



INSTITUTE FOR DEFENSE ANALYSES

DoD Net-Centric Services Strategy Implementation in the C2 Domain

P.J. Walsh, Project Leader

February 2010

Approved for public release;
distribution unlimited.

IDA Paper P-4549

Log: H 10-000075



The Institute for Defense Analyses is a non-profit corporation that administers three federally funded research and development centers to provide objective analyses of national security issues, particularly those requiring scientific and technical expertise, and conduct related research on other national challenges.

About this Publication

This work was conducted by the Institute for Defense Analyses (IDA) under contract DASW01-04-C-0003, Task BC-1-2526 for the Office of the Assistant Secretary of Defense for Networks and Information Integration [ASD (NII)]. The views, opinions, and findings should not be construed as representing the official position of either the Department of Defense or the sponsoring organization.

Copyright Notice

© 2010 Institute for Defense Analyses, 4850 Mark Center Drive, Alexandria, Virginia 22311-1882 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (NOV 95).

UNCLASSIFIED

INSTITUTE FOR DEFENSE ANALYSES

IDA Paper P-4549

DoD Net-Centric Services Strategy Implementation in the C2 Domain

P.J. Walsh, Project Leader

S.O. Davis

M. Kiefer

K.A. Morrison

J.W. Pipher

H.G. Potrykus

UNCLASSIFIED

UNCLASSIFIED

PREFACE

This study describes an approach for implementing the DoD Net-Centric Services Strategy (NCSS)¹ in the command and control (C2) domain. It complements a study also conducted by IDA in 2008 on implementing the Department's Net-Centric Data Strategy (NCDS) in the C2 domain.² It is anticipated that the results of this study will be used to develop guidance for implementing services in the C2 community and, if appropriate, will also be used as a model for implementing services across the entire DoD information enterprise.

The recommended approach for implementing C2 information support capabilities as services was formulated by developing a tiered organizational structure for implementing and managing C2 information services as part of an evolving Service-Oriented Enterprise (SOE); a C2 services CONOPS that ties together C2 service tiers, key implementation roles and actors; and a high-level governance concept. Additionally, a C2 Information Sharing Framework (C2ISF) is developed to highlight the design-time artifacts and run-time infrastructure services needed to accomplish NCSS goals and objectives. The recommended C2 services implementation approach emphasizes the importance of enabling and encouraging edge-user agility and innovation in developing and improving C2 services while implementing the SOE.

The study team consisted of Mr. Philip J. Walsh (Project Leader), Mr. Stanley O. Davis, Ms. MaryAnn Kiefer, Dr. Kyle A. Morrison, Mr. James Pipher, and Dr. Henry G. Potrykus. The team would like to thank Dr. David R. Graham, Dr. Richard J. Ivanetich, Dr. Davy Y. Lo, Dr. Margaret E. Myers, Dr. Douglas G. Shiels, Dr. Steve Warner, Ms. Patricia G. Phillips, and Mrs. Zelma B. Cameron from IDA for their critical reviews, helpful suggestions, and timely assistance.

¹ Department of Defense "Net-Centric Services Strategy," DoD CIO, May 4, 2007.

² IDA Paper P-4404, *Independent Assessment Team Report on C2 Data*, November 2008.

UNCLASSIFIED

(This page is intentionally blank.)

UNCLASSIFIED

UNCLASSIFIED

CONTENTS

EXECUTIVE SUMMARY	ES-1
A. Motivation.....	ES-1
B. Case Study Analyses.....	ES-2
C. Implementation Approach	ES-3
D. Conclusions and Recommendations	ES-6
 I. INTRODUCTION	 1
A. Background.....	1
B. Fundamental Definitions.....	2
C. Purpose and Objectives.....	3
D. Scope	3
E. Assumptions.....	4
F. Analytic Approach.....	4
G. Organization of the Report.....	5
 II. IMPLEMENTING A SERVICE-ORIENTED ENTERPRISE (SOE)	 7
A. Strategic Vision	8
1. Operational Component.....	8
2. Technical Component	10
B. Characterizing the Current Environment	12
 III. C2 INFORMATION SHARING FRAMEWORK (C2ISF)	 19
A. C2ISF Description	20
1. Design-Time Infrastructure.....	22
2. Run-Time Infrastructure	26
3. Relationship Between C2ISF and Services and Data Strategies	30

UNCLASSIFIED

B. Implementation and Management of the C2ISF	31
1. C2 Services Description Templates	32
2. C2ISF Run-Time Infrastructure Rules and Protocols	32
IV. C2 SERVICES CONOPS	35
A. Introduction the CONOPS	35
B. Overview of C2 IN AN SOE	36
C. C2 Service Tiers.....	37
D. Service Life Cycle.....	42
1. ITIL/COBIT	42
2. C2 Service Tier Characteristics.....	44
3. Roles in the Service Life Cycle	47
4. An Example of CONOPS Roles for C2 Services	50
5. Responsibilities in the Service Life Cycle.....	52
E. GOVERNANCE IN C2 SERVICES DOMAIN.....	55
1. C2 Services Portfolios.....	55
2. Portfolio Management Responsibilities.....	56
3. Proposed C2 Services Portfolio Governance Concept.....	57
F. GOVERNANCE ISSUES AND CONSIDERATIONS.....	60
V. CONCLUSIONS AND RECOMMENDATIONS	63
A. CONCLUSIONS.....	63
B. RECOMMENDATIONS	64
1. Guidance	64
2. Implementation Plans.....	64
3. Near-Term Implementation Actions.....	64

Appendixes

- A. C2 Services Case Studies
- B. C2 Information Sharing Framework (C2ISF)
- C. Glossary
- D. Figures and Tables

UNCLASSIFIED

EXECUTIVE SUMMARY

UNCLASSIFIED

EXECUTIVE SUMMARY

A. MOTIVATION

Since 2003, the Office of the Secretary of Defense (OSD) has provided strategic direction and policy guidance to the Department of Defense (DoD) that was intended to improve information access and sharing across the DoD enterprise over time. Much of this direction and guidance addressed the goal of achieving net-enabled¹ command and control (C2) capabilities. Two key issuances in this regard are the 2003 Net-Centric Data Strategy (NCDS) and the 2007 Net-Centric Services Strategy (NCSS).²

Implementing both the NCDS and NCSS has proven to be complex and challenging. In 2008, the OASD(NII)/DoD CIO sponsored a study to determine an approach for accelerating the implementation of the NCDS in the C2 domain. That study³ resulted in renewed efforts to hasten the development of C2 data-sharing mechanisms and standards (e.g., development of the C2 Core as an extension of the Universal Core) as well as the identification of sources of authoritative C2 data and the exposure of that data by those authoritative sources. The C2 data study also highlighted the need for a companion effort to address NCSS implementation in the C2 domain.

This study is that associated work. It provides a comprehensive approach for implementing C2 information support capabilities as *information services*,⁴ and thereby

¹ In this context, the term “net-enabled” is used in its most generic sense—i.e., C2 operations facilitated through the use of information technology (IT) systems interconnected via a communication network or network of networks.

² DoD Chief Information Officer Memorandum, “DoD Net-Centric Data Strategy,” May 9, 2003; Department of Defense “Net-Centric Services Strategy,” DoD CIO, May 4, 2007.

³ IDA Paper P-4404, *Independent Assessment Team Report on C2 Data*, November 2008.

⁴ C2 *information services* facilitate the provision of C2 data or data-processing functionality to C2 practitioners (called customers) without the customers owning or managing the mechanism for providing that support.

can potentially help achieve a key Department goal of migrating C2 capabilities from the current system-based implementation construct to a Service-Oriented Enterprise (SOE).⁵

Implementing C2 information support capabilities as services fundamentally involves three things: (1) continuously working to understand and answer the C2 customer's IT needs at all military command echelons; (2) creating an IT-enabled development, acquisition, and operation life-cycle management environment that can detect and rapidly respond to those needs; and (3) operating a 24/7 globally available IT infrastructure that allows C2 customers to discover, access, use, and rely on information support capabilities that are implemented as services.

B. CASE STUDY ANALYSES

To understand how information support capabilities are being implemented and used by operating forces today (particularly in Iraq and Afghanistan), the study examined four cases where operating forces are effectively leveraging IT capabilities to support C2.⁶ These four cases involve creating or improving operational processes and related software products associated with the following C2 mission areas:

- Specialized Tactical Ground Situational Awareness (SA) Services
- Command-Level Situational Awareness and Collaboration Services
- Air Operations Tasking and Control Services
- Force Deployment Planning and Execution Management Services.

These four case studies clearly indicate that the current operational environment is experiencing profound changes in how capabilities are obtained—from technical, procedural, and, importantly, organizational standpoints. Key observations gleaned from these case studies helped drive this study's results and include the following:

- C2 service implementation is currently taking place at three important levels: local C2 nodes, common clusters of C2 services, and enterprise infrastructure.

⁵ A *Service-Oriented Enterprise (SOE)* is an enterprise that combines a services-focused way of doing business with the latest technology in an operational culture where participating entities include both service providers and service consumers. This implies a broader and less technically prescriptive approach to providing and consuming services than is generally implied by usage of the term service-oriented architecture (SOA).

⁶ These four cases are discussed in detail in Section II.B, pages 12-17 and in Appendix A.

UNCLASSIFIED

- It is necessary to support and build on what is successfully deployed and is already operational.
- Final determinations on what capabilities work best and where to invest in improvements must be vested in the operational chain of command.
- IT capabilities generally require rapid improvement after initial fielding. This requires authoritative decision making to identify technical options, to apply funding, and to engage engineering/training/logistics support in both rear and forward areas in order to implement enhancements.
- Significant local innovation will occur and successfully deliver valuable capabilities to warfighters regardless of acquisition rules. In general, highly formalized processes for acquiring or improving C2 IT capabilities are not sufficiently responsive to emergent warfighter needs.
- Capabilities developed over decades through formal acquisition programs, once fielded, are being significantly enhanced by Web technologies.
- Services that emerge through local innovation can find a programmatic home, either by being adopted by an existing Program of Record (PoR) or through the creation of a new PoR by governing authorities at the C2 domain or enterprise level.

C. IMPLEMENTATION APPROACH

By using existing C2 policy and implementation guidance and observations from the four case studies as context, a C2 services implementation approach was developed. As illustrated in Figure ES-1 and amplified in the main body of this report, the resulting implementation approach consists of four major components:

1. A C2 Information Sharing Framework (C2ISF) to help the Department understand and implement an SOE that features run-time infrastructure services necessary to enable the creation, discovery, use, and management of C2 mission-related services. See Figure 3 and the associated discussion on pages 20-21.
2. A three-tiered organizational structure for implementing and managing C2 services over their life cycle, as part of an evolving SOE. See Figures 8 and 9 and the associated discussion on pages 37-42.
3. A C2 services concept of operations (CONOPS) that ties together C2 service tiers, key implementation roles and actors, and a top-level governance

UNCLASSIFIED

concept.⁷ See Figures 11, 12, and 13 and the associated discussion on pages 48-54.

4. A recommended C2 services governance concept. See Figure 16 and the associated discussion on pages 57-61.

The CONOPS developed in this study establishes a foundation for managing and governing the evolution of the SOE for C2 services. Consequently, it focused mostly on organization and responsibilities. However, it is not a comprehensive treatment of all issues and considerations that must be addressed to realize an SOE. Ultimately, governance of C2 IT and the SOE will have to also address issues such as:

- How can edge innovation be incentivized?
- How can operational situational awareness measures of service operations be better linked to resourcing?
- How can Departmental processes be modified to promote inter-dependency without undue burden?
- How can risk be effectively managed in a multi-tiered SOE?
- Accordingly, our study recommendations represent a first step in implementing a C2 services CONOPS and SOE.

⁷ The C2 services CONOPS is based on an adaption of Information Technology Infrastructure Library/Control Objectives for Information and related Technology (ITIL/COBIT) management concepts for commercial IT services. See pages 38-40 for a detailed discussion of these concepts.




Figure ES-1. Major Components of Proposed C2 Services Implementation Approach

UNCLASSIFIED

D. CONCLUSIONS AND RECOMMENDATIONS

The case study observations confirm that Internet-style services will be a progressively more vital means of providing C2 information support. These observations lead to the following key conclusions that guided the study's recommendations:

- Achieving an SOE for C2 requires long-term planning and effective oversight while supporting agility and innovation in the operational community.
 - Planning, oversight, and operations are dependent on having situational awareness information about that status of services and information products on the networks with focus on collecting and publishing usage and performance metrics.
- Implementation of C2 community services within an evolving SOE requires a time-phased “start-up” plan. Critical elements include:
 - Designation of categories or “tiers” of services and appropriate governance authorities for each category/portfolio
 - Development of portfolio-level governance processes for the identification, acquisition, and life cycle management of C2 services
 - Assignment of service management roles and responsibilities with emphasis on promoting edge-user innovation
 - Provision of CONOPS and implementation guidance to provide unity of effort for identifying, implementing, and managing information support as services within an evolving DoD Services-Oriented Enterprise
 - A commitment and plan of action to develop appropriate infrastructure services in a sequence and on a timeline that supports the simultaneous development of mission-oriented services.
- An aggressive program of work to identify and develop technical approaches and relevant standards is needed to implement a federated SOE and associated infrastructure services.

Based on these conclusions, it is recommended that:

- The DoD CIO, issue guidance to:
 - Clarify and institutionalize service-related terminology (to include appropriate ITIL/COBIT roles and definitions)
 - Adopt and advocate for a layered model for services and SOE governance that addresses agility and stability needs
 - Establish the necessary federation approaches and standards to support the evolution of a Department-wide SOE.

UNCLASSIFIED

- The DoD CIO, in coordination with U.S. Cyber Command as the NetOps mission owner, develop guidance for measuring and publishing service implementation and usage metrics on DoD networks.
- The DASD, C3S&S, in coordination with DASD, IMIT/Deputy DoD CIO, DoD Components and the COCOMs, initiate action to implement C2 services and the tiered C2 services structure and governance processes discussed herein.
- The DASD, C3S&S, in coordination with DASD, IMIT/Deputy DoD CIO, emphasize and build on specific near-term implementation actions at the C2 local, C2 common, and enterprise levels.⁸ Also consider:
 - Using the emerging “JC2 initiative” (i.e., NECC replacement/follow-on) as a pathfinding effort for C2 services-based development/acquisition
 - Codifying relevant C2 service and SOE implementation activities in the DoD C2 implementation Plan.

⁸ Recommended specific near-term implementations actions are listed on pages 61-63.

UNCLASSIFIED

(This page is intentionally blank.)

ES-8

UNCLASSIFIED

I. INTRODUCTION

A. BACKGROUND

Over the years, the Department of Defense (DoD) has issued strategic direction and policy guidance focused on enabling the warfighter through improved information access and sharing across the DoD enterprise. Key Department issuances in this area include:

- DoD Net-Centric Data Strategy (NCDS), May 2003
- Data Sharing in a Net-Centric Department of Defense, Dec 2004
- Guidance for Implementing Net-Centric Data Sharing, Apr 2006
- DoD Net-Centric Services Strategy (NCSS), May 2007
- DoD Command and Control (C2) Strategic Plan Version 1.0, Dec 2008
- Interim Guidance to Implement NCDS in the C2 Portfolio, Mar 2009
- DoD C2 Implementation Plan Version 1.0, Oct 2009.

Of particular relevance to this study is the NCSS issued in May 2007. Figure 1 depicts the highlights of this strategy.

The NCSS study, as a potential DoD pathfinder in implementing a Department-wide Service-Oriented Enterprise (SOE), is intended to build on the value propositions, goals, and key actions shown in Figure 1.

Although guidance related to implementing both the data and services strategies has been in place for some time, there has been limited measurable implementation progress in realizing the potential value of net-centric data and services approaches. Observations suggest that the Department still has inconsistent and incompatible technical implementation approaches, inconsistent and confusing terminology, inconsistent and under-defined concepts of employment and operations, undefined and under-defined implementation roles and responsibilities, and limited ability to execute effective governance and management due to very poor visibility into the evolving operational baseline and the effects of any given technology insertion. To the extent

possible, this study attempts to suggest a way ahead (at least in the C2 domain) that promotes coherence and uniformity in this difficult area.




Figure 1. Net-Centric Services Strategy (NCSS) Highlights

B. FUNDAMENTAL DEFINITIONS

The following fundamental definitions provide a useful starting point for discussing the implementation of capabilities as services:

- A service is a means of delivering value to customers by facilitating outcomes customers want to achieve without incurring the ownership of specific costs and risks.¹ Generally, information services provide access to data, computational or transactional functions, or management or orchestration

¹ This definition of *service* is taken from the Information Technology Infrastructure Library (ITIL) set of concepts and policies for managing information technology (IT) infrastructure, development, and operations as published by the UK Office of Government Commerce (OGC). An additional source of IT management-related definitions and concepts is the Control Objectives for Information and related Technology (COBIT), which is a framework for IT management that was created and is maintained by the Information Systems Audit and Control Association (ISACA) and the IT Governance Institute (ITGI). See Section IV.D for a more detailed discussion of how this study has adopted a combination of ITIL and COBIT terminology and concepts to address the life cycle management of C2 services in the context of a C2 Services Concept of Operations (CONOPS).

UNCLASSIFIED

functions without the customer owning the mechanism for providing that support.

