In: 8th Int. C2 Research and Technology Symposium, Washington, DC, June 2003

Building Information Systems for Network-Centric Warfare

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

Abstract

This paper examines the implications of network-centric warfare for information system development: How should we build C2 information systems for net-centric operations? We begin with six highly-probable predictions for the NCW future. From these we derive a number of present implications for system development: things we should do now, and problems we will have to solve along the way. Our answers touch on the information technology to be employed within the systems, the architectural principles that will guide and structure their development, and the acquisition process used to build and deploy the systems.

1. Introduction

Network-centric warfare (NCW) is a theory of military operations which holds that the seamless networking of the friendly force elements will bring about an increase in combat power [1]. This "networking" is not merely a communications network, the sort of thing that is implemented over physical cables, radio links, TCP/IP, and the like. These things are necessary, but are not among the key aspects. Instead, the "network" in NCW emphasizes a network of connections between people in the information and cognitive domains. It stresses the shared information and situational awareness that leads to increased speed of command and synchronized effects in the battlespace.

We will not argue the correctness of NCW in this paper. Instead, we assume the theory is correct, will be pursued, and will ultimately prove successful. Our interest is in the technology, architecture, and management needed to build net-centric C2 information systems. We are especially interested in how these information systems will collectively implement the seamless networking in the *information domain* – how they will supply the right information at the right time to the right decider so that he can make the right decision. As information technologists, our responsibility is to predict the technology developments as they emerge and to help consider how these might be applied to command and control problems. As information system architects, our responsibility is to help apply information technology (IT) to new mission capabilities, to serve as guides in exploring the search space of what is operationally desirable, what is technically feasible, and what is practically affordable in time and money. As information managers, our responsibility is to develop and implement the policy and procedures needed to ensure that the right information is collected, maintained, and made visible and accessible to the deciders who need it.

It is not instantly obvious how to best meet these responsibilities of information technologists, architects, and managers. But we can obtain some current directions by first considering what the destination will be like.

2. Predictions of the NCW Future

Assume that the DoD pursues the NCW concept and makes the best possible progress. What will the world look like in, say, fifteen years? We believe that the following six predictions are easily defensible, and fundamental to understanding the role of IT in net-centric operations:

1. We will have nearly all of the robust, seamless communications network connectivity we require. We will actually have a communications *internetwork,* composed of tactical radio nets, satellites, microwave and landline links, etc. We will call it seamless for two reasons: First, because we will almost always be able to transmit some data between any two participants. Second, because the technical difficulties of linking the separate network types will be hidden from most developers and users.

2. There will be very many participants on that single seamless network, on the order of 10^6, perhaps 10^7. Almost every battlefield entity will have a network presence. In addition to the weapon platforms, application servers, and C2 user information appliances, we will see a vast number of simple sensor devices on the network. Many of these participants will be fully automated. All will need to exchange information with some other participants.

3. Bandwidth limits will still be a problem, especially as we get closer to the combatants. We can always increase the capacity of the fixed landline segments of our network to meet increasing demand. This will not always be possible for satellite and especially tactical radio communications. In short, we will be able to get *some* data to everyone, everywhere… but not always *all* the data anyone could want, anywhere.

4. Information assurance concerns will still be critical. Our information systems will be a high-value target to any adversary.

5. Information technology (and the people who understand it) will become much less expensive and therefore widely available to adversaries. There will be little competitive advantage in IT *per se*. Advantage will come from knowing how to best employ the technology that will be available to everyone.

6. Working out the best ways to employ IT will be an iterative process; a co-evolution of technology, doctrine, and organization. Change will be the only constant. Making that iterative process go quickly will maximize our advantage.

3. Preconditions for the Predicted Success

These six predictions describe a world in which NCW has been successfully implemented. From those predictions we can derive certain useful implications, or preconditions for our success in building net-centric C2 information systems: things we will have to build, or processes which we will have to learn to perform, before we will be able to build the information systems we need.

