

2006 CCRTS
THE STATE OF THE ART AND THE STATE OF THE PRACTICE

My Two Cats Are a Community of Interest

Policy
C2 Architecture
Social Domain Issues

Scott Renner
The MITRE Corporation
202 Burlington Road, MS M338
Bedford, MA 01730
781-271-8625 / 781-271-7231 (fax)
sar@mitre.org

My Two Cats Are a Community of Interest

Dr. Scott Renner
The MITRE Corporation
sar@mitre.org

Introduction

It is usual for new concepts to change as they spread from the original core practitioners to a wider community. Much of this change is desirable, as new practitioners contribute new aspects overlooked by the originators. Less desirable is the loss of precision that occurs when people climb aboard the bandwagon with unrelated concepts and agendas. Readers familiar with the development of object-oriented programming will recognize the phenomenon. As this programming paradigm became popular, people started calling all manner of things “object-oriented” in an attempt to attract attention and approval – a state of affairs satirized by the title of the 1989 paper, “My Cat Is Object-Oriented” [1].

A similar loss of precision has occurred with the concept of *Communities of Interest* (COIs), widely introduced to the DoD in 2003 by the Net-Centric Data Strategy (NCDS) [2] and the Air Force Information and Data Management Strategy [3]. In the past three years, the COI term has become heavily overloaded. It now means so many different things to so many people that one might very well say my two cats are a COI. They certainly live in community – I am always finding them curled up together in a chair, or sofa, or anyplace else I want to sit. And they definitely share interests – just open a can of tuna fish in my kitchen, and you’ll see exactly what I mean. QED.

Unfortunately, a term that means almost anything turns out to mean nearly nothing at all. This is a problem when you want to get some work done. The people who must do the work need a better idea of what they are supposed to do. The people with money want to hear something a little more specific about what they will get. We should expect continued friction and slowed progress until we can tell a precise, coherent story about what COIs are and what they do.

In this paper I will briefly recount the history of the COI term, summarize the variant meanings of “COI” in current usage, and propose a direction for a more satisfactory understanding. The main point: it’s now time to think less about COIs *per se*, instead turning our attention to the COI *members*, and their roles and responsibilities.

Well, How Did I Get Here? [4]

Our history begins in 1990, when the DoD data administration program was established by Directive 8320.1. This program can be understood as part of a general drive towards interoperability through standardization. Many interoperability problems are caused by needless differences in the name, structure, and representation of data – eliminate those differences, and you eliminate the problems. The attempt to establish a single

comprehensive enterprise data standard followed conventional wisdom of the 1980s; however, this data standardization approach proved infeasible for enterprises approaching the scale of the DoD [5]. By the end of the decade, the DoD data administration program was at a standstill, with some recognized pockets of success, but no real hope of meeting its overall goals.

The next chapter of our history is the Rapid Improvement Team for DoD Data Interoperability, chartered by ASD/C3I and USD/AT&L in November, 2000. The team's report [6] marks the first use of the COI term in the context of data interoperability. The team observed that data interoperability depends on shared semantics – not necessarily a physical or logical model, but something with enough detail to establish semantic agreement, a common understanding of the data's meaning. Second observation: if you can't have one comprehensive data standard covering the whole enterprise, then you must have more than one standard, each covering part of the enterprise. The COI term was introduced to describe the *semantic community*,¹ the subset of people within the enterprise who develop and understand a shared data model, and then exchange information based upon it. This concept is described in more detail in [7].

Our story now moves to the Air Force and its work on information and data management, beginning with the 2001 AF Scientific Advisory Board study on “Database Migration” [8]. In the Air Force work, the COI term was again used to denote a semantic community, one which includes proponents, architects, implementers, and users. The term “common vocabulary” was introduced and described in terms of shared knowledge that must be developed *and understood* by the COI members. The necessity of managing data ownership (and not just data definitions) was also explored [9].