- In this NCSS study, we are addressing a subset of services called C2 information services (or C2 services). C2 services facilitate access to and the provision of C2 data or functionality to support C2 operational processes. C2 operational processes are also supported by services from other capability areas (e.g., logistics, intelligence), as well as enterprise-wide services, but such services will not be considered “C2 services.”
- Service-Oriented Enterprise (SOE) is an enterprise that combines a services-focused way of doing business with the latest technology in an operational culture where participating entities include both service providers and service consumers. This implies a broader and less technically prescriptive approach to providing and consuming services than is generally implied by usage of the term service-oriented architecture (SOA).

C. PURPOSE AND OBJECTIVES

The purpose of this study is to develop and describe an approach for implementing C2 information support capabilities as services. It is anticipated that the results of this study will be used to develop guidance for implementing services in the C2 community and, if appropriate, to also be used as a model for implementing services across the entire DoD information enterprise. In addition, the C2 services implementation approach is intended to satisfy the following objectives:

- Provide unity of effort for identifying, implementing, and managing required information support as services within an evolving DoD SOE
- Leverage and encourage ongoing efforts at all echelons within the operating forces to identify, implement, and continuously improve C2 information services
- Be compliant with and expand upon DoD’s May 2007 NCSS.

D. SCOPE

This study will examine the purposes, methods, and means of providing responsive IT for C2 of joint/interagency/multinational forces operating in current and anticipated future security environments. It focuses on services available on defense networks (in contrast to reliance on organic, stovepiped systems) at all military command echelons involved in C2, from local tactical command levels to global strategic command levels. Thus, the ultimate objective is to help facilitate the evolution of an SOE that can

UNCLASSIFIED

enable the provision of information services to commanders and their staffs throughout the C2 mission space.

E. ASSUMPTIONS

The following assumptions have been made to realistically constrain the scope of this study and provide a basis for formulating study recommendations:

- DoD will continue to work toward achieving net-centric, services-based information capabilities.
- The necessary scaling for a DoD SOE for C2 operations will be achievable. The complexity of operations, widespread interoperability, and scale needed for assured service orientation will be addressed through changes in technology, processes, and practices.
- Existing DoD capability needs, acquisition, and resource processes will be modified/adapted as necessary to achieve this.
- Where operationally viable and technically feasible, services will be made accessible to the widest possible user base.
- SOE implementation approaches will be constrained by expected limitations in transport capabilities—particularly at the “tactical edge.”
- C2 SOE implementation approaches and methods identified in this study may be considered for extension to other domains within the DoD enterprise.
- Capability portfolio management will continue to be used as a principal organizing construct for managing capabilities in the DoD.
- C2 services are part of the C2 capability portfolio.

F. ANALYTIC APPROACH

For this study, it is postulated that an approach for implementing C2 information services can be formulated by addressing the following topic areas in order:

- Identifying and defining services-related terminology and an understanding of DoD’s intended evolution to an SOE
- Developing a C2 Information Sharing Framework (C2ISF) to help the Department understand C2 services’ implementation and use, and in a run-time environment, can enable service discovery, access, and use. This includes defining Service Description Templates that facilitate service discovery, access, and use.

UNCLASSIFIED

- Defining a tiered framework for implementing and managing C2 services as part of an evolving SOE. This includes developing a C2 services CONOPS that ties together C2 service tiers, key implementation roles and actors, and a top-level governance construct.

The recommended C2 services implementation approach also:

- Addresses the importance of enabling and encouraging edge-user agility and innovation in developing and improving C2 services while implementing the SOE
- Includes near-term priorities
- Includes C2 infrastructure- and enterprise-related activities.

In developing the approach for implementing C2 capabilities as services, the study team placed special emphasis on the governance aspects of an SOE and a detailed examination of incipient “real world” case studies evolving in the Operation Iraqi Freedom (OIF) and Operation Enduring Freedom (OEF) wartime environments. In assessing the four wartime case studies, we paid particular attention to how organizational arrangements supported required activities for life cycle management of C2 services. These case studies are considered representative of the various ways that C2 prototype services are being implemented and used in the operating forces today. The four case studies are:

- Specialized Tactical Ground Situational Awareness (SA) Services
- Command-Level Situational Awareness and Collaboration Services
- Air Operations Tasking and Control Services
- Force Deployment Planning and Execution Management Services.

These four case studies, discussed in detail in Section II.B and Appendix A, focused on how capabilities are being used and evolved in Iraq and Afghanistan today. They clearly indicate that the current operational environment is experiencing profound changes in how capabilities are implemented—from technical, procedural, and, importantly, organizational standpoints.

G. ORGANIZATION OF THE REPORT

The remainder of the report includes our analyses and more detailed appendixes. The detailed analyses, which form the basis for the study results, are organized into four major topic areas:

UNCLASSIFIED

- Implementing an SOE
- C2ISF
- C2 Services CONOPS
- Results.

Finally, the first two appendixes provide additional detail on the following topics:

- C2 Service Case Study Analysis
- C2ISF.

The third appendix is a glossary containing definition of terms and acronyms and the fourth appendix contains the lists of figures and tables.

II. IMPLEMENTING A SERVICE-ORIENTED ENTERPRISE (SOE)

A major IT challenge facing the Department today is to manage the transition of stovepiped C2 systems to an interdependent, services-based, net-enabled enterprise. This challenge is both operational and technical.

Implementing both the NCDS and NCSS has proven to be a complex and challenging undertaking. In 2008, the Office of the Assistant Secretary of Defense for Networks and Information Integration [OASD(NII)]/DoD Chief Information Officer (CIO) sponsored a study to determine an approach for accelerating the implementation of the NCDS in the C2 domain. That study¹ resulted in renewed efforts to hasten the development of C2 data sharing mechanisms and standards (e.g., development of the C2 Core as an extension of the Universal Core) as well as the identification of sources of authoritative C2 data and the exposure of that data by those authoritative sources.

This study is a companion effort to the 2008 C2 data study—and addresses the implementation of the NCSS in the C2 domain. It provides a comprehensive approach for implementing C2 information support capabilities as *information services*,² and thereby can potentially help achieve a key Department goal of migrating C2 capabilities from the current system-based implementation construct to an SOE.

¹ IDA Paper P-4404, *Independent Assessment Team Report on C2 Data*, November 2008.

² C2 *information services* facilitate the provision of C2 data or data processing functionality to C2 practitioners (called customers) without the customers owning or managing the mechanism for providing that support.

A. STRATEGIC VISION

SOEs³ embody policies, procedures, organizational arrangements, and technologies that enable clusters of services to emerge and evolve on a network. A service is a means of delivering value to users by facilitating outcomes they want to achieve without assuming the costs and risks of actually acquiring the capability themselves. There are several categories of service; however, this study is focused specifically on information services. Information services are a means of facilitating mission outcome by providing information or data processing support as opposed to delivering materiel, for example. Generally, information services offer access to data or computational, transactional, management, or orchestration functions.

1. Operational Component

A DoD goal is to transform C2 information support capabilities into an interdependent, leader-centric and net-enabled⁴ IT portfolio. This means establishing a rich information-sharing environment in which trusted content and functionality are provided through assured services. This net-enabled vision requires migrating from the current, relatively closed, system-based implementation construct toward a far more open, service-oriented model similar to current Internet capabilities. For C2, the SOE must consist of numerous globally distributed services that are easily discovered, accessible, and operating 24/7 to provide the Department's decision-makers with the information they want and the functionality they need, when they want them, and wherever they are located (including mobile access). Service providers and consumers must also be able to rapidly evolve services by entering into cooperative arrangements with other providers and consumers to create new value-added capabilities in response to emergent customer demands.

³ The term "service-oriented enterprise (SOE)" is intended to describe an enterprise that combines a services-focused way of doing business with the latest technology in an operational culture where participating entities are both *service providers* and *service consumers*. This term implies a broader approach to providing and using services than the term "service-oriented architecture (SOA)," which implies the implementation of particular architectural or technical constructs. See Appendix D for definitions of both these terms.

⁴ In this context, the term "net-enabled" is used in its most generic sense—i.e., C2 operations facilitated through the use of IT systems interconnected via a communication network or network of networks.

UNCLASSIFIED

For C2, services will have to support a wide range of users from the local commander looking for information within his theater to regional and national commanders who require particular processes and functions to be executed. C2 services will consist of a variety of types: those that acquire and consume information, generate new information, and provide access to information; and those that provide value-added functionality for use by others. Generally, these services will be tailored to meet the needs of a primary group of known users, but increasingly they will be open to unanticipated users to be leveraged in new and unplanned ways. The freedom to expand service usage makes an SOE agile and adaptive, able to quickly and relatively easily morph existing capabilities to address new demands rather than developing capabilities from the ground up.

SOE agility implies that the nature and frequency of specific service interactions, the tightness of resulting federations, and distribution of content across the enterprise cannot be achieved with any precision in advance. SOE implementation and operation must be driven by a response to users' needs based on empirical usage data from network instrumentation and anecdotal feedback. Whether based on empirical data, user feedback, or validated needs, each service must be designed and developed around the basic mission and business processes and information needs of one or more organizations. Services are then combined over networks into even larger or more complex, loosely coupled processes or capabilities. Hence, services must have the potential to discover and communicate with each other.

Actors distributed across the enterprise, C2 community, and other mission communities are responsible for planning, designing, provisioning, and operating the SOE. The number of actors involved, the complexity of roles, and the need for flexibility requires a balance between rigid, top-down control and loosely coordinated user-driven actions. The concepts for operations of the SOE must account for this dichotomy.

DoD's SOE is not just about engineers, Web services, or technology. It entails a broad approach that exploits the service paradigm to engender agility in the Department's intra- and inter-organizational processes and to drive responsive IT. The SOE implementation and operation approach involves knowing emerging user needs; understanding what is happening on the networks; leveraging commercial technology; continually monitoring, analyzing, reporting, and formulating responses; and adopting

best practices in IT services management. The *C2 Services CONOPS* in Section IV describes an approach to managing and governing the SOE for C2.

2. Technical Component

The DoD C2 Strategic Plan and DoD C2 Implementation Plan⁵ characterize future net-enabled C2. Figure 2 depicts a framework by which C2 capabilities are provided as services delivering both data and functionality throughout the SOE. This SOE is capable of information sharing via Web-based services with other U.S. Government organizations and coalition partners.




Figure 2. Conceptual C2 Services Framework

The SOE vision for C2 shown in Figure 2 includes the support environment (lower half), networks, and user-invoked mission capabilities (upper half). The SOE is interdependent with C2 processes and procedures. Importantly, C2 capabilities are supported by access to data from a variety of sources including those managed by the C2 community as well as other communities (e.g., Battlespace Awareness and Logistics). In

⁵ The DoD C2 Strategic Plan, Version 1, dated December 18, 2008, and the DoD C2 Implementation Plan, Version 1.0, dated October 1, 2009.

UNCLASSIFIED

the foreseeable future, these sources will be a mix of services and systems with the former gradually coming to predominate. C2 systems, platforms, and facilities with reliable and robust access to a network will be the initial implementers of services, beginning the migration toward an SOE. However, some capabilities will remain traditional point-to-point information exchange solutions, particularly where required to support time critical sensor-to-shooter exchanges or disconnected, interrupted, and low bandwidth (DIL) operational environments.

The SOE depicted in Figure 2 requires C2 capabilities to leverage robust and persistent infrastructure services (included in the box labeled “SOA Foundation”). The infrastructure services define the environment—they enable information sharing and interoperability across missions and organizations. They are content or mission-neutral, general-purpose capabilities that provide for basic communications, collaboration, publication, discovery, security, and information and service management. These infrastructure services are not optional. An SOE can only exist when these services are operational. Some of this infrastructure must be provisioned and operated at the Department level, and communities will offer some infrastructure where specialization or localization is required. The SOE will have adaptive growth when sufficient infrastructure exists to support discovery, access, understanding, and use of services within C2 boundaries and from other communities.

Service design begins with an organization’s business processes concept (e.g., planning, tasking, monitoring the effects of operations) that identifies the critical components of the process (e.g., unit identification, task orders, unit status reports). These components define the parameters of stand-alone pieces of software instantiated as services. Services can be written to serve specific purposes (e.g., define the identity of a unit) and be shared with other programs (e.g., borrow unit identity definition from the Army and apply it to joint/coalition forces). Lessons learned from development of the services can be used to revise the business practices within communities and across the enterprise. This permits the evolutionary development of progressively larger clusters of services-based capability that can be rapidly combined and recombined to create new capabilities and processes.

In an SOE, multiple services can be combined into composite services or composite applications. The interfaces to these composite services can be exposed and invoked in the same manner as a stand-alone “atomic” service. A composite service

UNCLASSIFIED

relies on the various services it invokes; therefore, the functionality and successful operation of the composite service is dependent upon the state of those underlying services at the time of any given transaction. Many DoD services, particularly those supporting large complex processes, such as air operations management or force deployment and monitoring, appear to fit the composite service model.

The more complex processes yield clusters of composite services that support or execute specific operations. These clusters can involve complex sets of interrelations as the services are linked together loosely or more tightly. A loosely coupled approach to design, development, and implementation of services and services-based processes is possible because of Web technologies. The loosely coupled approach has proved to be the only effective way to accommodate Internet or very large intranet scales as is found in DoD. Loose coupling resists attempts to define and provision IT through top-down direction, rigid planning, tight end-to-end engineering, and extensive compliance enforcement mechanisms that are hallmarks of the typical Program of Record (PoR). However, current operational IT capabilities developed as part of existing systems and PoRs are often amenable to service evolution and can act as a nucleus for rapidly clustering services. The NCSS does not call for wholesale replacement of these systems and their operational capabilities. It advocates for existing capabilities to be adapted to a services-based paradigm, particularly those from high-value warfighter systems. Most of these services will eventually have both a human interface using standard Web capabilities (i.e., common browsers) and a machine interface leveraging Web-services technology.

B. CHARACTERIZING THE CURRENT ENVIRONMENT

In formulating an approach for evolving C2 information capabilities as services within an SOE, it is useful to understand the current environment. We examined four existing C2 capabilities (Appendix A) that are operational, currently available on SIPRNet, and amenable to service-oriented evolution. Examples of pure and purposeful service implementation remain few and far between in DoD. Accordingly, the cases selected are not paragons of service development and operations. Rather, they lead to observations concerning current behaviors in fielding, operating, and evolving C2 capabilities. The specific case studies are relevant both to the Department's overall migration toward an SOE and to the implementation of C2 services in particular:

UNCLASSIFIED

- **Specialized Tactical Ground Situational Awareness (SA) Services.** "Situational Awareness" here refers to knowledge of previous observations, lessons-learned, issues, and solutions experienced by others in similar circumstances. The IT is actually a family of Defense Advanced Research Projects Agency (DARPA)-sponsored, Web technologies that have been inserted into Iraq and Afghanistan over the past several years with positive results. The most noteworthy of these are the Tactical Ground Reporting system (TIGR) and the Combined Information Data Network Exchange (CIDNE).
- **Command Level Situation Awareness and Collaboration Services.** This case looks at Command Post of the Future (CPOF), which was originally a DARPA technology demonstration in the late 1990s that became an Army PoR in 2006. It has been tightly federated with the Army's Maneuver Control System (MCS) and consumes data from the Global Command and Control System (GCCS).
- **Air Operations Tasking and Control Services.** This case assesses key C2 business processes, organizations, and IT used by Combatant Command (COCOM) Air Component Commanders to manage air operations on a regional basis. From an IT standpoint, the principal focus of this case is on Theater Battle Management Core Systems (TBMCS) resident in globally distributed Air Operations Centers (AOCs). The TBMCS program has a decades-long developmental and operational history.
- **Force Deployment Planning and Execution Management Services.** The Joint Operational Planning and Execution System (JOPES) is the Department's principal tool for designing, monitoring the progress of, and managing force deployments for generally large-scale deployments and rotations. It is considered part of the GCCS family of systems comprising over 200 Government and commercial off-the-shelf (GOTS/COTS) capabilities. It is currently operational at the national level and in a variety of other flag-level command centers.

These four case studies clearly indicate that the current operational environment is experiencing changes in how capabilities are implemented—from technical, procedural, and organizational standpoints. A significant aspect is the migration of C2 capability down-echelon, particularly the establishment of company-level command posts (CPs). It appears that even smaller units distributed among villages will be connected to the maximum extent possible with larger, regional rapid response force HQs that are linked to higher command in a progressively expanding Web. Of note, the lower echelons in this network increasingly include other U.S. Government representatives (e.g., State,

UNCLASSIFIED

Agriculture) and multi-national DoD coalition partners. C2 collaboration across the expanding Web will be significant.

These cases suggest that much of the IT supporting current operations is being developed and provisioned from multiple sources that are essentially uncoordinated at the enterprise level and only marginally coordinated within specialized warfighting communities. The sources include PoR products (generally large system hardware and software combinations requiring long-term development); demonstrations or prototypes that use cutting-edge technology; and well-proven COTS technologies. Thus, present day IT provisioning is characterized by what might be termed “natural growth,” which, if shaped by some governance, could be far more effective and efficient. But Information Age governance requires transparency. There is little ability to detect and track the particulars of evolving operational C2 IT support much less to intervene to achieve efficiencies and synergies.