3.1 A Foundation Layer of Enterprise Services (implied by predictions #1, #4, #6)

Robust, seamless connectivity between C2 information systems depends on a set of common foundation enterprise-wide services. The bottom layer of this foundation is the single network service that can transfer data between any two participants. Other such common services will include identity management, authentication, and authorization. All these enterprise services must work throughout the whole environment – any observable seam will result in a barrier across which information cannot flow.

However, we must not purchase this seamless operation at the price of a rigid infrastructure and inflexible mission applications. The ability to quickly implement new and changed operational capabilities is essential to quick coevolution. The key is to avoid unnecessary coupling between C2 applications and the common infrastructure; otherwise, changes needed locally require change throughout the enterprise, which is difficult.

The *C2 Enterprise Reference Architecture (C2ERA)* is an Air Force construct which attempts to find the best tradeoff between seamless integration and flexibility [2]. The C2ERA is mandated for use in all C2 information systems [3]. It is a technical concept of operations for C2 enterprise integration. It tells developers what they should to today to build C2 systems that will fit into the net-centric world tomorrow. The key elements of the C2ERA are shown in the following diagram.

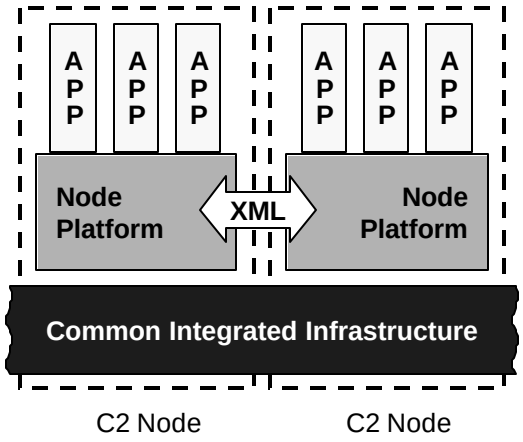

Figure 1: The C2 Enterprise Reference Architecture

- Mission applications (software which directly implements the capabilities desired by the users) are separated from infrastructure services. The infrastructure is further divided into the *Common Integrated Infrastructure,* which is the same across the enterprise, and into *node platforms,* which may vary.
- Mission applications which support related operational activities are gathered together and managed as a *C2 Node.* The C2 Node Manager is responsible for delivering and sustaining integrated capability as a weapon system to operational users.
- Information exchange between C2 Nodes are implemented using least-common-denominator, XML-based data exchanges. The goal is to preserve the independence of the C2 Node

implementations, so that a change to one C2 Node will not necessarily force all others to change.

At the DoD level we see a similar construct known variously as *Net-Centric Enterprise Services (NCES)* or *GIG Enterprise Services (GES)*. Like the C2ERA, NCES also identifies a set of core enterprise services, which are separated from mission applications and "edge" user clients, all connected by a seamless communications network backbone. These core services are depicted in the following diagram. We expect that C2ERA and NCES will be completely compatible.