The next episode is the DoD Data Management and Interoperability Broad Area Review, followed closely by the DoD Core Data Management Working Group, both established by the ASD/C3I in early 2002. The working group developed the Net-Centric Data Strategy (NCDS) during 2002 and 2003. At the beginning of the NCDS work, the COI concept was entirely about establishing shared semantics to enable information exchange. Over time, the concept was expanded to include two new dimensions: *institutional* vs. *expedient*, and *functional* vs. *cross-functional*. While the COI term still includes communities that exist to establish shared semantics, it expanded to include discussion groups and collaborative teams with some shared purpose. According to the COI FAQ [10], “every task-oriented workgroup (e.g., the bomb damage assessment cell at the Air Operations Center) can be a COI. Any collection of people with a declared interest (e.g., in biological warfare) can also be a COI.”

The story now moves to 2004, with COIs acquiring a role in IT portfolio management for the GIG [11]. Portfolio management establishes mission areas (e.g., “warfighter”) and subordinate domains (e.g., “battlespace awareness”) to govern investments in information capabilities and services. Mission area and domain managers are assigned responsibilities with respect to COIs. Some COIs will be created top-down by the domain managers. The members of these COIs are not the people exchanging

¹ The term “semantic community” was considered, but rejected as too technical.

information at runtime. They are not primarily trying to establish common semantics at build-time. Instead, these people are most interested in directing and influencing the creation of new and modified capabilities. They are largely concerned with the governance of acquisition. Their role in IT portfolio management is still unspecified, but might go something like this:

- Domain owners approve COI charters, adjudicate COI disagreements
- Programs of record (PORs) participate in approved and *ad hoc* COIs
- COIs establish agreements for data sharing capabilities (including vocabularies)
- Milestone decision authorities evaluate and enforce POR compliance with the agreements of approved COIs

The latest story chapter unfolds through 2005 with the development of several “pilot COIs” sponsored by OASD/NII to demonstrate and build experience with the NCDS.² These COIs are intended to develop and demonstrate a new, useful net-centric data sharing capability within nine to twelve months. Almost all of the funding is provided by the PORs, who implement the new capabilities not for their current deliverables, but must be included in the *next* acquisition spiral. High-level COI leadership is considered very important, and so the pilot COIs are expected to have an executive board of flag/general officers (O-8/O-9 level).

The Current Meanings of “COI”

I have found the following four definitions of “COI” in widespread use. They correspond roughly to the steps of the history above. When you hear somebody else mention COI in a conversation, they probably have one of these four definitions in mind.

- **Stovepipe COI:** An existing functional community with a large, established, and detailed community vocabulary. That vocabulary is completely optimized for their view of the world. They aren’t much interested in any possible overlap with other communities and other views of the world. When forced to share outside the COI, their typical reaction is, “we’ll extend our model, and then you can adopt it”. Some people think that *every* COI is a stovepipe COI – so if you use the term, they assume that’s what you’re talking about.
- **Vocabulary COI:** A community that has (or will develop) a common vocabulary for some subject-area domain of interest. As a result, they are *able* to exchange information, regardless of whether they all *actually do* exchange with each other. The level of detail and rigor in the common vocabulary will match the COI’s purpose; it could be a dictionary, or a taxonomy, or an implementation-level data model. Some vocabularies will be broad and shallow; i.e. containing a few definitions which can be used by almost everyone; these are sometimes called “loose connectors”, or “loose couplers” [12]. Some vocabularies will be narrow and deep; i.e., with many

² These pilot COIs include Blue Force Tracking (BFT), C2 Space Situational Awareness (C2-SSA), Global Strike (GS), and Maritime Domain Awareness (MDA), among others.

definitions used by a few people. (Broad and deep vocabularies are lovely, but are typically too expensive to develop.)

- **Sharing COI:** A community that *actually is* exchanging information with each other. This sharing depends on a common vocabulary, and so in a theoretical sense a sharing COI is always a subset of some semantic community. This information sharing also requires data producers who create and post the shared information. And it requires some form of shared information space from which consumers pull the data they need.³
- **Proponent COI:** A community that seeks to establish *new or improved* information sharing capabilities, often through their influence over acquisition. Typically these aren't the people who actually share the operational data. They usually aren't the people who develop and learn a common vocabulary. Instead, they collaborate to identify the desired capabilities, and direct (or influence) their own organization to implement the necessary change. These COIs serve as a catalyst for transformation.