Given the present lack of insight, local commanders acting in the context of theater C2 arrangements are in the best position to understand and control what is going on in their respective operational environments. Healthy SOE-like behaviors already have a foothold in the operating forces. Cooperation, via Web-based IT, acts as a catalyst to federate a variety of supporting data sources (e.g., creating “mashups”) that bear on mission information needs. Operational forces have demonstrated the ability to quickly assemble new capabilities by using experimental products, commercially available Web technologies, and small cadres of forward-deployed supporting engineers and trainers. The cases show that with some granular knowledge of specific theater needs and conditions, new locally innovated C2 Web-enabled capabilities and services can be provisioned to meet emerging mission requirements. Once these new capabilities and services are determined to be effective, their usage can be rapidly expanded.

When local capabilities achieve rapid adoption or when local capabilities cluster to form a common capability, the responsibility for management changes. Experimental capabilities assembled by DARPA, for example, migrate to PoR support where they become common capabilities as opposed to single instances. For common capabilities, synergies emerge that can lead to improved implementation such as physical or virtual consolidation of data storage and management.

SOE-enabling infrastructure is coming into place, whether provided locally, regionally, or globally. This infrastructure allows users to find, access, and exploit

UNCLASSIFIED

common services and data sources. Despite the availability of an enterprise service registry, there is scant registration of DoD service offerings (planned or actual). Hence, it is difficult to find information on what services are being operated in support of any given theater. To truly characterize the environment, DoD must implement mechanisms to automatically discover all available IT resources, identify the specific users they are serving, and catalog their content. Visibility into what services are operational on the network or where services are in their development stage is crucial for assigning roles and responsibilities in DoD's IT life cycle management model. C2 services would be visible and efforts to manage and govern services and the SOE would be improved by adopting and implementing the C2ISF (see Section III).

Our case study analyses led to some general observations for NCSS implementation for the C2 community:

- To implement C2 capabilities in an SOE, DoD requires agile and collaborative governance that embraces the full range of IT engineering and operations activities. The governance must accommodate highly variable C2 node-specific arrangements at three levels: 1) those arising from local needs; 2) clusters of C2 capabilities that form common C2 services, and; 3) enterprise infrastructure that must be robust and stable with carefully planned changes due to large-scale dependencies. *(Case studies show that critical implementation action is occurring at all three levels and that all three are needed to make agility versus stability tradeoffs.)*
- DoD decision makers need a robust capability to know what specific IT is actually operating on DoD networks and to monitor its usage and performance. This requires instrumentation to collect metrics and feedback mechanisms that allow users to publish comments on the IT products and services they use. *(Case studies highlight a requirement to support and build on what is successfully deployed and operational. Implementation must build on what is being used and cannot significantly disrupt existing capability. This requires intimate knowledge of the current operational environment.)*
- The most effective control points appear to lie within the operational chain of command where C2 facilities, available IT, content, tactics, techniques, and procedures (TTPs) are constantly being assembled and adjusted to answer pressing requirements. *(Case studies reflect operational chain of command making final determinations on what capabilities to use or not use and where to invest in improvements in-theater.)*

UNCLASSIFIED

- DoD needs new processes for acquiring, managing, operating, and continuously improving C2 information services. The dynamic nature of an SOE drives a requirement for faster processes with more transparency with crisp lines of authority and accountability. *(Case studies show that IT capabilities generally require rapid improvement after Initial Operational Capability (IOC). This entails rapid decision making to identify technical options, apply funding, and engage engineering/training/logistics support in both rear and forward areas to implement enhancements.)*
- Knowledgeable developers with only moderate resources can successfully engage in-theater to formulate and implement significant C2 capability improvements. Local innovation can bring real and timely benefit to the warfighter. DoD authorities at various echelons can encourage innovation or repress it. *(Case studies show that significant local innovation will succeed in delivering valuable capabilities to warriors. Enemy forces, unconstrained by bureaucracy, will exploit commercial innovation faster than formal acquisition processes.)*
- Much of the system capabilities currently deployed is amenable to rapid and relatively inexpensive service-oriented adaptation. *(Case studies show that capabilities developed over decades through “Big A” acquisition are being significantly enhanced by Web technologies, once fielded.)*
- Services that emerge through local innovation can find a programmatic home either by being adopted by an existing PoR or by decision authorities at the common or enterprise level creating a new PoR to provide resources for them. *(Case studies show that locally developed capabilities need more robust support processes and resources once they become more widely accepted.)*

The SOE evolution is really a journey, not a destination. It is clear that without appropriate direction, oversight, and changes to DoD processes, the value of an SOE is not likely to be achieved. Efforts to evolve the C2 SOE will need to build from the evolving operational baseline, ensuring no loss of capability to engaged forces and senior leadership. Managing the transition of stovepiped C2 systems from their present largely client-server environment to an interdependent services-based, net-enabled enterprise is the major IT challenge facing the Department today. The existence and use of effective SOE infrastructure (both enterprise and community) is necessary to realize combinative processes and service clusters. The infrastructure is critical to the management and the transparency of the SOE.

UNCLASSIFIED

Our case studies have identified three key levels of this SOE evolution. The first is *local*: the activities involved in satisfying C2 IT requirements of a single organization, generally within a single area of responsibility (AoR), often within or among a very small number of physical command facilities. The second is *common* (or community): clusters of C2 mission capabilities that have evolved into more comprehensive and complex sets of interrelations. These clusters of capabilities are more widely distributed both geographically and organizationally than locally. As a result, they rely on a common set of information-sharing infrastructure and a more structured approach to operational control. The third is *infrastructure* to support interoperability, collaboration, and interaction of mission services. The infrastructure may be supplied through sets of community-specific services or enterprise services or a combination of the two. At present, the case studies suggest that very few C2 service-based capabilities are supported by enterprise infrastructure beyond the transport networks.

In the SOE, services must be discoverable, accessible, and useable by consumers in an operational (i.e., run-time) environment. This requires C2 and other mission services to be hosted and operated within a highly networked enterprise accessible to both service providers and service consumers. Important components of an SOE are federated infrastructure services and information-sharing templates that specify C2 services. The next section on the *C2 Information Sharing Framework (C2ISF)* describes the functionality of these SOE components and their interrelationships.

UNCLASSIFIED

(This page is intentionally blank.)

UNCLASSIFIED

III. C2 INFORMATION SHARING FRAMEWORK (C2ISF)

In a mature SOE, data and computational support are encapsulated in services, based on relatively inexpensive and widely available technologies. Data sources are abundant and do not require much sophistication to access and use. Participants in the SOE interact with services via published interfaces. Services in the SOE are available across the entire enterprise, limited only by security restrictions and network reach and capacity. Participants in the SOE can be expected to create and operate services for the enterprise, and they will become dependent on services located anywhere on the network. Intended users will consume the services, as will unanticipated users, who may use the services in unpredicted ways. As an SOE evolves, its participants and processes come to rely on stable collections of services that are created and operated independently, but that cluster in mutually beneficial federations.

These SOE characteristics imply the need for standards and infrastructures of general-purpose, utility services that enable users and services to navigate, access, trust, and understand the complex and changing information environment in which they operate. The infrastructure required to support the users and services is referred to as an Information Sharing Framework (ISF). An ISF is principally a set of services that provides information on the status, operation, management, and evolution of services in the enterprise. The ISF primarily involves the collection, processing, and publication of metadata—structural, organizational, and operational—to provide the following:

- Assistance in finding data and computational services (e.g., through taxonomies, registries, search engines)
- A small set of standards for networks and interfaces to reduce the complexity and increase the predictability of service interaction and implementation
- Assistance in assessing the trustworthiness and reliability of services
- Assistance in guaranteeing the identity of participants in communications
- Means to determine the rights and privileges of potential users
- Means to measure and publish data concerning the activity and performance of services

UNCLASSIFIED

- Publication of the structure and semantics of content available on the networks
- Assistance in assessing the effects and risks of currently operating portfolios of services.

A key aspect for federating DoD's SOE is the need to determine the correct balance between centralized control of, and enforced adherence to, the ISFs and the flexibility to introduce new elements to a framework to accommodate emerging technologies or edge innovations. Real-time interoperability of C2 services with the enterprise and other communities is dependent upon a federated SOE. The governance and management approach defined in the CONOPS (Section IV) must account for this need for balance.

A. C2ISF DESCRIPTION

The C2ISF is a collection of services and information used for creating, finding, using, and managing C2 services and information. It includes, for example, registries to hold C2 service descriptions so that potential users can easily discover services and learn how to communicate with them; information assurance capabilities to safeguard content; and monitoring/metrics collection mechanisms to inform governance.

This study principally addresses the portions of the C2ISF that are used for finding and understanding C2 services and information.¹ For example, the C2ISF includes registries to hold C2 service descriptions and templates for producing those descriptions. The C2ISF will not be implemented as a stand-alone suite of services; it will exploit existing enterprise services that provide related functionality. The C2ISF must be accessible by users anywhere. We expect there will be different (albeit federated) implementations of the C2ISF to provide the desired cross-community capability.

Information-sharing frameworks similar to the C2ISF are needed for other mission domains and for enterprise services. The C2ISF is distinct in that it will require C2-specific detail in the descriptions of the C2 services and information (different

¹ By "information," we mean data that have an associated meaning imparted by incorporated labels (such as XML tags) that use terms defined in readily available (published) artifacts (in this case, the C2ISF artifacts), or by being associated with published ontologies in a standard resource-description-tagging fashion.

attributes in the metadata from other communities), and it will be subject to performance constraints peculiar to C2. The additional detail will ensure that C2 data structures can carry the information commanders require. The detail will help C2 service users discriminate among service and information offerings when searching, improving the users' chances of obtaining what they want.

The C2ISF is illustrated in Figure 3 and has two major parts: the “Design-Time” infrastructure that includes all necessary data definition artifacts, and the “Run-Time” infrastructure that includes all the operational services needed to perform the discovery and information-sharing and management functions.




Figure 3. C2ISF Concept Diagram

C2 Service Owners use the Design-Time infrastructure when designing C2 services, and C2 Service Users use the C2ISF's Run-Time Infrastructure to share and discover information about those C2 services. The C2 Service Owners effectively advertise their services via the Run-Time Infrastructure to provide information such as what the services produce, who is responsible for them, and how to access them. A more detailed diagram of the C2ISF is provided in Figure 4.




Figure 4. Detailed C2ISF Diagram

1. Design-Time Infrastructure

The Design-Time Infrastructure includes artifacts, which define the terms used when describing data and services handled by the C2ISF, and design documents, which are used to build the Run-Time Infrastructure. The artifacts for the C2ISF include the Universal Core Semantic Layer (U-Core SL) and the C2 Core. The U-Core SL is an ontology² that describes the relationships among the general groupings of terms in the high-level U-Core schema. The C2 Core will extend the terms in the U-Core and the terms and relationships in the U-Core SL to more precisely describe C2 services. The extension of U-Core SL in the C2 Core and adoption of a small number of other schemas and ontologies will effectively constitute the development of the Joint C2 Conceptual

² The relationships described by ontologies are called “statement,” and have an abstract syntax. A simple example of an abstract syntax is the Resource Description Framework (RDF), which expresses semantic statements as subject, predicate, object combinations. This abstract syntax is easily expressed in Extensible Markup Language (XML), and any of the subject, predicate, or resource object parts are simply Universal Resource Indicators (URIs) (hence, “net-ready”).

Model and Vocabulary,³ which must define the relationships between the terms listed in the individual schemas (the Conceptual Model is, thus, an ontology). In addition to the Joint C2 Conceptual Model and Vocabulary, other schemas, taxonomies, and ontologies will likely be assembled to make the data encoded as C2-relevant information more readily shared between disparate groups in the C2 community. As the set of artifacts applicable to C2 is further embellished, with additional terms and with relationships between the categories of terms, the C2 services and information can be more accurately described and the C2 services can exchange more detailed information with new services.

The resources in the C2ISF Design-Time Infrastructure include the C2 Service Description Templates and the C2ISF Run-Time Infrastructure Rules and Protocols (see Appendix B, Section B for more detailed discussion). The C2 Service Description Templates, as schemas themselves, are similar to forms that Service Owners fill out to describe each service they create. The completed forms are stored in a common Service Registry that may be searched (e.g., via browsing, query, or faceted search). The templates are formally derived from the descriptions of C2 concepts defined by the C2ISF artifacts. Examples of the types of artifacts and corresponding data that would be useful for crafting the templates are listed in Tables 1 and 2, respectively.

³ There is currently an organization developing the C2 Core that has defined a C2 Core XML schema, though no organization has explicitly extended the semantic relationships in U-Core SL for the Joint C2 Conceptual Model yet.

UNCLASSIFIED

Table 1. C2ISF Service Description Templates' Artifacts and Concepts

Root Artifact Category	Concepts Encoded from Artifact	Concept Explanation/Comment
Dublin Core/DoD Discovery Metadata Specification and Extensions	Coverage	Spatial or virtual (cyberspace) coverage types included
	Topical coverage	Informally a description, more formally a semantical model (cf. Operational/Functional Artifact Category)
	Owner/Author/Creator	POC information
	Owner/Operator's relationships in the organization	
Governance Information	Acquisition status	JCIDS-related information, e.g., funding source and requirements
	Operational employment	For instance, what commands are using it
Technical Information	Technical service specification (inputs and outputs)	To lead into WSDL, if a (mature) Web Service is being described
	Dependencies	Dependencies on other services, expressed in elementary fashion (e.g. "is_dependent_on") or more formally, cf. the below
	Formal description of processes	Including workflow, orchestration, and even algorithms, as possible and appropriate
Operational/Functional	Performance/Use	Also Maintenance and Provisioning
	Dependencies	Data sources, but also needed hardware, software, bandwidth, and communication infrastructure
	C2-Objects, Core Taxonomy, Joint Common System Function List/JCA-derived concepts	
	Joint IC/DoD Enterprise Services Registry Taxonomy (Appendix C) concepts	
Upper Ontology		

UNCLASSIFIED

Table 2. Example Data Requested by C2ISF Service Description Templates

Data Category	Data Item
General Data	
	Service name
	Contact info (POCs) (Service Owner)
	Version
	High-Level Description (including motivation)
	Location (URI)
	Operational Status
	Classification
	Other Security Data
	Other joint IC/DoD Enterprise Services Registry Taxonomy's (Appendix D) required data: Namespace, Creator, Publisher, Creation Date, Effective Date, Validation Date, End of Life Date, Geographic Coverage, URI of Related Data Resources, Rights to Data (copyright, etc.), Classification Data (including dissemination and access controls).
Governance Data	
	Current JCIDS milestone achieved
	Link to source of JCIDS documents
	JCIDS schedule and POCs
	Funding sources
	Requirements
	Which authorities ensure it is useful (Process Owner)
	What commands are using it
	Access control policies
Technical Data	
	Technical Service Specifications (Input Format, Output Format, call procedure, Standards and Technologies used) (for a Web service, this is a WSDL)
	Dependencies on other services (e.g., list of services and data required for full operation)
	Description of Algorithms (a detailed explanation of how the service works)
Operational/Functional	
	Performance and Usage Metric measurements (usage statistics)
	Maintenance and Provisioning Data: when and where it can be expected to work, how much traffic it can handle (calls/hour)
	Dependencies: needed hardware, software, bandwidth, data and data sources, communication infrastructure, schemas
	Categorization according to Joint IC/DoD Enterprise Services Registry Taxonomy (Appendix C)

UNCLASSIFIED

The terms used to describe an individual service's attributes will ideally be taken directly from controlled vocabularies in the collection of C2ISF artifacts (c.f. Appendix B for a discussion on Knowledge Organization System employment, in this context). For example, a C2 Service Description Template may have a concept for what region the service is applicable and another for what level of support the C2 Service Owner can provide to users. At the most elementary level, the template then would only allow the options "Global," "CONUS," "USSOUTHCOM," "USEUCOM," "USCENTCOM," "USPACOM," "USAFRICOM," and "Local" for the region concept, and "Full Support," "Helpdesk only, 9–5 Eastern Standard Time," and "No Additional Support" for the level-of-support concept. Constrained descriptions derived from the controlled vocabulary make it easier to refine a search to return only services applicable for "USEUCOM" or "Global" and providing "Full Support."

An engineering-level description of how the C2ISF operates is maintained in the C2ISF Run-Time Infrastructure Rules and Protocols. They specify how C2 Service Owners should interact with the C2ISF and how the services in the C2ISF interact with one another. For example, one of the rules would be that all C2 Service Owners must provide descriptions of their services to one of the C2ISF service registries (i.e., "register" their services), regardless of service complexity. An example of a protocol would be a detailed description of how the service registries are federated, including the form of the queries that a service search engine should use to automatically search one of the registries in the federation.

2. Run-Time Infrastructure

The C2ISF Run-Time Infrastructure has critical, necessary, and optional sets of services. The critical minimum set of services is:

- One service registry
- One artifact registry
- One information registry
- A URI/Universal Resource Locator (URL) management service
- Other basic C2 or enterprise infrastructure services, such as Identity and Cryptographic Services.