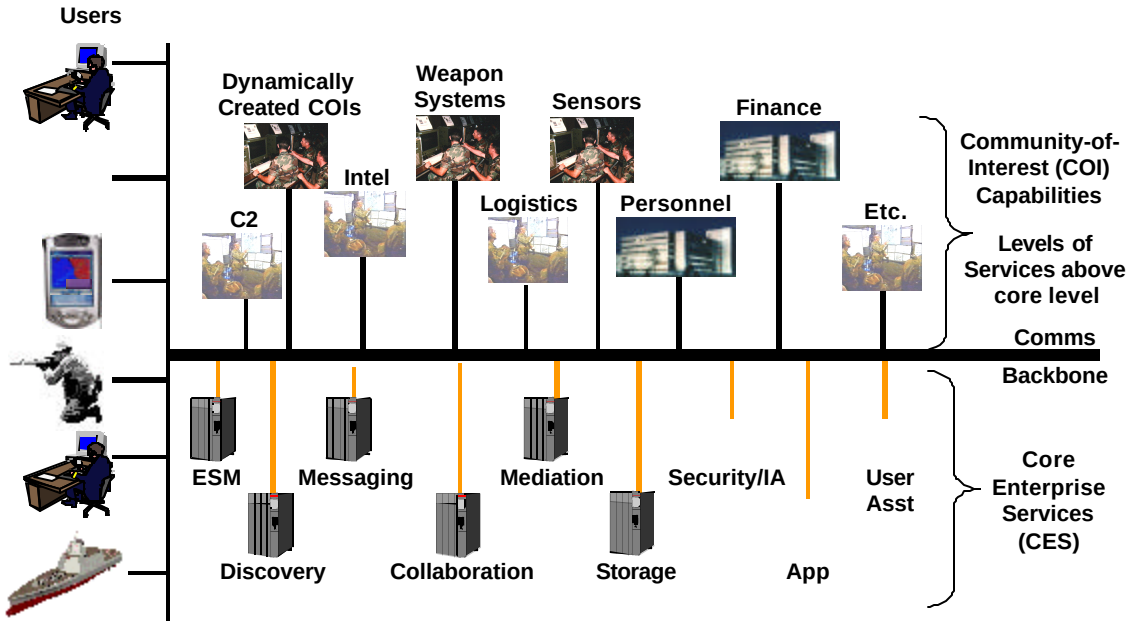

Figure 2: Net-Centric Enterprise Service Taxonomy

3.2 Information Object Publish/Subscribe Services (implied by #2, #6)

We must have great flexibility in the information exchange arrangements between network participants. Pairwise connections, tediously arranged by people, will no longer suffice, because there will be far too many participants to consider pair-by-pair. Instead, we will need a publish/subscribe (or post/pull) architecture, in which producers publish their data, making it available on the network, while consumers describe their information requirements, pulling data from the network. In the middle is an *information object service* layer that matches up the descriptions of what producers have posted with descriptions of what consumers need, and delivers data as required. This information object service architecture is shown in figure 3 below.

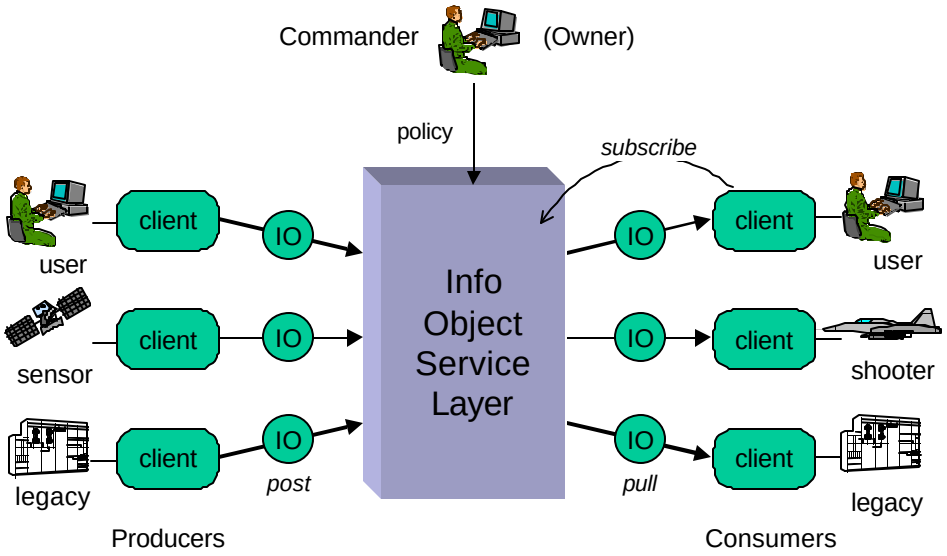

Figure 3: Information Object Services

Information object services are at the heart of the *Joint Battlespace Infosphere (JBI)* concept, developed by the Air Force Scientific Advisory Board [4,5]. The Air Force Research Laboratory is presently researching technology for implementing the information object service layer [6]. We can see the beginning steps towards the full IOS layer in the discovery service contained in NCES, and the new DoD Discovery Metadata Standard (DDMS) being developed by the DoD CIO [7].