One often encounters something named or proposed as a “COI” that has nothing to do with data sharing, something that is no more a “real” COI than my two cats. These are caused by bandwagon-climbers and by the sincerely confused. Both causes will diminish once we have a better understanding of the COI term.

What Next?

The meaning of the “COI” term has been enlarged a great deal over a relatively short time. It began as an equivalent for “semantic community”. It quickly grew to include the people actually sharing information. It expanded again to include proponents desiring new and better information sharing. At each stage, some aspect essential to net-centric data sharing was added. From a theoretical point of view, the *ad hoc* nature of this process, together with the loss of precision in nomenclature, is an annoyance. Viewed practically, this is all progress, was probably inevitable, and in no way should be undone.

There are two common elements in all this COI activity. First, COIs are formed to solve a problem – not just any problem, but a data sharing problem.⁴ Second, COI members cooperate to solve the problem. In a sense, COIs don't *do* anything – the *members* of the COI do everything.⁵ Based on these observations, I propose the following definition of a community of interest:

Within the context of the NCDS, a “community of interest” is defined as a group of people who cooperate to solve a problem in net-centric data sharing. To solve their problem, the members of a COI may need to form and promulgate a common vocabulary, which may be used for discovery metadata, machine-to-machine data exchanges, and architecture descriptions. The members may actually exchange information with each

³ There are many ways to implement this shared space; e.g., a relational database, a shared mailing list, a portal. The key aspects of a shared space are its content, and its control.

⁴ An observation made by others; for example, in [13].

⁵ Also observed by others, for example, in [14].

other on a subject of common interest. They may develop new capabilities to support new and improved information sharing. COIs do not overtly control resources or direct activities; instead, to accomplish their goals, they act entirely through the cooperation of their members and their respective organizations.

In this formulation, COIs solve problems through consensus and cooperation. How, then, to handle those decisions that require authority and direction? The pilot COIs formed by OASD/NII handle this problem through their executive boards. These pilots are relatively small efforts, and so, as a rule, if the flag and general officers agree, getting consensus and cooperation from everybody else will not be one of your big problems.

This is a fine way to get things started. But in the long run, we will need something more formal, if for no other reason than O-9s are scarce and may have other things to do. Developing a steady-state governance structure for information management begins by identifying the kinds of decisions to be made, and the roles of the people who make them. I have argued that in large enterprises the information management responsibilities need to be divided three ways, and governed by three kinds of entities:

- *Information owners and data producers* are organizations that exercise authority over data. They are responsible for the production of data in the enterprise. They acquire and operate information systems to carry out their ownership responsibilities. Their authority is to say what data will be collected and how it will be represented and stored. They are accountable for maintaining their data, and ensuring that it is protected against unauthorized access.
- *Shared information spaces* (or *infospaces*) are collections of data intended to suit the needs of different groups of data consumers. This data is exposed through information services⁶, which may be implemented in various ways. Each infospace has a *controller*, who serves the interests of and may have authority over the consumers. Infospace controllers have authority to say which producers may post to their infospace, what kinds of data they may post, and which of these sources will be authoritative for their infospace consumers. They also establish attributes for access control and the rules which specify the privileges of infospace consumers in terms of those attributes.⁷
- *Semantic communities* are groups of people who establish a shared understanding of terms and definitions in some subject-area domain; these are captured in a *common vocabulary*. The detail in these vocabularies may range from simple dictionaries used in architecture descriptions, to the detailed data models and elements required for machine-to-machine data exchanges. The authority of the semantic community is to

⁶ This term is defined in the DoD *Shared Service Strategy*, in draft as of this writing. It refers to the kind of service found in a service-oriented architecture (SOA).

⁷ Earlier drafts spoke of “roles” here; however, role-based access control (RBAC) is rightly giving way to attribute-based access control (ABAC).

control the definitions in the common vocabulary. They don't compel you to use their definitions – but someone else may do so (an information owner, perhaps).⁸ These are the three essential governance entities, in the sense that combining any two will lead to conflicts [15]. How should these responsibilities be assigned within the enterprise? I have speculated that governance of information owners should follow the existing organizational hierarchy; i.e., the components, services, and agencies. Shared information spaces should be controlled by people with authority over information consumers: combatant commanders, or their equivalent in the business mission area. Common vocabularies should be developed by consensus within semantic communities; some of these might be created via top-down authority, and some form with bottom-up spontaneity. These ideas are still tentative, and require further consideration.