A service registry is a service that stores and provides descriptions of services and how to use them. Those descriptions should, ideally, include everything the user needs to

UNCLASSIFIED

know to find and use the service. Developmental, non-operational services require less information to describe them, but they should nonetheless be registered. An artifact registry is a service that stores and provides schemas, data definitions, ontologies, and other artifacts the C2 service owners need to consult when designing their services. An information registry is a service that stores and provides descriptions of information sources, similar to popular indexing engines like Yahoo!'s SearchMonkey. It is not necessary, neither as a consequence of using services nor a limitation of current technology, that these different registries be separate in practice. Nor is it necessary that all the services' description data be stored in the registry if a static address to the information can be made available. A URI/URL management service is used to ensure that every service and site for information is given a unique network address and name. It will be used, in particular, to lend URIs to newly registered services and data. The other basic infrastructure services are those necessary for communications and security. These include the identity service that ensures each user has a unique identity⁴ (for the employment of subsequent cryptographic functionality) and other cryptographic services that are used to assure interactions among the C2 service users.

The additional services that complete the set of C2ISF run-time infrastructure services that are necessary, in practice, are:

- Monitoring/Metrics Collection
- Role Management
- Policy Enforcement
- Data Mediation/Schema Transformers
- Other Essential Services, such as Schema Verification
- Additional Registries.

Monitoring and Metrics Collection services record the use of services and information sites automatically (either by having the services themselves report all calls they receive or by parsing network logs) and compute metrics that are descriptive of the services' or information's use. The Role Management service allows authorities to define roles for users that will give them access to services and information, and assign those roles to C2 Service Users, as appropriate. The Policy Enforcement service checks

⁴ In the form of a certificate derived from a public/private key pair.

that a particular C2 Service User has an assigned role that allows the user access to a service or information before it allows the service or information site to respond to that user. It is important to note that a C2 Service User can also be another service that is calling the C2 Service in an automated fashion. Data Mediation is a service that represents data that are stored according to one schema in the terms of another schema that the C2 Service User understands. This is accomplished through the use of Schema Transformers, which are used to translate each piece of information into information conforming to the new schema. Schema Verification is a service that checks that input information conforms to a given schema, which is important for ensuring consistent descriptions in registries. Lastly, given the geographically distributed nature of the C2 mission, using only one large registry for all C2 services and information products can cause undue access, performance, and management problems, so implementing a federated registry system is a practical necessity.

Some services in the C2ISF that are optional, but useful, include:

- Crawlers
- Tagging Engines
- Feedback Services
- Reputation Service.

Crawlers are programs that can recursively follow URLs to find all the URLs associated to a given Web resource. This service is useful for filling information registries, especially when combined with a tagging engine. Tagging engines take documents or services and infer metadata from them. Feedback services are automated ways of collecting and organizing user feedback on the performance of services registered in the C2ISF (including the C2ISF services themselves). Feedback services help the C2 Service Owners and C2 services portfolio authorities determine the needs of their users. A Reputation service is a data service that provides data describing the authoritative attributes of a particular C2 Service Owner or information author, so C2 Service Users may better choose the service or data source that is right for their application.

Several of the key interactions between C2ISF entities are illustrated in Figure 6 with arrows. These are not a comprehensive set of likely interactions but illustrate ones of particular importance. In Figure 5, the dotted, red arrows represent invocation actions

(requests). The dashed, green arrows represent registration actions (service or information description submissions), and the solid, black arrows represent Product Delivery Actions (responses). These arrows are briefly described in Table 3.




Figure 5. Detailed C2ISF Diagram with Information Exchanges

Table 3. Representative C2ISF Interface Exchanges

Arrow Number and Description	Example Actions
1. User to C2 Service	Calling the service
2. User to Metadata Capture Services	Requesting collection of monitoring data or collection of metadata from an input location or document set
3. User to Metadata Interpreting Services	Requesting translation, mediation of registry metadata
4. User to Service Registry	Sending a query to the registry to find a service
5. User to Feedback Services to Design-Time Infrastructure	Sending feedback to the people updating the designs
6. C2 Service to User	Replying to the user request
7. C2 Service to Other Essential Services	Checking the user's identity, encrypting response to user, requesting of a URI by an instance, etc.
8. C2 Service to Service Registry	Performing automatic registration of URI as endpoint in service description entry
9. C2 Service Owner to C2 Service	Providing the C2 service

UNCLASSIFIED

Arrow Number and Description	Example Actions
10. C2 Service Owner to Metadata Interpreting Services	Checking available mappings between schemas
11. C2 Service Owner to Service Registry	Submitting the C2 service's description to the registry
12. Artifacts to C2 Services Description Templates	Constraining or informing the templates
13. C2 Services Description Templates to C2 Service Owner	Providing the templates to the Owner to submit the service's description
14. C2 Services Description Templates to Service Registry	Submitting a description of the template in the registry, so Owners can find the templates they need
15. Service Registry to Metadata Interpreting Services	Sending requested metadata to metadata interpreting services
16. Service Registry to Service Registries; Service Registries to Service Registry	Sending a query to the service registries to ensure the C2 service isn't already registered somewhere else; Sending a query to the registry to see whether it has services similar to what the user wants.
17. Metadata Capture Service to Service Registry	Storing collected metadata in the corresponding registry entry
18. Other Essential Services to Service Registry	Sending confirmation of identity, information about reputation of C2 Service Owner and C2 service that are trying to register, sending decrypted documents to registry, identifying which metadata follow the appropriate schema, etc.

3. Relationship Between C2ISF and Services and Data Strategies

Implementation of the NCDS is closely interrelated with implementation of the NCSS. Services can provide computational processes as well as access to information products. Much of the NCDS was premised on using services to provide accessibility to data or information sources. Figure 6 illustrates the relationship between some of the implementation activities for both the NCDS and the NCSS and the C2ISF. As discussed earlier in Section III.A.1, the ongoing C2 Core development activities will influence the service description templates. Thus, Figure 6 illustrates how a conceptual Data Service “X” would be defined and created in accordance with the C2ISF design-time artifacts (on the left side of the figure) and then fielded and exposed for use via the C2ISF run-time infrastructure services (on the right side of the figure). Departmental guidance for implementing both of these strategies must recognize the interrelationship between the activities and the dependence on the C2ISF.




Figure 6. Harmonizing NCDS and NCSS Implementation Activities

B. IMPLEMENTATION AND MANAGEMENT OF THE C2ISF

The run-time environment of the C2ISF is a logical collection of services. One important management activity for the C2ISF will be to determine what set of functionality must be assembled to allow C2 information sharing. It is likely that part of the needed functionality will be covered by enterprise infrastructure services, which will be operated independently. A suite of services designed specifically for the C2 community will provide the remainder of the functionality. We call the services designed for the C2ISF “C2ISF-specific services.” C2 services portfolio governance will have to commission and lead the development of the C2ISF-specific services, which will be managed primarily as C2 infrastructure services. It is also possible that services from other communities will be used to provide some of the needed C2ISF functionality. All services that are used for the C2ISF, regardless of the source, should be managed through a well-defined life cycle process (see, for example, the CONOPS in Section IV).

Another important C2ISF life cycle management activity is to define the performance requirements for the C2ISF at each host. In accordance with the Section IV CONOPS, the Customer/Operational Process Owner (OPO) will be responsible for defining performance requirements of the canonical set of C2ISF-specific services. Similarly, the Customer, OPO, and Service Owner will be responsible for testing the C2ISF performance against requirements. The organizations that manage any host setup, provisioning of network resources, and monitoring of C2ISF functioning will be working on behalf of the Service Owner.

UNCLASSIFIED

The schemas in the design-time environment of the C2ISF will be maintained by a C2ISF librarian, who will maintain the latest versions of the C2-relevant ontologies, taxonomies, schema, etc., and derive the C2ISF Service Description Templates from them. In doing so, the librarian must compile and act upon the needs of C2ISF users and customers for refinement of the template schemas.

1. C2 Services Description Templates

A regular structure and vocabulary for search improves the ability of C2 Service Users to find information and services. A more detailed discussion of C2ISF template creation and management is given in Appendix B. Formal systems for the storage and management of such data are generally known as Knowledge Organization Systems (KOSs). A more detailed description of KOSs is given in Appendix B.

To keep the templates for describing C2 services current, so they can effectively describe and differentiate between services, the organizing system for the types of data in the templates must be updated to include new terms and new relationships between terms. This change-request handling must be performed by a set of skilled librarians who will have the authority to institute the changes to the official version of the templates.

For services in different C2 service tiers (see CONOPS, Section IV.C), there will naturally be different data required when filling out the Service Description Templates. For users to effectively discover the services they need, there is a minimal amount of metadata required for each C2 service: the service's URI (URL where possible), Service Owner, Point of Contact, High-Level Description (of its functionality), Virtual Coverage (what network domains on which it operates), Access Procedure, Operational Status, and Classification Level. One fundamental C2ISF rule is that all C2 Service Owners must provide descriptions of their services to one of the C2ISF service registries (i.e., register their services). This requirement is necessary to allow visibility into the state of the SOE.

2. C2ISF Run-Time Infrastructure Rules and Protocols

The C2ISF run-time infrastructure rules and protocols are the engineering-level description of how the C2ISF operates. They are the critical components of the design of the C2ISF that explains how the various software, information, and standards (e.g., W3C standards, C2ISF artifacts) will be coordinated to enable C2 information sharing. The C2ISF run-time infrastructure rules and protocols must be designed to provide the information needs of the C2 services while accounting for the constrained network

UNCLASSIFIED

capabilities and characteristics. The Service Owner and C2 infrastructure portfolio governance authority must provide technical management of the C2ISF in accordance with the rules and protocols. A more detailed discussion of C2ISF run-time infrastructure rules and protocols is given in Appendix B.

UNCLASSIFIED

(This page is intentionally blank.)

UNCLASSIFIED

IV. C2 SERVICES CONOPS

A. INTRODUCTION TO THE CONOPS

This study will use the C2 services CONOPS as an organizing construct that ties together C2 service categories, key implementation roles and responsibilities throughout the service life cycle, and a high-level governance construct. The CONOPS identifies tiers of C2 services based

The CONOPS identifies tiers of C2 services and associated portfolios, provides a structured method for examining life cycle activities and roles, and assigns actors to roles for each C2 service tier.

on their characteristics. The CONOPS serves to illustrate how C2 services can be developed, provisioned, operated, and used in the SOE. The CONOPS recognizes the dependence of C2 services on other parts of the SOE including the enterprise services and other community or domain services. The CONOPS also recommends governance authorities for C2 services portfolios.

The CONOPS was developed through a process (Figure 7) that uses industry standards, tailored activities, defined roles, and comparisons with real case studies. We postulated a service life cycle and key provider and consumer roles based on industry standard frameworks. For each stage of the life cycle, we defined some key activities for that phase. The roles and activities were compared to our example scenarios and the defined characteristics of each tier. That comparison was used to update elements of the CONOPS. Finally we have identified example actors to fulfill the key roles.




Figure 7. Developing the C2 Services CONOPS

B. OVERVIEW OF C2 IN AN SOE

Evolving the DoD to an SOE has the potential to significantly reduce the C2 OODA (observe, orient, decide, act) loop cycle time for both operations and IT acquisition. By unleashing Web technologies and an Internet-savvy warfighter, the approaches and processes used in C2 will be continually evolving toward collaboration, cooperation, and coordination. In the future SOE, any commander would have access to a wealth of information resources through Web-based capabilities, tools to generate and exploit information, and the ability to use available content and tools without having to wait for an acquisition organization to design, buy, build, and install a system. Information flows horizontally and vertically.

Local command capabilities are constantly being born and shared on the networks. These local services are developed quickly and at low cost. Warfighters begin to use these services; the high-value services achieve rapid and broad-based adoption. Soon, demand outstrips the ability for local commands to support the capability, and a regional command or a military Service assumes responsibility for continued evolution, sustainment, and support. For the foreseeable future, programs of record will provide that support. Commanders and developers have a full range of services available to them

to enter into cooperative arrangements with other providers and consumers to quickly create new value-added capabilities and processes in response to emergent mission demands. A robust global infrastructure exists that is composed of enterprise, regional, and local elements. The C2ISF is adopted and implemented. The run-time infrastructure enables seamless operation of services across geographic and organizational boundaries.

The future SOE implementation and operation approach for C2 involves knowing emerging user needs, understanding what is happening on the networks, leveraging commercial technology, continually monitoring, analyzing, reporting, and formulating responses, and adopting best practices in IT services management guidance. Decision-making about resource allocation is informed by access to accurate and current information about what IT services exist, which are useful, and what improvements are needed. Acquisition, delivery, and deployment occur in more “commercial” timeframes versus DoD timeframes. Achieving an SOE entails adjustments in both operations and acquisition.

C. C2 SERVICE TIERS

Organizing for an evolution to an SOE entails adjustments in acquisition *and* operations. Effective governance of the C2 SOE evolution requires (1) assigning responsibilities and authorities for the globally distributed activities, and (2) addressing the agility-stability tradeoff between highly responsive local needs and the need for more stability for common and enterprise services. A multi-tiered C2 services governance structure can accommodate the agility demands of local C2 nodes and the reliability and control needed to federate services across the C2 community and the rest of the enterprise. The multi-tiered structure can facilitate provisioning both specialized and general-purpose capabilities without seriously sub-optimizing the former or over-restricting the latter. In summary, a tiered C2 services management and governance structure:

- *Enables Appropriate Operational Controls*—The goal is to have all services visible across the enterprise, but it is not practical or desirable for all services to be available for use by anyone in the enterprise, especially when considering bandwidth or degraded operations. Since the model is intended to address authorities across the entire service life cycle, it is necessary to consider the implications of different levels of operational control based on mission needs.

UNCLASSIFIED

- *Provides Unity of Effort*—An enterprise the scale of DoD cannot be managed and governed monolithically. A key aspect of implementing an SOE is to have well-defined and established governance. A tiered model allows for governance of “enterprises” within the DoD enterprise. A well-defined model for addressing the full range of services across the DoD will yield unity of effort in implementation, operation, and management of the SOE.
- *Promotes Local Innovation*—The tactical or edge user needs to have the ability to respond to a rapidly changing environment. That need for agility drives a “lighter” or less formalized approach to service life cycle management. Recognizing a local tier of services allows for a less formalized implementation and management model for that specific category of services.

For purposes of analyzing the necessary control and governance structures throughout the service life cycle, we have categorized C2 services into tiers based on characteristics and usage (Figure 8).




Figure 8. C2 Service Tiers

This model supports increasing flexibility and agility as one moves up the “stack” with a need for increasing degrees of robustness and stability in service operations as one moves down the stack. Service dependencies in terms of both human and machine users decrease as one moves toward local and increase as one moves toward enterprise. This model requires explicit recognition of governance at these three levels but does not

UNCLASSIFIED

require any particular management construct for either enterprise services or other domain services such as personnel or logistics. For purposes of this report, we have split enterprise services into core enabling enterprise services, such as authentication, enterprise service management, and discovery, and all other enterprise services. Core enabling enterprise services (also called infrastructure services) are especially important for allowing federation of services across all domains.

We define the three tiers of services for the C2 domain or community in Table 4.

Table 4. Definitions of C2 Services Tiers

Tier	Definition	Examples
C2 Local	Mission-oriented information services that are tailored to meet the needs of a limited group of users, e.g., specific organizations or entities, usually within a single organization or AOR.	• Initial deployments of TIGR, CIDNE
C2 Common	Capabilities that fulfill data or functionality requirements inherent in multiple C2 missions but that are not expressly tailored to, or necessarily useful for, supporting other mission areas. C2 common services will typically be available for use by multiple commands in one or more AORs (i.e., regional) or globally.	• Planning capability such as JOPES • Air tasking planning tools • Blue Force Tracking
C2 Infrastructure	Mission-specific, general-purpose capabilities, configured expressly to address C2-community-specific performance requirements, business processes, or behavior characteristics, which provide for basic communications, collaboration, publication, discovery, security, and information and service management. C2 infrastructure services may be instantiated in conjunction with C2 local and C2 common services as required. C2 infrastructure, much of which will exist as part of the run-time elements of the C2ISF, is critical to enable interoperability across the C2 domain.	• Mission-specific access control services such as policy enforcement, role management • Data mediation • Extension for specialized search • CPOF C2 specific collaboration

Most services in the C2 common and C2 infrastructure tiers will continue to be delivered through PoRs. However, most PoRs are not designed to provide only mission functionality or only infrastructure functionality. Any given PoR is likely to be developing new mission services, possibly some new infrastructure services, and likely some non-service-based capabilities. PoRs should also be looking to reuse existing services from the C2 domain, the enterprise space, or other domains. The use of PoRs to deliver services or capability built around services does not conflict with the tiered C2 services model. The model represents a way to describe the characteristics of services and to examine the management and governance of those services. Program managers

for PoRs will have to adhere to the appropriate life cycle roles and responsibilities for their specific service tiers. Governance for those service tiers will include assessments of the PoR plans and status. The tiered structure provides a way to begin migrating from stovepiped C2 systems to an interdependent, services-based, net-enabled enterprise.

The tiered C2 services model provides the basis for examining the differences in the characteristics of the tiers. Services will fall into tiers across the continuum of this model and will move across tiers as service usage, scope, and implementation considerations drive the need for more or less control (Figure 9).




Figure 9. C2 Services Tiers and Varying Characteristics

All C2 services are registered and discoverable across the enterprise. However, we expect service accessibility to be greater for C2 common services than for C2 local services. It may not be desirable or practical for C2 local services to be accessible to everyone in the enterprise because of operational and technical limitations such as mission priorities, bandwidth limitations, and server capacity issues. Similarly, we would expect user feedback to have a more immediate and direct effect on C2 local services than on C2 common services. By recognizing multiple tiers, an appropriate management and governance model can be developed for each tier.