3.3 Dissemination Optimization (implied by #3)

Bandwidth limits in some parts of the enterprise mean that we cannot always afford the overhead of a separate network transmission for each information object delivered. Instead, we will need to employ caching, multicasting, and other mechanisms to optimize information dissemination. Information object services necessarily collect a great deal of data about the information that participants have and need. We can exploit that data for this purpose.

3.4 Subject-Area Vocabularies for Communities of Interest (implied by #2, #6)

Information object services cannot work unless producers and consumers have common vocabularies to describe the information that producers have and consumers need. For example, if the producer describes an information object as pertaining to "military ships", while the consumer asks for objects about "naval vessels", the infrastructure will not return a match.[1] These common vocabularies are necessary for users, who have to understand how to look for the data they want to pull, and what that data means when they get it. They are also necessary for system builders, who have to understand what the data means so their software can use it.

[1] The infrastructure *could* return a match if it is armed with a thesarus telling it that the two phrases have the same meaning. This does not eliminiate the vocabulary agreement problem; it just moves the problem to a different place.

We say "common vocabularies" because it is impossible to settle on a single comprehensive vocabulary within an enterprise on the DoD scale. (The attempt to produce a universal data model is one of the known shortcomings of the now-defunct DoD data administration program [8].) Instead we must form several (overlapping) information *communities of interest* (COIs), each with its own subject-area vocabulary. The COI concept is an important part of the new DoD data strategy [9].

Producing a COI's vocabulary is a knowledge management problem: the vocabulary is the body of knowledge, the members of the COI are the people who need to learn it. We offer some suggestions on effective ways to develop COI vocabularies in [10].

3.5 Operational Architecture for Directing Co-Evolution (implied by #5, #6)

Architecture, especially operational architecture, helps the leadership direct the co-evolution of technology, doctrine, and organization. The DoD Architecture Framework specifies the form of certain architecture products in order to ensure that different architecture descriptions can be compared and related [11]. These architecture products include descriptions of information flows. The people who construct and use the architecture products must have a common vocabulary for describing the information in these flows. Furthermore, they need to employ the same common vocabularies used by system builders and users. In this way, architecture descriptions become the linkage between the mission capabilities we want, the systems we build, and the users who employ them. Architecture then becomes the lubrication that makes the co-evolution process turn more quickly, maximizing our advantage.

3.6 Intelligent System Degradation (implied by #3, #4)

We want our C2 information systems to degrade gracefully as their component elements fail. Among other things, this entails the ability to suspend less important information flows so that more important flows still occur. Of course, importance is a matter of doctrine and commander's intent, and these are described in operational terms. We want the commander's priorities – expressed in the terms he uses – to be automatically enforced by individual systems and the communications network. The linkage between operational architecture, systems, and data providers (based on COI common vocabularies) can make this possible, if we capture this linkage in formal, machine-processable terms.

3.7 Information Preplanning (implied by #5)

The information object service approach gives us the flexibility to quickly arrange new, unanticipated information flows. However, many information flows will be known in advance, discovered through operational architecture analysis. We will want to perform deliberate information preplanning in addition to arranging ad hoc information flows. Every mission capability depends on certain essential resources: people, material, facilities. Information must be treated as another mission-essential resource. We will need to plan for its availability.

3.8 Accountable Data Owners (implied by #5)

Plans for information must eventually be grounded in known, identified data owners. We must assign authority and responsibility for creating and maintaining data. This will require something of a culture change: there must be real accountability for ensuring that the right information is available and in fact delivered to the right people... in the same way that today there is real accountability for ensuring the availability of people, material, and facilities. Some thoughts on the responsibilities that are entailed by data ownership are available in [12].