COIs are best understood as problem-solving teams working at the intersection of information owners, shared information spaces, and semantic communities. As depicted below, these governance entities will have interests that extend beyond any single COI. We can expect COIs to form around high-value data sharing opportunities, to bring about this intersection of the participants that need to share. We should expect COIs to ascertain the governance decisions they should make. We should not expect the COIs to make and enforce these decisions.

For example, net-centric data sharing must involve access control decisions, some concern for who is allowed to see what. In general, we cannot support net-centric data sharing if we have to make that decision for every single producer-consumer pair [16]. We should expect COIs to *help negotiate* community agreement between producers and

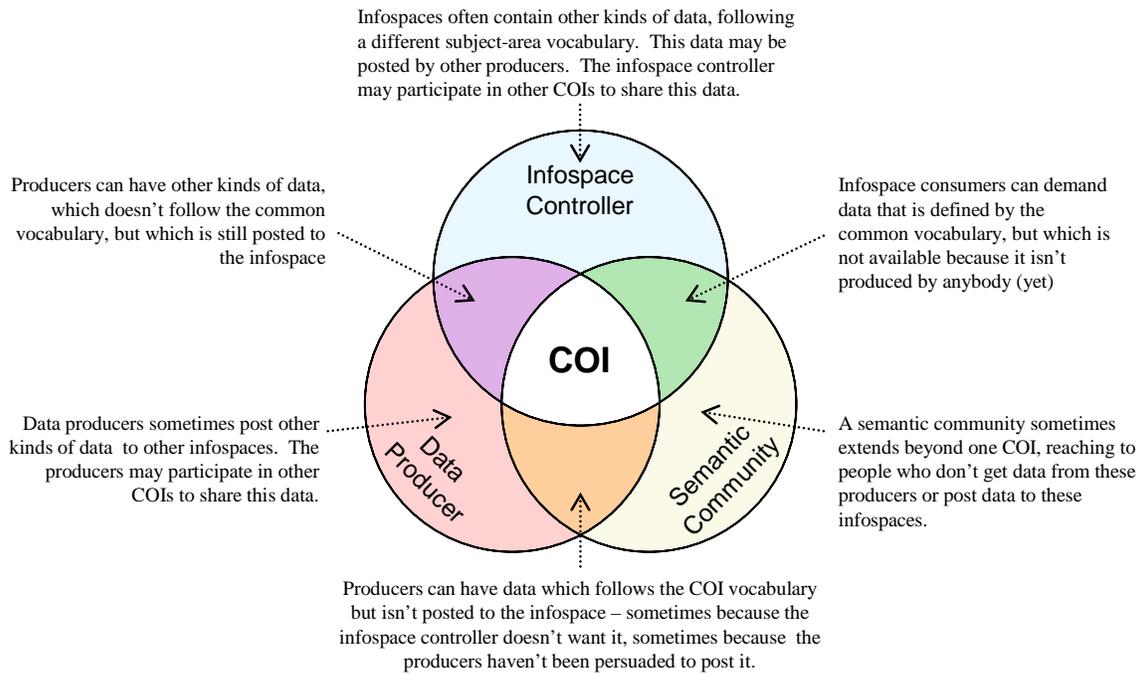


Figure 1: COI members and their data interests outside the COI

infospace controllers. However, the principals in this agreement will be COI members; we should not expect the COI *per se* to be one of the principals involved.

The next step is to move beyond speculation, and establish how the governance entities should be mapped to the enterprise organization. Who are the information owners and the infospace controllers? What do those people do? Who makes those assignments? These answers must be detailed in roles and responsibilities, captured in policy and procedures.

Conclusion

Over time, the COI concept has evolved, expanded, and is now reaching a stable equilibrium. COIs work to solve problems in net-centric data sharing; not as an authority giving orders, but instead as a group cooperating to reach solutions. This area of cooperation lies at the overlap of those entities that must operate through authority and direction: the information owners, and the infospace controllers.