UNCLASSIFIED

A key characteristics summary of the C2 service tiers is identified in Table 5.

Table 5. Key Characteristics of C2 Services by Tier

Tier	Key Characteristics
C2 Local	• Requirements driven by local mission need • Service level guarantees minimal and based on local demand • Short delivery and upgrade cycles driven by local decisions • Access for users outside the specific organization or AOR may be very limited (but services should still be visible across the enterprise) • Fewer architecture, engineering artifacts developed • Change control decisions made locally • Funding may be from operations budgets
C2 Common	• Requirements driven by multiple AORs and may be common across all commands and/or all C2 missions • Implementation architecture driven by need for distributed use • More rigorous enforcement of service-level provisions • More rigorous planning for release and upgrade cycles • Acquisition and implementation driven through PoRs • Services accessible to global C2 users • Operational support to meet user needs in potentially very different locations • Change control decision authority will affect multiple commands/AORs
C2 Infrastructure	• Requirements largely derived from mission service needs and implementation considerations • Significant engineering analyses required to architect for performance, scalability • Rigorous and well-planned change control process • Well-understood and enforced service-level guarantees • Visible and accessible to end users (as required) and service developers

Another reason to consider multiple tiers of services in the C2 community is to recognize the uniqueness of the C2 mission from other missions. The C2 mission requires maximum agility to respond to local commanders' needs and rapidly evolving situations in theater. Hence, we would expect more services in the local tier for the C2 community than if this same model were applied to other communities such as finance or logistics. A similar model in the financial community would have many more services in the common layer than in the local layer because of standardized financial reporting and accounting regulations. In the financial community, the goal would be to provide those as common services. Our case studies (Appendix A) illustrate where service-based capabilities initiated for local command use have been deployed and enhanced for broader use by the C2 community. These real-world examples drove the need to develop a C2 domain services model that has flexibility and a continuum of tiers.

In this model, services are not "static" with respect to their tier. In a well-functioning SOE, services would come about as a result of local mission need, with an

associated “informal” governance structure, and experience a rapid and broad-based user uptake making them more common. The increase in user dependency would then require that the service migrate to a management and governance structure associated with C2 common services. As indicated in Table 5, C2 common services support a broader user base and therefore, need to be resourced for more robust application and more rigorous change control to avoid broad mission disruption. Similarly, infrastructure services may migrate as extensions of enterprise services, or elements of the run-time C2ISF may eventually become an enterprise service. Similarly, legacy enterprise service versions can devolve to become domain specific in the course IT capability migrations. Services can, and will, move within the tiers; that is, tiers do not have distinct boundaries but represent a continuum of service types in the C2 space. The variability of characteristics of service tiers is more interesting as we look across the service life cycle and the governance of portfolios of these services.

D. SERVICE LIFE CYCLE

Implementing an SOE will result in a far richer and more complex information environment and therefore requires special attention to management of mission-specific and infrastructure services. Industry experience shows that it is easier and more effective to implement services and an SOE when the organizations involved have adopted an ITIL/COBIT-like governance process.

1. ITIL/COBIT

The level of complexity in the SOE demands a structured framework to describe and analyze the roles and responsibilities for life cycle management. The two most widely used frameworks in this respect are the Information Technology Infrastructure Library (ITIL)¹ and the Control Objectives for Information and related Technology (COBIT).² ITIL provides a cohesive set of best practices for the management of IT service provision. ITIL is being adapted for use across a variety of DoD components. The Defense ITIL³ effort is developing IT Service Management (ITSM) process

¹ ITIL® Home, www.itil-officialsite.com/.

² COBIT, www.isaca.org/cobit/.

³ Defense ITIL, https://www.intelink.gov/wiki/Defense_IT_Infrastructure_Library.

guidelines and specifications to integrate DoD CIO objectives for common enterprise-level, Department-wide processes.

COBIT provides a consensus of good practices across a process framework and presents recommended activities in a manageable and logical structure. The practices are focused more strongly on control, less on execution. These practices will help optimize IT-enabled investments, ensure service delivery, and provide a measure against which to judge when things go wrong. ITIL and COBIT have been mapped to show the correlation and overlap between the frameworks.⁴

Both industry standards take a complete life cycle view of IT services. Both ITIL and COBIT prescribe a set of life cycle activities that are needed to manage and govern the actions of service providers and those who govern and control the use and provisioning of services. By using these industry standard life cycle models, we take advantage of a substantial body of best practices to tailor for our C2 analyses. For our analyses, we have drawn from both COBIT 4.1 and ITIL V3 models to define the life cycle phases, shown in Table 6, for C2 services.

Table 6. C2 Services Life Cycle for CONOPS

Life Cycle Phase	Representative Functions/Processes
Plan and Organize ^a (Service Strategy ^b)	Business environment, financial management, demand management, customer management
Acquire and Implement ^a (Service Design ^b)	Catalog management, service-level management, capacity management, availability management, service continuity, security management, sustainability management
Deliver and Support ^a (Service Transition ^b)	Change management, configuration management, release and deployment, service evaluation, service validation
Operate and Control (Service Operation ^b)	Content production management, knowledge management, event management, incident management, request fulfillment, problem management, and access control
Monitor and Evaluate ^a	Feedback mechanism management, instrumentation management, thresholds and metrics definition, effects vs. capabilities correlation
Continuously Improve ^a (Continual Service Improvement ^b)	Service adjustments to mitigate negative effects and improve service functionality or performance
^a COBIT 4.1 ^b ITIL V3	

⁴ “Mapping of ITIL v3 With COBIT® 4.1,” IT Governance Institute, 2008.

2. C2 Service Tier Characteristics

With a life cycle model established, we can further understand the C2 service tiers by considering the differences between the tiers across the life cycle. Table 7 identifies characteristics of each C2 service tier across the life cycle.

Table 7. Service Tier Characteristics by Life Cycle Phase

	C2 Local	C2 Common	C2 Infrastructure
Plan, Organize	• Ad hoc response to mission need • Requirements, user base “validated” by local command • Typically low dollars for local implementation • May be funded out of command budget • Less emphasis on business case, RoI	• Mission needs identified by multiple units/commands • Requirements “validated” through formal process; if based on existing service (e.g., C2 local service or pilot), then requirements emphasize “as is” capability • Funded out of PoR budget • Business case, RoI based on theater validation	• Requirements derived from engineering analyses of mission needs and existing service/network performance • Typically funded out of PoR budget for a <i>mission capability</i>; however, infrastructure “need” may not be solely driven by the PoR mission capability • Business case, RoI difficult to do as stand-alone services • Portfolio view required to address derived C2 infrastructure “needs”
Acquire, Implement	• May be developed in field or with available technical support • Short timelines (weeks-months) for implementation • Heavy user involvement in design, testing • Extensive use of existing data sources, systems, and COTS products • Little documentation for service implementation • Registration in most locally available service registry; discoverability <i>may</i> be limited based on user attributes (e.g., role, command)	• May be evolved from C2 local service, pilot capability, or acquired as a new requirement • Built for scalability to support broad and potentially large user base (e.g., multiple roles, multiple AORs) • May be implemented on an existing C2 or enterprise infrastructure • Implementation may be multiple short timeline (months) releases • Significant end user involvement in testing • Service implementation well-documented • Registered and widely discoverable but access may be limited	• May be derived, evolved from C2 local or common service or existing enterprise infrastructure service • Built for scalability and to support range of mission services • Implemented on an existing enterprise infrastructure service framework • Limited end-user involvement in design, implementation, testing through mission capabilities • Must also provide test environment for mission service developers • Service implementation well-documented • Registered and widely discoverable and available

UNCLASSIFIED

	C2 Local	C2 Common	C2 Infrastructure
Deliver, Support	- Delivered to specific command or AoR - May be configured, installed to run on locally available HW, OS (vs. new) - Change, configuration management more ad hoc and based on local user needs - New releases on short timelines based on user feedback - User support provided by development team	- May be deployed as single instance or replicated in multiple locations for performance - Change, configuration management handled through a single operational command authority - New releases planned and delivered based on operational user feedback - Clearly defined operational support plan for problem resolution, incident management	- May be deployed as single instance or replicated in multiple locations for performance - Change, configuration management handled through a single operational command authority; impacts to other C2 services and from enterprise infrastructure services must be assessed - Delivery of new releases based on technical and performance needs - Clearly defined operational support plan for problem resolution, incident management; may be tied to other infrastructure support plans
Operate, Control	- Content management at local level - Even if widely discoverable in service registry, access may be limited - Authority to limit access or install changes rests with specific command or AoR commander - May be service-level assertions; limited to specific user base with others using at own risk - Service management (i.e., fault and performance reporting) may be limited to local network monitoring	- Content defined beyond local level; local commands contribute - Visible to all; access may still be limited based on user role, authorities - Authority to limit access or install changes rests with single operational command authority; ad hoc change to access limitations unlikely - Service-level agreements or assertions apply to all authorized users - Service management implemented across the network - Service performance visible across network	- Typically content agnostic; may have mission-specific rules for content handling - Visible to all; access may still be limited based on user role, authorities - Authority to limit access or install changes rests with single operational command authority; unlikely to change access limitations on ad hoc basis - Service-level agreements apply to all authorized users - Service management implemented across the network - Service performance visible across the network
Monitor, Evaluate	Enhancements/fixes prioritized based on specific end user feedback involved in requirements, development, testing; others accommodated as practical	Service performance, user demand reported and used to influence actions in acquire/implement, deliver/support, operate/control life cycle stages	- Service performance, user demand reported and used to influence actions in acquire/implement, deliver/support, operate/control life cycle stages - Portfolio manager uses service performance, user demand to influence resource allocations

UNCLASSIFIED

	C2 Local	C2 Common	C2 Infrastructure
Improve	- User drives upgrades - Short timelines for improvements - Limited tech refresh upgrades under commander authority - Expanding user base and demand for enhancements can shift responsibility for improvements outside immediate command authority	- Upgrades, enhancements done as mission needs change - Plan for upgrades developed by acquiring authority based on user needs and approved by single operational command authority - Improvements can be based on user feedback, performance limitations, or emerging requirements	- Plan for upgrades developed by acquiring authority based on engineering analyses and funding constraints - Schedule for upgrades based on potential effects on mission users - Improvements can be based on performance limitations, emerging requirements, technology refresh

Figure 10 graphically summarizes the different characteristics of the service tiers. This figure demonstrates the continuum of services within the tiered model. The variability of service tier characteristics across the life cycle means that we need dynamic management and control structures to govern the SOE. In order to be dynamic, those structures and decision bodies must have access to continually published information about the status of services on the networks and plans for new services.




Figure 10. Variation in C2 Service Tier Characteristics Across Life Cycle

Figure 10 illustrates the gradation of service tier variability and the continuum of service tiers. The color gradation color indicates changes in characteristics as the color shifts from darker to lighter. For example, in looking at the plan and organize life cycle phase, three characteristics change as services move through infrastructure, common, and

local tiers. In the case of the user base characteristic, for C2 infrastructure services, the user base is generally unknown, but as services move up through common and local tiers, the user base becomes more known than unanticipated. This is because C2 infrastructure and C2 common services offer to a broad user base, while C2 local services support a distinct user base. Similar analyses apply across the life cycle for various characteristics.

3. Roles in the Service Life Cycle

The complexity of the SOE drives the need to have well-defined roles for executing and managing life cycle activities. Governing the C2 domain in the SOE requires that life cycle roles be defined and the identity of the actors responsible for executing those roles within in each C2 service tier are published. Based on the ITIL/COBIT models, we have identified a set of key Government roles (Table 8) that require specific organizational and individual assignments. These roles, typically held by Government personnel, identify those responsible for (1) the process for defining/refining service requirements (Customer); (2) the service operational process (Operational Process Owner, OPO); (3) the process for physically providing IT to support service capabilities (Service Owner); and (4) use of the service (Users). A key feature of this CONOPS is to explicitly acknowledge the need for an OPO who has a critical role in defining, managing, and executing an operational mission process that uses the service-oriented IT. This role is responsible for defining how a service might be used in a process and developing or modifying TTPs to incorporate the service in the mission process.

Table 8. Principal Roles in C2 Services CONOPS Life Cycle

Role	Description
Customer	- Provides requirements and potentially resources for local instantiations of C2 services - Negotiates Service Level Agreements (SLAs) with Service Owner (when appropriate) - Advocates for capability needs
Operational Process Owner (OPO)	- Source of operational process definition and TTPs and determines how information services will be used to support operational processes - Source of authoritative capability needs for services (and required improvements thereto) in support of C2 missions - Advocate for resources—and, in some cases, provider of resources
Service Owner	- Responsible for creating, acquiring, fielding, managing, supporting, improving the information service—throughout the life cycle - Responsible for identifying and managing the service provider - Negotiates SLAs with customers and functional requirements with OPO - Provides support to the OPO and users throughout the life cycle
Users	Personnel that use the service in executing missions

Basically the roles can be thought of in the context of provider functions (Service Owners) and consumer functions (customer, OPO, users). The customer and the OPO are owners of the technical and functional requirements for services. The difference between an OPO and customer role would be highlighted in the example of a C2 common service such as JOPES. In that case, each Combatant Commander would be a customer for a service of JOPES; his or her representative would be responsible for stipulating performance and integration requirements for the particular AoR to the Service Owner (in the case of JOPES, the DISA GCCS-J Program Manager). However, the OPO for JOPES is the JFCOM J3 who defines the common process and develops TTPs needed to execute the planning function. The Service Owner must be able to synthesize the requirements to guide the service provider implementation, provisioning, and support activities. The principal role relationships are illustrated in Figure 11.




Figure 11. Key C2 Services CONOPS Roles

While the formality of designating these roles will vary by C2 service tier, the functions need to be accomplished whether expressly designated or not. In the case of C2 local services, the customer and OPO might be the same actor, but for C2 common and infrastructure services, that is less likely to be the case. Most importantly, it is necessary that any designation or delegation of roles and responsibilities be published on the

UNCLASSIFIED

network for any service to ensure that status of actions can be quickly determined. A process must be established for determining and publishing the actors responsible for executing life cycle functions within each C2 service management tier.

Another characteristic of the C2 CONOPS is that it can accommodate the changing nature of the tiers across the life cycle of the services. The real value of the tiered model for C2 services is to ensure that the appropriate type and level of management and governance is applied for each tier across the life cycle. To identify appropriate levels of management and governance, we looked at our case studies and identified actors to execute the life cycle roles for each tier of C2 services. Table 9 captures typical actors for each tier and provides examples for consideration.

Table 9. Typical Actors for C2 Service Roles

Tier	Typical Actors for C2 CONOPS Roles				Comments
	Customer	OPO	Service Owner	User	
C2 Local <i>Example: TIGR</i>	Unit commander <i>1st Calvary Division</i>	Unit commander <i>1st Calvary Division</i>	DARPA, JCTD, or ACTD <i>DARPA Info. Proc. Tech. Office</i>	Unit personnel <i>Platoon Commdrs</i>	• Tight interaction between all roles • Customer and OPO are typically the same organization or under the same local chain of command • Customer/OPO is primary sponsor of funding and resource advocacy • User base likely known • Service Owner may be external entity but service development initiated by local personnel
C2 Common <i>Example: JOPES</i>	Regional commanders <i>Commdrs, COCOMs J3</i>	Operational organization responsible for functional process <i>JFCOM J3</i>	PEO <i>DISA GCCS-J PM</i>	Unit personnel <i>Planners in COCOMs</i>	• OPO more likely to be functional organization • Customers likely to represent regional commanders • Likely to include many unanticipated users • Service owners from PEOs and services provided via programs of record (PoRs)

UNCLASSIFIED

Tier	Typical Actors for C2 CONOPS Roles				Comments
	Customer	OPO	Service Owner	User	
C2 Infrastructure	Program manager	Functional organization responsible for defining rules for infrastructure service operation (e.g., access control policies)	PEO	Unit personnel (if user-facing service) or another service (if machine to machine)	- Customers and Users are predominantly other Service Owners (for machine-machine interfaces) - Service Owner is most likely advocate for resources - May be more service-level assertion than negotiation - Multiple OPOs may be involved; OPOs may have enterprise-level responsibility to define standards and rules
<i>Example: Authorization Service</i>	<i>TIGR PM would choose to build on authorization service</i>	<i>NSA or Army accrediting organization would work with authorization Service Owner to ensure correct policy implementation</i>	<i>DISA or MilSvc PM is responsible to provide the service</i>	<i>TIGR log-in would use the service</i>	

4. An Example of CONOPS Roles for C2 Services

To illustrate the four key roles, Customer, OPO, Service Owner, and User, we have depicted an example of provisioning of a weather service and use by a C2 service in Figure 12. CONOPS analyses that involve all four key roles, even if performed by the same actor, enable a thorough assessment of the life cycle responsibilities.

In this example, weather sensor operators are the Service Owners of services that provide weather data for others to use. The Air Force Weather Agency (AFWA) is a Customer of the sensor data access services as well as the Service Owner of a service that provides weather forecasts to commanders in their AoRs. The CENTCOM JFACC is the Customer of an ATO Generation Service that actually uses the AFWA provided weather forecast service. Air Combat Command is the OPO for the generation process since they have determined what sources and processes will be executed in support of the mission. In this case, the OPO is different from the Customer since the Customer establishes the performance requirements for the Command while the OPO has established a consistent process with identified authoritative sources for use across all commands. Electronic Systems Command is the ATO Generation Service Owner because they ensure the service is provided for use. Finally, planners are Users of the ATO generation service.