3.9 Need-To-Hide, Not Need-To-Know (implied by #2, #4, #6)

Access to information, security certification and accreditation cannot be based on pairwise end-to-end assurances. These will be impossible for precisely the same reasons we must switch to an information pub/sub architecture: too many participants to consider pair-by-pair. We must instead rely on an information publish/subscribe infrastructure that enforces policy constraints on information flows.

4. Current Problems and Opportunities

Starting with six predictions about the future NCW environment, we have derived nine information technology preconditions for building the future net-centric C2 information systems. Of those nine, which demand the most immediate attention?

- *Enterprise services:* These are well in hand, including the communications network layer. We observe constant progress in stitching different network substrates into a seamless internetwork. We expect NCES to become a funded program in FY04. While there is plenty of hard work to be done, we believe this item is on track for success.

- *Information object publish/subscribe:* The need for these services is widely accepted; however, this acceptance is technically shallow. One thing needed soon is a consensus reference architecture for the information object service layer. This would help to separate concerns, expose requirements for other services and standards, and direct research and experimentation. People could talk about "publish and subscribe" with some confidence that they were talking about the same thing.

- *Community of Interest vocabularies:* The DoD data strategy team is very active, and understands the need for COIs, but is just now coming to grips with the problems of developing the COI subject-area vocabularies. Some of the problems are not well understood. Some of the solutions are still in the research laboratories. We expect good progress, but slower progress than with enterprise services overall.

- *Accountable data owners:* This is outside the scope of the current DoD data strategy. It is a part of the Air Force strategy. We believe that progress will be difficult, because this is largely not a problem of technology. The subject will require more attention than it receives at present.

5. Conclusion

Several important problems of information technology, architecture, and management must be solved as we build the net-centric C2 information systems of the future. We must build C2 information systems that work together and are easy to change. We need architecture descriptions that are useful for directing coevolution and also for understanding and controlling the collection of C2 information systems. We need information management procedures (and the supporting infrastructure) to ensure that the right information is available for the right decider to make the right decision. This paper describes the problems, shows how they are related to each other and to the NCW future, and identifies those problems that at present are most in need of attention.

References

[1] D. Alberts, J. Gartstka, F. Stein, *Network Centric Warfare, 2nd Edition*, August 1999. http://www.dodccrp.org/Publications/zip/ncw_2nd.exe

[2] *Air Force C2 Enterprise Technical Reference Architecture, Version 3.0*, December 2002.

[3] *C4ISR Enterprise Directive 008, Technical Architecture for C4ISR Enterprise Integration,* January 2003.

[4] Air Force Scientific Advisory Board, *Information Management to Support the Warrior*, Dec. 1998. http://www.sab.hq.af.mil/archives/recommend/index.htm

[5] Air Force Scientific Advisory Board, *Building the Joint Battlespace Infosphere,* Dec. 1999. http://www.sab.hq.af.mil/archives/recommend/index.htm

[6] Air Force Research Laboratory, *Joint Battlespace Infosphere: Mercury Project,* March 2003. http://www.rl.af.mil/programs/jbi/mercury.cfm

[7] DoD Chief Information Officer, *Department of Defense Net-Centric Data Management Strategy*, March 2003.

[8] S. Renner, *Improving 8320.1 Data Administration*, Federal Database Colloquium '98, San Diego, September 1998.

[9] DoD Deputy Chief Information Officer, *Department of Defense Discovery Metadata Standard (DDMS), Review Version 1.0.* April 2003.

[10] S. Renner, *A "Community of Interest" Approach to Data Interoperability*, Federal Database Colloquium '01, San Diego, September 2001.

[11] *DoD Architecture Framework, Version 1.0,* Jan. 2003. http://flrc.mitre.org/dodfw/

[12] AF-CIO Chief Architects Office, *Discussion Paper DP-009: Data/Information Management*, October 2002.