Broker pattern could be leveraged to discover pertinent static and dynamic information across the Enterprise, reusing CDR RA Broker Component patterns to achieve distributed CCR Component functionality.

However, there are some downsides to this approach. For instance, consumers may need to cache data to achieve better performance characteristics in their implementations, reducing the benefits of information dynamically obtained by the CCR Component information. Additionally, such an implementation may be considered overkill by providers that have static data or highly consistent dynamic information that may be expressed in relative terms. For instance, extending on the concept of data freshness, a provider may prefer to define their data as always containing a relative freshness of less than 24 hours old, rather than providing something like a timestamp of "most fresh" in their content collection.

CCR Realization Pattern 3: Via Query Management Component for Alerting

An additional option for maintaining accurate static and dynamic CCR Component information is via the Query Management Component's Alerting capability. In contrast to a manual registration model or a pull-based model, Alerting provides a mechanism to push data to interested subscribers. It has many of the same implementation benefits of the Search Component model, without some of the potential drawbacks, such as inaccurate caching or frequent querying for updates if none are applicable.

While this model for pushing information to subscribers has many benefits, it may be not be well accepted by providers that have static data or prefer a "relative" mechanism for defining their dynamic data.

Choosing Realization Patterns

Selecting the most applicable realization pattern for the CCR Component will vary across users. Within enclaves, the cost and benefit of the approaches laid out in this Appendix must be considered to find an optimal solution for the enclave's needs. At the Enterprise level, a combination of approaches may be needed to allow different enclaves to participate without creating undue burden on any group of providers. Additionally, alternative realization mechanisms exist outside those described in this Appendix, and may prove more applicable in some cases.