Figure 12. Example of CONOPS Roles

Some key observations from this example include:

- The same organization or person can and will play different roles within the SOE and for C2 services. The nature of services is such that multiple services can be exploited in the conduct of any mission and an organization's role may change from provider to consumer.
- The Customer of a C2 mission service is typically a commander (or his/her representative); he/she is responsible for the mission in theater. In some cases, the Customer and the OPO will be the same organization or person. Typically in C2 local service cases, the Customer stipulates the performance requirements and defines the mission process.
- Services lend themselves to a many-to-many relationship between a Service Owner and service Customers and Users. Similarly, the OPO may have multiple services and sources to choose from in building the mission process. The many-to-many relationship means that OPOs, Customers, and Users can take advantage of available services without a new development program.
- There is a strong demand for OPOs and Customers to collaborate on service needs. Service Owners need to respond to both Customer and OPO

UNCLASSIFIED

requirements; collaboration across all three roles increases probability of successful service implementation, delivery, and operation.

- Service Owner is a provider role; the Service Owner is responsible to the Customers, OPO and Users to deliver the service. The Service Owner takes responsibility for managing the efforts of a service provider such as a contractor. Typically a Service Owner, such as the weather sensor operator, will be providing the service to a variety of Customers and Users in the C2 community and beyond.

5. Responsibilities in the Service Life Cycle

The Department has an ongoing effort to tailor the ITIL processes for Defense use. Consequently, in this study, we did not attempt to analyze every activity and assign responsibility for activities in the life cycle. Instead, we chose to limit our analyses to some high priority activities related to goals of the Services Strategy and the DoD Net-Centric Data Strategy (Table 10).

Table 10. Analyzed Activities by Life Cycle Stage

Life Cycle Stage	Key Activities
Plan, Organize	Identify existing services and information sources before creating new ones
	Develop vocabulary and structure for service output (i.e., COI activities)
	Identify resource needs
	Identify funding sources
Acquire, Implement	Implement service
	Obtain certification and accreditation to operate the service
	Register the service for discovery in the appropriate service registry
Deliver, Support	Negotiate service-level performance or understand service-level assertions
	Provision IT that is required to deliver the service
	Provide user support or work with existing support organizations to address user support functions
	Assess and approve changes and releases to operational service
Operate, Control	Develop access control rules for service users and content consumers
	Enforce operational control rules including prioritization
	Manage content by deciding what gets posted, who can post, rules and processes for content updating, and organizing and analyzing information
Monitor, Evaluate	Assess operational demand for services and content
	Assess service performance vs. guarantees or assertions; identify performance issues
	Measure user satisfaction; provide feedback to improve processes or service capabilities
Improve	Identify high-value improvements in processes or service capabilities

UNCLASSIFIED

The Department focuses heavily on managing the acquisition stages of the life cycle, but the SOE requires the same degree of attention be paid to the operational stages of the life cycle, while recognizing the inherent differences between service tiers. When each C2 service tier is analyzed, we can identify the life cycle roles responsible for these activities. We have not attempted to conduct a thorough RACI (Responsible, Accountable/Approver, Consulted, Informed) analysis for this study; hence, multiple roles may be assigned responsibility for each activity. The result of our analysis for the C2 services domain is summarized in Table 11.

Table 11. Key C2 Service Roles and Responsibilities

LC	Key Activities	C2 Local Responsibilities	C2 Common Responsibilities	C2 Infrastructure Responsibilities
Plan/Organize	Identify existing services and information sources needed for mission	Customer, Users	OPO, Users, Customer	OPO, Customers
	Develop vocabulary and structure for service output (i.e., COI activities)	Customer, Users, Service Owner	OPO, Users, Service Owner	OPO
	Identify resource needs	Service Owner	Service Owner	Service Owner
	Identify funding sources	Customer	OPO, Customer	OPO, Customers
Acquire/Implement	Implement service	Service Owner	Service Owner	Service Owner
	Obtain certification/accreditation	Service Owner	Service Owner	Service Owner
	Register the service for discovery	Service Owner	Service Owner	Service Owner
Deliver/Support	Negotiate service levels	Customer, Service Owner	Customer, Service Owner	Customer, Service Owner
	Provision IT for the service	Service Owner	Service Owner	Service Owner
	Provide user support	Service Owner	Service Owner	Service Owner
	Assess and approve changes and releases	Customer	OPO, Customer	OPO
Ops/Control	Develop access control rules	Customer	Customer	OPO
	Enforce operational control rules	Customer	Customer	Customer
	Manage content	Customer, Users	Customer, Users	Customer, Users
Monitor/Evaluate	Assess operational demand	Customer	OPO, Customer, Service Owner	OPO, Customer, Service Owner
	Assess service performance vs service levels	Customer, Service Owner	Customer, Service Owner	Customer, Service Owner
	Measure user satisfaction	Customer	Customer, OPO, Service Owner	Customer, OPO, Service Owner
Cont./Improve	Identify high value improvements	Customer, Users	Customer, OPO, Users	Customer, OPO, Users

Note that since typically for C2 local services the Customer and the OPO are the same, they are not separately identified in the C2 local column. In looking at the responsible roles across the life cycle, it is clear that Customers are heavily engaged in the later life cycle stages, such as operate and control.

The next challenge is to identify candidate actors for these roles for each service tier. It would be impossible for this study to identify specific organizations or personnel

UNCLASSIFIED

responsible for each role for every C2 service. Table 11 indicates that there is only minimal difference between the roles and responsibilities across the tiers, but in looking at the candidate actors, it is clear that there may be significant differences in assigned organizations or personnel for each role.

When we combine the typical actors in Table 9 with the life cycle role responsibilities in Table 11, we can develop a summary of the principal actors across the life cycle for services in the C2 domain (Figure 13). As the figure illustrates, the person or organization that has the primary responsibility during each phase of the life cycle may change. While all actors are involved and collaborating across the life cycle, one entity has the most important role based on the primary activities during that phase. We also show representative actors if the same life cycle model was applied to enterprise services.

C2 Domain Services

	C2 Local	C2 Common	C2 Infrastructure	Enterprise*	Enterprise* (core enablers)
Plan/ Organize	Local cmdr control	Regional cmdrs, process owners	PEOs/PoRs with process owners	CPMs	DoD CIO
Acquire/ Implement	Local cmdr, service owner	PoRs		PEOs/PoRs	PEOs/PoRs with process owners
Deliver/ Support					
Operate/ Control	Local cmdr control	Regional cmdrs	Global, regional, or local cmdrs	Global or regional cmdrs	Global Cmdr
Monitor/ Evaluate		Regional cmdrs, process owners	Cmdrs, PoRs, process owners	Global or regional cmdrs, process owners	Global Cmdr, PoRs, DoD CIO
Improve				CPMs, cmdrs, process owners	

*Shown only to demonstrate application of CONOPS to enterprise services

Figure 13. Summary of Principal Actors for C2 Services Across the Life Cycle

Realizing C2 services in the SOE requires many entities and organizations working together across the enterprise. Managing and governing the plethora of service-related activities require a governance approach that allows the services to be grouped in ways to enable a portfolio view, ensuring someone is looking across related services to maximize efficiency and effectiveness.

E. GOVERNANCE IN C2 SERVICES DOMAIN

Many services will exist in the DoD SOE and an overall portfolio governance approach should be in place to manage and oversee implementation and operation. The scope of the DoD or even the C2 community is such that a monolithic approach to portfolio management is impractical.

1. C2 Services Portfolios

A portfolio approach to governance of services in the C2 domain enables specific portfolio managers to maintain a broader view of services available, needs to be fulfilled, and status of ongoing operations. For example, when we look at C2 local services, it is clear that many services are operating in support of different C2 missions in different theaters. Hence, we would expect there would be multiple portfolios of C2 local services. Conversely, for C2 infrastructure services, it is critical that those services be engineered so that collectively they provide the capability for mission services to interoperate within C2 and across other mission domains. In that case, we would expect that a single C2 infrastructure portfolio would include the C2ISF. Figure 14 illustrates the potential C2 service portfolios. At the local level, there will be many portfolios. At the common level, there will be regional portfolios and a global portfolio. Finally, C2 infrastructure services should be governed as a single portfolio.




Figure 14. C2 Services Portfolios

Each of these portfolios requires a governance structure to provide oversight and management of the service activities within the portfolio. Figure 15 illustrates some

governance characteristics. Local C2 service portfolios will most likely have minimal formalized governance structures. C2 common portfolios will have differing governance structures depending on the reach and scope of the portfolio. Due to the dependence of the other C2 services portfolios on C2 infrastructure, that portfolio will need to have strong governance. C2 infrastructure, much of which will exist as part of the run-time elements of the C2ISF, is critical to enable interoperability across the C2 domain.




Figure 15. C2 Services Portfolios Governance Characteristics

2. Portfolio Management Responsibilities

To assess governance actors and structures, we need to examine the responsibilities that fall to each portfolio manager. Regardless of the formality and/or structure of the portfolio governance, the services portfolio manager has several key functions to fulfill. The service portfolio manager has a responsibility to:

- Evaluate, advocate for, and monitor the services within the portfolio
- Oversee and enforce compliance of actors in service life cycle roles.

In effect, the portfolio governance revolves around understanding the service states, and ensuring the service actors (OPO, Customer, Service Owner, Users) are all fulfilling their responsibilities.

In the evaluate/advocate/monitor role, portfolio governance advocates for resources for services within his/her portfolio—whether it is to create new services, enhance existing services, or expand access to operational services. The portfolio manager should track what services are being planned, deployed, and operated; and

understand what services are being used, as well as the performance levels so advocacy is appropriately framed. In this role, the portfolio manager should look for opportunities to promote services from local to common or common to enterprise. Clearly, to fulfill this role, any portfolio manager must have access to information about services on the networks. Current DoD attempts to manage even single services on the networks are hampered by a lack of broadly visible information about operational services.

In the oversee/enforce role, the portfolio manager needs to ensure that a Service Owner, who understands his/her responsibilities across the life cycle, is identified for every service. The portfolio manager also needs to make sure that OPOs are identified, and if not, that the responsibilities of the OPO are being carried out by another actor (i.e., Customer). This implies the portfolio manager needs to ensure that each service is examined from the perspective of “how can this be used in our operations,” or “how can our operations be improved by using various services,” and “how do we institutionalize this new process (e.g., make it part of TTPs).” Finally, each portfolio manager must identify an organization or entity to provide technical and engineering support to ensure the C2ISF standards and rules are followed to enable federation; C2 and enterprise infrastructure are appropriately used/reused; and deployed computing platforms are rationalized. While each Service Owner will have technical resources to fulfill his/her responsibilities, the coordination across the portfolio requires a broader perspective.

Given the portfolio responsibilities, multitude of service life cycle actors, and the dynamic nature of the SOE, it is important that the governance concept allow for “lightweight” structures to support the agility of the local environment, while having the formality needed to enforce infrastructure considerations.

3. Proposed C2 Services Portfolio Governance Concept

Based on the need for varying layers of structure, we are recommending a governance concept (Figure 16) that assigns an individual with overall governance authority for each C2 service portfolio. Each portfolio governance authority will then designate the individual(s) or organization(s) to execute his/her responsibilities. The portfolio governance authority will also stipulate the governance processes to be used for his/her portfolio. In the case of the C2 common-global and C2 infrastructure portfolios, we have recommended individuals who should be designated by the governance authority to perform monitoring, evaluation, and advocating for services within that portfolio. In

UNCLASSIFIED

addition, we suggested specific organizations to perform technical oversight and support on behalf of the portfolio authority.

C2 Domain Services

Portfolio Layer	Portfolio(s) Governance Authority	Monitor/ Evaluate/ Advocate	Ensuring Life Cycle Compliance			Applicable Governance Process
			Specific Service Owner	Specific Service OPO	Service Technical Oversight	
C2 Local	Local Cmdrs	Local Cmdrs to designate responsibilities within portfolios				Defined by Local Cmdrs
C2 Common-Regional	COCOM Cmdrs or Component leads	Portfolio governance authority to designate responsibilities within portfolio				Prescribed by Governance Authority
C2 Common-Global	ASD/NII	C2 CPM	As designated ASD/NII		JSIC	Prescribed by ASD/NII
C2 Infrastructure	ASD/NII	C2 CPM	As designated by ASD/NII		DISA	Prescribed by ASD/NII
DoD Enterprise* (other)	DoD CIO	As designated by DoD CIO				Prescribed by DoD CIO
DoD Enterprise* (core enablers)	DoD CIO	As designated by DoD CIO		DISA	DISA	Prescribed by DoD CIO

** Enterprise governance construct shown only to demonstrate application of C2 model to enterprise portfolios*

** Enterprise governance construct shown only to demonstrate application of C2 model to enterprise portfolios*

Figure 16. C2 Services Portfolios Governance Concept

In the case of C2 local portfolios, the local commander is the logical governance authority for all services developed and deployed for use by that command. To provide as much agility as possible, the local commander can designate or delegate governance responsibilities and authorities using his/her locally defined process for managing the portfolio. Note that a local command may also be using C2 common and C2 infrastructure services as well as enterprise and other domain services, but the local commander does not have governance responsibility and authority for those services.

In the case of C2 common-regional, a logical governance authority is the regional Combatant Commander or alternatively, a component lead. Consider the example of a C2 common-regional portfolio where a Combatant Commander is the governance authority. In that capacity, the commander would designate who will conduct monitor/evaluate/advocate functions on his or her behalf. In addition, the Combatant Commander will designate the individual or organization responsible to be the Service

UNCLASSIFIED

Owner and OPO for the services under his/her governance authority. Finally, the Combatant Commander will identify the individual or organization responsible to provide technical oversight of services in his/her portfolio. The Combatant Commander can prescribe his/her governance process.

For comparison, the ASD(NII) is the portfolio governance authority for the single C2 common-global portfolio. For that portfolio, the C2 capabilities portfolio manager (CPM) is responsible for monitoring, evaluating, and advocating for the portfolio. The ASD(NII) can designate the Service Owners and OPOs for services in the portfolio. We recommend that the Joint Systems Integration Center (JSIC) provide technical oversight for the C2 common-global portfolio.

The C2 infrastructure portfolio, with the inclusion of the C2ISF, represents a special governance case. Here, we propose the ASD(NII) as the governance authority, and the C2 CPM to monitor, evaluate, and advocate for infrastructure services. The C2ISF (see Section III) represents a logical collection of design-time elements (artifacts, templates, rules, and protocols), and a run-time infrastructure to allow federation across the C2 mission services and with other services in the SOE. However, it is likely that the C2ISF will be provided by multiple organizations. Hence, the ASD(NII) must stipulate Service Owners for each of the services of the run-time infrastructure, and the “owners/managers” of the design-time elements. For this portfolio, it is critical that overall technical oversight be provided by a single organization (e.g., DISA) to ensure that engineering and technical implementation and operation trades are conducted in a holistic manner.

Finally, at the bottom of Figure 16, we illustrate a potential governance authority for enterprise service portfolio(s) as an example of implementing the same C2 service portfolios governance concept. The DoD CIO is currently developing governance approaches and structures for enterprise services.

For purposes of looking across the entire C2 domain of services portfolios, the ASD(NII) has the responsibility to designate portfolio governance authorities. As services become more mission critical or if they experience broader, more significant usage, the services will migrate into different portfolios. In that event, the receiving portfolio governance authority has to accept the responsibilities for services that move into his/her portfolio (e.g., as a C2 local service migrates into C2 common portfolio). The ASD(NII) has to establish the overall guidance for how service governance is handed

off as a service moves into a different portfolio. Finally, as the C2 domain authority, the ASD(NII) is also responsible to work with other domain and enterprise governance authorities to ensure successful SOE evolution.

The need for agility across the C2 domain within the SOE drives a more complex or layered approach to governance. It is unlikely that a single governance authority and a single governance process could be flexible enough for, and responsive to, the demands for mission-essential services in theater, while ensuring that needed infrastructure is in place to enable interoperability. The recommended approach recognizes the diversity of the C2 domain.

Elements of this CONOPS are already in existence or coming into practice as part of ongoing operations. There is a need to accelerate SOE management and governance processes for C2 services and enterprise services.

F. GOVERNANCE ISSUES AND CONSIDERATIONS

This CONOPS has established a foundation for managing and governing the evolution of the SOE for C2 services. Consequently, the CONOPS development focused mostly on organization and responsibilities. However, it has not been a comprehensive treatment of all issues and considerations that must be addressed to realize an SOE.

Current Departmental decision support processes, such as planning, programming, budgeting, and execution (PPBE), and acquisition, are not structured to promote or enable SOE evolution. The SOE is all about building an inter-dependent enterprise, rather than the stovepiped DoD of the past. Ultimately, governance of IT and the SOE will have to address issues such as:

- How can edge innovation be incentivized?
 - Edge innovators need seed funding with the recognition that some new ventures fail. The good ideas that spread need to attract resources and talent.
- How can we better link operational situational awareness measures of service operations to resourcing?
 - The lack of visibility into operational status of services on the networks limits situational awareness. Having SA implies that there may be a way to direct funds toward high-value, high-usage services.