Those entities are not clearly defined at present. COIs are able to operate for now with a partial, sometimes vague understanding of information management governance. However, we are nearing the point where detailing the roles and responsibilities of IM governance becomes necessary to further progress. With these roles captured in policy and procedures, more COIs will spring up to produce more data sharing – and will then be far more productive than my two cats.

Acknowledgements

I am grateful to Shelby Sullivan and Phil Barry for long conversations hammering out the details of the Air Force Information and Data Management Strategy. The material in this paper was greatly influenced by that work. Thanks also to Kevin Gunn, Mary Ann Malloy, Bob Miller, Arnie Rosenthal, and Dan Winkowski for their helpful review and suggestions.

References

- [1] R. King, “My Cat Is Object-Oriented”, in *Object-Oriented Concepts, Databases, and Applications*, F. Lochovsky, ed. New York: ACM Press, 1989. pp 23-30. <http://doi.acm.org/10.1145/63320.66469>
- [2] DoD Chief Information Officer, *DoD Net-Centric Data Strategy*, March 2003. <http://www.defenselink.mil/nii/org/cio/doc/Net-Centric-Data-Strategy-2003-05-092.pdf>
- [3] U.S. Air Force, *Air Force Information and Data Management Strategy Policy*, March 2004. https://www.my.af.mil/gcss-af/USAF/cms/AFMC/files/R2L-mZ0_0+63+-y-rOJ_AF_Information_Data_Management_Policy.pdf
- [4] Talking Heads, “Once In a Lifetime”, on *Remain in Light*, Warner Brothers, 1980. <http://www.amazon.com/exec/obidos/tg/wma-pop-up/-/B000002KO3001004/002-8020646-3258408>
- [5] S. Renner, “Improving 8320.1 Data Administration”, in *Proc. Federal Database Colloquium*, San Diego, 1998.
- [6] Department of Defense, *Data Interoperability Rapid Improvement Team Report*, Nov. 2000.

- [7] S. Renner, "A Community of Interest Approach to Data Interoperability", in *Proc. Federal Database Colloquium*, San Diego, 2001.
http://www.mitre.org/work/tech_papers/tech_papers_01/renner_community/renner_community.pdf
- [8] Air Force Scientific Advisory Board, *Database Migration for Command and Control*, SAB-TR-01-03, Nov. 2002. <http://handle.dtic.mil/100.2/ADA412520>
- [9] Air Force Chief Information Officer, *Data/Information Management*, Chief Architect Office Discussion Paper DP-009, Oct. 2002.
- [10] DoD Chief Information Officer, *Communities of Interest in the Net-Centric DoD: Frequently Asked Questions*, 2004. https://disain.disa.mil/nctdesc/library/COI_FAQ.pdf
- [11] Deputy Secretary of Defense, "Information Technology Portfolio Management", memorandum, Mar. 2004.
- [12] R. Byrne, "Getting DoD Linked", The MITRE Corporation, Jan. 2005.
http://www.mitre.org/work/tech_papers/tech_papers_05/04_1260/
- [13] S. Chamberlain, M. Boller, G. Sprung, and V. Badami, "Establishing a COI for Global Force Management", in *Proc. 10th Int. C2 Research and Technology Symposium*, McLean, VA, June 2005.
<http://www.dodccrp.org/events/2005/10th/CD/papers/151.pdf>
- [14] E. Ramming, "COI Primer", USSTRATCOM, July 2005.
<https://gesportal.dod.mil/sites/stratcom-codification/default.aspx>
- [15] S. Renner, "Net-Centric Information Management", in *Proc. 10th Int. C2 Research and Technology Symposium*, McLean, VA, June 2005. <http://www.dodccrp.org/events/2005/10th/CD/papers/348.pdf>
- [16] S. Renner, "Building Information Systems For Network-Centric Warfare", in *Proc. 8th Int. C2 Research and Technology Symposium*, Washington, DC, June 2003.
http://www.dodccrp.org/events/2003/8th_ICCRTS/pdf/078.pdf