UNCLASSIFIED

- How can we modify Departmental processes to promote interdependency without undue burden?
 - Multiple efforts have looked at streamlining the acquisition process to accommodate more commercially driven IT solutions. Fixing acquisition without a holistic approach to budgeting, funding, and requirements development is only a partial solution.
- How can we effectively manage risk in our multi-tiered SOE?
 - Local commands have an increased risk tolerance for IT solutions due to urgency of need. Interdependency means that a local risk can quickly become an enterprise risk. Balancing risk, agility, and control requires an operational, technical, and process framework.

A comprehensive SOE CONOPS that addresses the significant organizational, policy, and process issues requires much more development. Our recommendations in the next section represent a first step in implementing a C2 services CONOPS.

UNCLASSIFIED

(This page is intentionally left blank.)

UNCLASSIFIED

V. CONCLUSITONS AND RECOMMENDATIONS

A. CONCLUSIONS

Study observations confirm that services will be a vital means of providing C2 information support. Based on case study analyses, and development of the C2ISF and CONOPS, we have drawn the following key conclusions:

- Achieving an SOE for C2 requires long-term planning and effective oversight while supporting agility and innovation in the operational community.
 - Planning, oversight, and operations are dependent on having situational awareness information about that status of services and information products on the networks with focus on collecting and publishing usage and performance metrics.
- Implementation of C2 community services within an evolving SOE requires a time-phased “start-up” plan. Critical elements include:
 - Designation of categories or “tiers” of services and appropriate governance authorities for each category/portfolio
 - Development of portfolio-level governance processes for the identification, acquisition, and life cycle management of C2 services
 - Assignment of service management roles and responsibilities with emphasis on promoting edge-user innovation
 - Provision of CONOPS and implementation guidance to provide unity of effort for identifying, implementing, and managing information support as services within an evolving DoD SOE
 - A commitment and plan of action to develop appropriate infrastructure services in a sequence and on a timeline that supports the simultaneous development of mission-oriented services.
- An aggressive program of work to identify and develop technical approaches and relevant standards is needed to implement a federated SOE and associated infrastructure services.

B. RECOMMENDATIONS

The recommendations span guidance and implementation planning actions while highlighting near-term activities.

1. Guidance

DoD CIO, issue guidance to:

- Clarify and institutionalize service-related terminology (to include appropriate ITIL/COBIT roles and definitions)
- Adopt and advocate for a layered model for services and SOE governance that addresses agility and stability needs.
- Establish the necessary federation approaches and standards to support the evolution of a Department-wide SOE.

DoD CIO, in coordination with U.S. Cyber Command (USCYBERCOM) as the NetOps mission owner, develop guidance for measuring and publishing service implementation and usage metrics on DoD networks.

2. Implementation Plans

DASD, C3S&S, in coordination with DASD, IMIT/Deputy DoD CIO, DoD Components, and the COCOMs, initiate action to plan for and implement C2 services and the tiered C2 services structure and governance processes. Critical elements of such a plan should be:

- Designation of types or categories of services (i.e., service portfolios) and appropriate governance authorities for each category/portfolio—the service portfolio construct discussed herein should be considered.
- Determination and promulgation of portfolio-level governance processes for the identification, acquisition, and life cycle management of C2 services—assignment of service management roles and responsibilities with emphasis on promoting edge-user agility and innovation as discussed herein should be considered.
- A commitment and plan of action to evolve appropriate infrastructure services in parallel with mission-oriented services.

3. Near-Term Implementation Actions

DASD, C3S&S, in coordination with DASD, IMIT/Deputy DoD CIO, consider:

UNCLASSIFIED

- Using the emerging JC2 initiative (i.e., NECC replacement/follow-on) as a pathfinding effort for services-based development and acquisition. Consider aiming efforts at timely, effective responses to user-driven requirements emerging from operations in SW Asia.
- Codifying relevant C2 service and SOE implementation activities in the DoD C2 implementation plan.

Additional near-term actions should be carried out at three levels: (a) DoD Enterprise; (b) C2 Common; and (c) C2 Local.

a. DoD Enterprise

Emphasize core enablers:

- Implement and operate registries, directories, catalogs, search capability, authentication service, and feedback and measurement services
- Develop an authorization service for implementation by PoRs and local capabilities
- Increase rapid provisioning of Web-based computing capabilities for edge (local) use (e.g., cloud computing, Web development tools)
- Provide guidance to PoRs and local service owners to implement visibility/registration requirement
- Require PoRs to resource engagement with local users and to capitalize on what is occurring at the edge.

b. C2 Common

Emphasize Web-enabled access to existing PoR data:

- Implement data access services to major fielded data stores maintained by PoRs
- Register data access services (identities and network locations)
- Register vocabularies of data sources; participate in communities of interest (COIs) as needed to rationalize vocabularies; develop mediators where appropriate
- Advertise community content in widely available catalogs
- Develop C2ISF and require adherence for PORs.

c. C2 Local

Emphasize innovation and registration:

UNCLASSIFIED

- Identify priority data for which edge users want Web-enabled access and widely publish their requirements for data
- Register in-use or developing local capabilities in a registry/catalog so that their existence can be discovered
- Work with enterprise and common governance authorities to ensure connections to enabling infrastructure are in place so local edge users can find, access, and exploit common services and data sources.

UNCLASSIFIED

Appendix A
C2 SERVICES CASE STUDIES

UNCLASSIFIED

Appendix A

C2 SERVICES CASE STUDIES

A. FOUR CASE STUDIES DESCRIBED

We examined four existing C2 capabilities that are operational, currently available on SIPRNet, and amenable to service-oriented evolution. Examples of pure and purposeful service implementation remain few and far between in DoD. Accordingly, the cases selected are not paragons of service development and operations. Rather, they provide important clues and observations concerning current behaviors in fielding, operating, and evolving C2 capabilities, which are relevant both to the Department's overall migration toward an SOE and to the implementation of C2 services in particular.

The analysis included an initial categorization of these cases loosely based on the Joint Capabilities Assessment (JCA) C2 taxonomy, which affirmed that they are within C2 boundaries. There are multiple JCA-described C2 capabilities inherent in each case. Indeed, few "course grained" service-like activities of the kind found operational in the field today equate to a single JCA capability except at the highest levels of generality. However, two of the cases are almost entirely within the top-level JCA C2 category "understand," which includes collaboration and situational awareness. The other two lie mainly in the "planning" category although linked importantly to the JCA's "organize," "direct," and "monitor" functions. Our research indicates that clear-cut and consistent service categorization continues to be a major challenge throughout the Department.

1. Specialized Tactical Situational Awareness Services

Situational awareness here refers to knowledge of previous observations, lessons-learned, issues, and solutions experienced by others in similar circumstances. IT is actually a family of DARPA-sponsored, Web technologies that have been inserted into Iraq and Afghanistan over the past several years with positive results. The most noteworthy of these are TIGR and CIDNE. These SA engines are essentially, electronic pass-down logs "on steroids" that allow operators (plus some supporting rear-echelon analysts) to share facts, suppositions, lessons-learned, outstanding issues, and solutions via geo-referenced multimedia. While CIDNE and TIGR have some significant

UNCLASSIFIED

differences in implementation, they are becoming progressively interconnected among themselves and with other operational capabilities. Both cases demonstrate the following:

- The relatively simple, inexpensive, user-friendly, and COTS-with-minor-enhancements nature of successful edge innovation
- What start-up, expansion, and rapid evolution of capabilities entail in a "hot war" operational environment.

2. Command-Level Situational Awareness and Collaboration Services

Essentially, this case examines CPOF, which was originally a DARPA technology demonstration in the late 1990s that transitioned to an Army PoR in 2006. It has been tightly federated with the Army's MCS and consumes data from the GCCS. CPOF is basically a real-time collaboration technology with a powerful military mapping feature. It constitutes a virtual "sand box" where commanders can publicly depict situations, plan potential courses of action, offer ideas, refine tasking and approaches to its execution, and self-synchronize plans. Its users are flag officers, unit COs, and their senior staff at several echelons. In geographically dispersed locations such as Iraq and Afghanistan, CPOF allows commanders to share battle update assessments, obtaining real-time feedback on a regular basis. They can communicate, collaborate, and brief each other without leaving their operations centers. This enables synchronization of activities at the tactical and strategic levels while avoiding the hazards of travel.

The CPOF case suggests that widely adopted collaboration capabilities, which are technically open to federation with new data sources, may prove to be one of the most powerful organizing forces in the Department's emerging SOE. In addition, this case illustrates that, once an experimental capability is recognized as important to current operations, incorporating it into an existing PoR has many immediate benefits. This includes enabling Initial Operational Test and Evaluation (IOT&E) short cuts; piggybacking on mature training and logistic activities; and focusing already programmed resources on evolving federations among important capabilities.

3. Air Operations Tasking and Control Services

This case deals with key C2 business processes, organizations, and IT used by COCOM Air Component Commanders to manage air operations on a regional basis. From an IT standpoint, the principal focus of this case is on TBMCS resident in globally

UNCLASSIFIED

distributed AOCs, which has a decades-long developmental and operational history. TBMCS is the system of record for generating and promulgating the Air Tasking Order (ATO). The TBMCS program management activity is an element of the Air Force's ESC at Hanscom AFB. This, along with the system described in the next case study, provide IT for essentially large, complex, scheduling services accompanied by some capacity to monitor event execution, formulate adjustments, and proffer control inputs. TBMCS is actually an integration of six to eight "legacy" systems or subsystems, and today the business process engages a half dozen Web technologies that are not part of the formal program to improve information sharing. This case illustrates how high-tempo around-the-clock operations can drive supporting IT to rapidly evolve through local and community innovation, using both PoR and other resources to become far more efficient internally and far more responsive in terms of product packaging and delivery. It also suggests that a mature system implementation, evolved over many years to address all steps in a complex business process, might be transformed into a composite service, an amalgam of multiple service offerings, each of which can support a more diverse customer base.

4. Force Deployment Planning and Execution Management Services

JOPES is the Department's principal tool for designing, monitoring the progress of, and managing force deployments, generally large-scale. The Defense Information Systems Agency (DISA) manages development and computing services support for operations. It is considered part of the GCCS federation of technologies, over 200 GOTS and COTS capabilities that are currently operational at the national level and in a variety of other flag-level command centers. Although JOPES features powerful database synchronization technology supported by joint strategic servers located at enterprise computing centers, this system executes a highly complex business process that requires accurate data entry from multiple, widely distributed human and machine sources in multiple theaters. Under CENTCOM staff (G-4) auspices, this system is actively engaged in support of logistic evolutions for both Iraq and Afghanistan, and there is strong TRANSCOM involvement. JOPES data are a vital part of formulating large-scale movement orders, but the system is also employed to manage small-unit and materiel transportation. JOPES, by itself, cannot perform key expected, commonly understood, commercial transportation functions such as "track shipment," which requires data from a Global Transportation Network (GTN) confederate. This case illustrates circumstances

UNCLASSIFIED

and considerations surrounding how a large, complex, global system can fit into the service-oriented paradigm providing more comprehensive and flexible support for key military business processes.

B. OBSERVATIONS BASED ON CASE STUDIES

We have made some observations from detailed analyses of the individual case studies (see Section C).

The family of capabilities that includes TIGR and CIDNE is basic in terms of the business process it supports, yet rich in responsiveness to user requirements. Maps, forms, and other common user interfaces have been tied together via readily accessible databases that are distributed in some instances and centralized in others. TIGR capabilities directly support the lowest echelon decision makers, i.e., personnel engaged in recon patrol activities. CIDNE, which was designed primarily to manage brigade and above Sons of Iraq (SoI) contacts and coordination efforts, has evolved into the IT centerpiece of counter-Improvised Explosive Device (C-IED). These ostensibly simple capabilities, which are friendly graphic user interfaces coupled with impressive multimedia storage and retrieval engines, required several years to field once relatively mature in the lab setting. After initial experimental fielding, these capabilities experienced rapid evolution that responded directly to ongoing operational experience. Although tailored for data-sharing conditions far more primitive than the Internet, Web technologies were incorporated to support the possibility of rapid user base expansion, which has occurred.

CPOF, as a map-centric battlefield management capability based on collaboration technology, is well suited to support Information Age behaviors. It has often been described as a “John Madden whiteboard” that allows senior officers to share virtual game plans on the battlefield. It is one of the best examples of a technology demonstration that has gained widespread acceptance and that has been successfully melded with a mainline PoR capability (MCS). Moreover, demands of CPOF user-collaborators are stimulating technical connections with more data sources. This effort received the 2009 Network Centric Warfare Award for Outstanding U.S. Government Program. CPOF is currently deployed to SW Asia with an associated in-theater training program, mainly for use at Army division and brigade levels, but it is rapidly becoming a mainstay of joint/coalition C2.

UNCLASSIFIED

JOPES (force deployment case) and TBMCS (air operations case) represent the traditional, “Big A” acquisition requirements, development, and fielding process as updated with spiral methodology and selected Web technologies. The currently fielded versions represent more than 20 years of evolution. Both systems have become confederates of the GCCS. These cases are similar in that they exhibit two important user behaviors:

- Planners (command staff) input their operational requirements into the business process in order to get the events their commands need scheduled and formally assigned to specific units
- A whole host of command users-recipients of the process’s authoritative scheduling products (op orders, task orders, control orders) provide organizational assignments and control direction covering what commands are expected to do when, where, and with what.

These two behaviors point to a wiki-like characteristic that most C2 capabilities exhibit; viz., the primary users of the service are also the primary content producers and consumers.

Despite their similarities, these scheduling/planning cases seem to reflect different models from the service implementation approach and roles standpoint. The air operations case is far more decentralized with many site-specific “tailored” technical/business process arrangements in the various regional centers and no large-scale aggregation activity. JOPES is intended to be a more uniform global capability that can affect planning at many echelons in a coordinated manner from national on down. Both systems support complex industrial (assembly line) business processes that have been defined in detail over many years. JOPES exemplifies a powerful global materiel and personnel movement-scheduling capability that has nonetheless *not* been designed to address the entire package-schedule-transport-track-manage business process. TBMCS is a development effort that attempts to address the whole job of air operations planning as well as orders formulation and promulgation. To do this, a half-dozen, formerly independent system capabilities have been integrated in the traditional, tightly coupled, proprietary style. In the new service-oriented paradigm, those TBMCS modules would each become services that are capable of supporting not only air operations planning but also other business processes. JOPES, on the other hand, would become one atomic service within a large composite service. Enterprise, C2 common, and other community

UNCLASSIFIED

common services would be incorporated, with a strong potential for C2 local service innovation to better exploit the planning products.

The details in these cases suggest that much of the IT supporting current operations is being developed and provisioned from multiple sources that are essentially uncoordinated at the enterprise level and only marginally coordinated within specialized warfighting communities. These sources include PoR products (generally large system hardware and software combinations requiring long-term development), demonstrations or prototypes that use cutting edge technology, and well-proven COTS (especially common Web technologies). Thus, present day IT provisioning is characterized, almost in its entirety, by what might be termed “natural growth,” which, if shaped by a modicum of governance could be far more effective and efficient. But Information Age governance, requires transparency. There is scant high- or even middle-level ability to detect and track the particulars of evolving operational C2 IT support much less to intervene to achieve efficiencies and synergies.

Given the present lack of insight, local commanders acting in the context of theater C2 arrangements are in the best position to understand and control what is going on in their respective operational environments. They have manifested the ability to quickly assemble new capabilities by using experimental products, commercially available Web technologies, and small cadres of forward-deployed supporting engineers and trainers. Our studies show that with some granular knowledge of specific theater needs and conditions, new locally innovated C2 Web-enabled capabilities and services can be provisioned to meet emerging mission requirements. In addition, once determined effective, their usage can be rapidly expanded.

Our analyses across these case studies led to some general observations for NCSS implementation for the C2 community:

- To implement C2 capabilities in an SOE, DoD requires agile and collaborative governance that embraces the full range of IT engineering and operations activities. The governance must accommodate highly variable C2 node-specific arrangements arising from local needs, clusters of C2 capabilities, which form common C2 IT, organizational, or procedural constructs, and enterprise infrastructure, which must be robust and stable with carefully planned changes due to large-scale dependencies. *(Case studies show that critical implementation action is ongoing in all three arenas and that all three*

UNCLASSIFIED

perspectives are needed to make agility versus stability tradeoffs and prioritization.)

- DoD governors need a robust capability to know in real time what specific IT is actually operating on DoD networks and to monitor its usage and performance. This requires instrumentation to collect metrics and feedback mechanisms that allow users to publish comments on the IT products and services they employ. *(Case studies highlight requirement to support and build on what is successfully deployed and operational. Improvements must chase demonstrable utility. Implementation cannot significantly disrupt existing capability. This requires intimate knowledge of the current operational environment.)*
- The most effective control points appear to lie within the operational chain of command where C2 facilities, available IT, content, and TTPs are constantly being assembled and adjusted to answer pressing requirements. *(Case studies reflect operational chain of command making final determinations on what capabilities to use or not use and where to invest in improvements in theater.)*
- DoD needs new processes for acquiring, managing, operating, and continuously improving C2 information services. The dynamic nature of an SOE drives a requirement for faster processes with more transparency with crisp lines of authority and accountability. *(Case studies show that IT support generally requires rapid improvement after IOC. This entails short-fuse, authoritative decision making to identify technical options, apply funding, and engage engineering/training/logistics support in both rear and forward areas to implement enhancements.)*
- Knowledgeable developers with only moderate resources can successfully engage in theater to formulate and implement significant C2 capability improvements. Local innovation can bring real and timely benefit to the warfighter. DoD authorities at various echelons can choose to accommodate and encourage that behavior or to repress it. *(Case studies show that significant local innovation will occur and succeed in delivering valuable capabilities to warriors regardless of rules. Enemy forces unconstrained by bureaucratic roadblocks will exploit commercial innovation faster than highly formal acquisition processes.)*
- Much of the client-server and other “system” capability currently deployed is amenable to rapid and relatively inexpensive service-oriented adaptation. *(Case studies show that capabilities developed over decades through the “Big A” acquisition approach, once fielded, are being significantly enhanced by Web technologies.)*

- Services that emerge through local innovation can find a programmatic home either by being adopted by an existing PoR or by governors at the common or enterprise level creating a new PoR to provide resources for them. *(Case studies show that locally developed capabilities need more robust support processes and resources once they become more widely accepted.)*

C. CASE STUDY—SPECIALIZED TACTICAL-LEVEL SITUATIONAL AWARENESS SERVICES

Key Takeaways:

- C2 local service-based capabilities are arising in-theater with greater frequency due to the proliferation of low-cost Web-based technologies, counter-insurgency mission demands, and the distribution of C2 responsibilities to numerous small units down echelon.
- C2 local capabilities that experience rapid growth because of strong user acceptance and demand will need to be migrated to C2 common capabilities via PoRs to provide more robust operations, sustainment, and support. Existing logistics and training activities will embrace innovative capabilities or initiatives when they are widely and rapidly adopted in the operational setting.
- Capability developers can significantly increase the effectiveness of their service implementation efforts by maintaining some forward presence and tight linkage with field users and by aligning improvements with unit rotation cycles.
- Deployed capabilities can be effectively repurposed in-theater. Their content, once widely published using service technologies, can be exploited and augmented with value-added data by globally distributed activities.

1. Tactical Ground Reporting (TIGR) System

TIGR is a case study that illustrates local services expanding to C2 common service. TIGR's graphical, map-referenced user interface is highly intuitive and allows multimedia data such as voice recordings, digital photos, and GPS tracks to be easily collected and searched. The system also uses a state-of-the-art data distribution architecture that minimizes the load on tactical networks while allowing digital imagery and other multimedia data to be exchanged. TIGR provides users with a sort of pass-

UNCLASSIFIED

down log “on steroids” in which experiential information, multi-source empirical observations, and analyses are continually collected, stored, shared, and organized for discovery. This functionality is made available as a service to users operating in similar circumstances.

TIGR’s mapping capability, which links still imagery, audio, video, and text to geography, offers a multi-media information product that junior officers and NCOs can study before patrolling and adding to upon return.

This multimedia post-patrol Web-logging activity records specific information about individuals, mission activities, facilities, equipment, and dangers encountered during operations that is helpful to ensuing teams. Text entries are made by junior officers or leading non-commissioned officers with detailed geographic data in forms, map annotations, plus any supporting imagery or even audio appended. Each entry is dated and gives enough information to clearly communicate whatever threats or other problems were encountered, sequences of events, and answering measures enacted during a given patrol. Information can be routine, or it can include tactical action and casualty-related details. This kind of log, if properly maintained, is an invaluable tool, enabling patrol leaders to understand what can happen in any given locale and how others have responded. Traditionally, pass-down logs are not considered formal documents and therefore had no designed-for-the-purpose IT support.

By offering TIGR as a service, soldiers can aggregate experiences, conduct trend analysis in theater or even globally, assess performance, track recurring problem-solution sets, and highlight unanswered challenges for any number of innovators to offer solutions. Warfighters are able to improve their situational awareness and to facilitate collaboration and information analysis. Using TIGR, patrol leaders can conduct company- and patrol-level intelligence preparation of the battlefield (IPB) both pre- and post-mission. By clicking on icons and lists, they can see the locations of key buildings (such as mosques, schools, and hospitals) and retrieve information (such as location data on past attacks, geo-tagged photos of houses and other buildings, and photos of suspected insurgents and neighborhood leaders). They can listen to civilian interviews and watch videos of past maneuvers. The aim of TIGR’s developers was to leverage the power of multimedia information and reduce vertical stovepipes that slow or diminish the ability to

UNCLASSIFIED

share best practices rapidly. TIGR was expressly created to support horizontal information sharing at relatively low echelons of U.S. ground force operations.

“It is a bit revolutionary from a military perspective when you think about it, using peer-based information to drive the next move.... Normally we are used to our higher headquarters telling the patrol leader what he needs to think.”

—Quote from staff officer in First Brigade Combat Team on using TIGR

More than a thousand TIGR instances are supporting users at the company level and below for planning patrols. The program is now scheduled for fielding in most of the brigades in Iraq and many C2 facilities in Afghanistan by 2010. In the meantime, the software and its capabilities are continuing to evolve.

Troops in theater are creative in their approaches to the challenge of generating, organizing, storing, and sharing data to support C2 decisions. One of the more innovative and well-resourced approaches was the development of CavNet, a precursor to TIGR. Developed entirely in-house by the 1st Cavalry, CavNet was essentially a collection of blogs and forums that allowed junior leaders down to the squad level to share information with one another across the entire division.

Troops in theater today are highly creative in using technology to manage and share information. One of the more creative and well-resourced approaches is essentially a collection of blogs that allow junior leaders down to the squad level to share information with one another across an entire division.

Although CavNet and its successors improved information sharing, they lacked a robust database for multimedia and reports, and a friendly, well-integrated human-machine interface. Seeing a clear operational need, soldiers from the 1st Cavalry teamed with DARPA to work on what would be later called TIGR. A DARPA PM began interacting directly with soldiers returning from Iraq and was able to refine her appreciation of the operational need. She also confirmed that battalion-level leaders were open to quickly getting useful tools for counterinsurgency in theater.

A team of programmers was assigned to work directly with soldiers in developing specific TIGR features in both the “must have” and “nice to have” categories. As new

UNCLASSIFIED

versions were developed, 1st Cav personnel tested them in exercises at the unit's home station and at Army training centers. By working directly with soldiers who would actually use the software on deployment, developers were able to meet a 1-year development schedule and to create a system meeting or exceeding most user requirements.

The rapid fielding schedule and unorthodox development method meant that TIGR had not gone through all the normal development channels. Compelling operational needs demanded its presence in theater, and commanders in the field made the decision to employ it. TIGR was not a program of record. It did not have Army acquisition support for fielding, and it was not initially sanctioned for use over wireless tactical networks. In an initial compromise agreement, the system would only be used within 1st Cav, and network use would be limited to a few base camps in Iraq. In addition, resource allocations were limited. However, there was enough backing for the capability within 1st Cav so that the division helped fund the program, developed standard operating procedures (SOPs), and encouraged its use down to the squad level. To ensure the system could be maintained in Iraq, DARPA teamed with the Rapid Equipping Force (REF), which provided critical support and funding to send a training and engineering team to theater.

Because of its popularity with troops in the field, TIGR gained support within the larger Army establishment. In addition, the capabilities that make it popular, principally the ability to share multimedia information across all echelons, are somewhat at odds with SOPs for sharing classified information. In addition, TIGR does not easily interoperate with the mainline C2 systems at the battalion level and above. Continuing collaborative development was needed to overcome both procedural and technical roadblocks to this capability's effectiveness.

The continuing improvement process, the foundation of agility, is less orderly than with rigidly engineered systems because of the large active user base, which provides feedback to requirements and priorities. Developers listen to what users want and try to build to modern open standards, while introducing rapid, responsive changes without disrupting operations.

These circumstances led decision makers in the operational and acquisition chains to collaboratively deploy in-theater teams of field service representatives and system

UNCLASSIFIED

engineers responsible for fielding and maintaining TIGR instances throughout Afghanistan and Iraq. These “forward” managers handle all logistics and personnel requirements including living quarters, vehicles, work hours and tasks, team goals, and special projects as requested by CONUS-based program management. The program manager coordinates actions with and supports theater government representatives. The program manager ensures a helpdesk is maintained and a robust training program is executed in support of theater operations.

2. Combined Information Data Network Exchange (CIDNE)

CIDNE was developed by a small team of software engineers working directly with troops in the field to fulfill pressing operational needs. It was sponsored by DARPA in collaboration with the Army’s III Corps, and championed by CENTCOM. It is a Web-based system with special counterinsurgency enhancements such as an indigenous “leadership engagement” tool. Like TIGR, CIDNE incorporates COTS technologies in a novel way. With four major releases since fall 2006, CIDNE has become the “gold standard” for enemy IED activity reporting in Iraq with a suite of tools generally used at brigade and above. Actions are underway to expand linkages among CIDNE and TIGR instances, addressing needs of both the higher-echelon staff users that CIDNE serves and the patrol leaders that TIGR serves.

CIDNE began operational life as a local C2 capability inserted at a division command center and several brigade HQs in Iraq to enable information management and sharing plus storage for access by replacement units. CIDNE offers a capability for tracking three types of entities—people, facilities, and organizations—specific entities that influence operations in a region or population cluster.

The engagement tool in CIDNE was designed for anyone interacting with people, facilities, and organizations. It establishes a persistent product to familiarize organizations that are new to the operating area.

Later, the 445th Civilian Affairs Battalion (CA BN) gave the technology to their deploying troops and pushed use of the product. Their goal was to evolve the design of a platform that enhanced the civil affairs mission by providing soldiers conducting SoI engagements with a powerful knowledge store. Previously, civil affairs did not have a data-sharing capability designed to deal with SoI information. CIDNE wound up

UNCLASSIFIED

bringing together disparate communities by providing a standardized reporting framework across the intelligence and operations disciplines. This common framework allows structured operational and intelligence data to be correlated and shared as part of user-defined workflow processes that collect, aggregate, and vend information to troops and commanders in theater. In this expanded role, CIDNE's capability to track people, facilities, and organizations proved to be critical in the burgeoning effort to counter IEDs.

In an example of repurposing, CIDNE is now used primarily to support C-IED operations. Its capabilities facilitate both defeating IEDs and attacking the terrorist networks that employ them—from initial threat reporting through device exploitation, target development, and evidence tracking. CIDNE is continuing to expand support for operational missions in SW Asia.

As of 2009, CIDNE training was widely available in both Iraq and Afghanistan where importantly it has grown to include a database of all IED activity. Development and fielding of CIDNE in Afghanistan and Iraq is continuing.

Organizations collaborating to address IEDs were adapting quickly to the ever-changing threat. Not only did the Soldiers incorporate new technologies that were constantly offered to them, but also they came up with innovative ways to push the equipment to the limit and sometimes beyond the original design concept.

As mentioned, CIDNE is the database of record for IED information, and it provides users both in theater and stateside with tools to support the diverse and complex analytical processes contributing to this mission. It also constitutes a common information bridge among various communities that, while working the same problem sets from different perspectives, might not otherwise be able to share data. Specifically, the Web-enabled Temporal Analysis System (WebTAS) is a suite of generic analytical tools that allows organizations to quickly fuse, visualize, and interpret disparate sources, including databases, data streams, and other structured information. WebTAS is designed to help users uncover trends, patterns, and relationships in their data through a number of visualization options. Using WebTAS to mine the CIDNE database, users are able to obtain associated data on explosive hazard events throughout the theater in near real time. This enables them to create accurate and up-to-date explosive hazard overlays for analysis at both the tactical and operations management levels.

3. Wartime Fielding Lessons from TIGR and CIDNE

Current DoD processes for identifying and validating operational needs from the field can take longer than the typical “Web” development timelines for new capabilities. The result is technology solutions that can significantly lag the operational need.

SOPs have arisen to support wartime fielding of information technologies like TIGR and CIDNE. Emergent requirements are first documented as an Urgent Operational Need (UON), which may be joint or Service-specific services. Today, a commander at any echelon who determines that he has a gap in capability can publish an UON detailing the requirement and, in some cases, describing some proposed IT to fill the gap. Once validated by Army G-3, resources are allocated to fulfill identified needs. The process from identification of an operational need to resourcing generally takes 60 days or more, depending on the priority of the requirement.

Once an UON is resourced, a process determines if a PoR will adopt the UON or if it will become a new PoR. For IT, the Army G-3 and G-6 are involved in this process. IT used on networks must be certified to ensure interoperability with information infrastructure and systems. The process can take a year or more, and it does not seem able to prioritize requests by units actually engaged in combat.

Commanders are empowered to assume risks and allow the use of uncertified systems in theater, generally with approval of the cognizant COCOM. This exemplifies the operational chain of command authority trumping acquisition rules. However, resource constraints will usually prevent wider adoption of a technology outside of the theater. Like TIGR, CIDNE was able to reach theater C2 nodes rapidly because commanders approved “at risk” operation. Both are still bogged down with formal process demands for further certifications, recommendations, and validations. However, because of their relatively low cost and ease of use, these capabilities were able to dramatically expand in terms of usage among the operating forces without formal programmatic support.

The Army has made an effort to streamline its acquisitions process, but bureaucracy continues to slow responsive IT fielding and sustainment. The DoD remains unable to balance the importance of protecting networks, ensuring unity of effort, enabling IT platforms to share information, and rapidly supplying units in combat with

capabilities they need. Substantially improved governance is required to overcome this impedance.

4. Lessons Learned from Edge Innovation Cases

Our enemies, particularly the non-state actors we face today, have access to wireless communications, satellite communications, the Internet, computers, and many of the same software tools used within the military. In some cases, their IT is superior because it is more current due to the lack of bureaucratic constraints. They have relatively flat hierarchies, so information can flow quickly from one group to the next. Attacks can be planned and coordinated via e-mail, executed with the assistance of cell phones and reported for propaganda value on a Web site. Detailed damage assessments can be obtained through readily available media reports and blogs.

Lacking a hierarchical bureaucracy and rules for IT acquisition, small groups of terrorists or insurgents simply procure what they need and upgrade their capabilities as new technology becomes available. Information can be disseminated relatively freely without necessarily having to be cleared through a hierarchy.

The IT employed by insurgents and terrorists is easy to use, inexpensive, and readily available via commercial channels. Because these are relatively small organizations, they are able to upgrade or replace capabilities rapidly. Terrorist and insurgent organizations have shown themselves to be extremely adaptable in their use of IT: Propaganda has been spread via YouTube; digital videos have been emailed to major news networks; cell phones are used to detonate IEDs; and Google Earth has been used to plan missions and target artillery attacks. Adoption, training, and fielding can be accomplished in much shorter periods of time and with far fewer resources. To maintain an information advantage, the Department must be able to keep pace with these much smaller, nimbler organizations. In these areas today, they have an advantage over the United States and its coalition partners. Introducing greater agility and responsiveness into DoD's IT acquisition system is therefore a matter of the highest priority.

UNCLASSIFIED

Examining what has actually transpired within the Department's small business equivalent, viz., local, lower echelon C2 nodes, can identify approaches to addressing the bureaucratic, hierarchical disadvantage. The local C2 nodes are where the new technologies discussed herein have been adopted outside the traditional defense acquisition paradigm.

TIGR and CIDNE indicate that viable approaches can be formulated and institutionalized for shortening bureaucratic processes. Moreover, the U.S. military has its own advantages that can be exploited. In particular, the Department has an enormous research and development budget and direct access to developers. With TIGR and CIDNE, troops in the field identified a way to operate more effectively and worked directly with software engineers to rapidly develop the means to enact responsive improvements in IT support for C2.

The case studies demonstrate that even when software is developed quickly and inexpensively, bottlenecks in the acquisition system may unnecessarily slow fielding and further improvement. For instance, no special priority seems to be applied for the certification of networked hardware or software specifically designed for troops in current combat situations. It can take months or years of testing and evaluation to officially clear systems for the field. In contrast, the Marines and Special Forces have created fast-track certification for wartime IT. These could serve as models for broader use by DoD and other departments.

Another lesson learned is to maintain teams dedicated to rapidly upgrading software in constant contact with users in the field. TIGR and CIDNE met this challenge by having no bureaucratic middlemen in the requirements process plus a small team of programmers and support personnel dedicated to continuously upgrading the software based on direct input from troops. These personnel were split into two groups. The first group forward deployed to provide direct support to fielded capabilities and collect feedback. The second group provided updates and patches to the software from stateside as it was needed.

These case studies also demonstrate the value of user-friendly Web-based front ends and social software coupled with databases focused on individuals, organizations, and facilities in support of the counterinsurgency effort. Indeed, operational effectiveness analyses being carried out in theater include sophisticated tracking of

societal trends. Various developers independently arrived at similar Web-based solutions to these non-traditional requirements because of their low cost, availability, flexibility, ease of upgrade, relatively low bandwidth requirements, data protection on physically secure servers, and ease of distribution. An added bonus is that units can monitor events in theater from stateside while preparing for deployment.

Finally, to be successful, information must be allowed to flow freely among peers and up or down echelon, not only among units engaged in counterinsurgency operations but also among their supporting acquisition organizations. Tools and SOPs such as those engendered by TIGR and CIDNE remove some of the bottlenecks to information flow, enabling U.S. personnel to learn and adapt as rapidly as their adversaries.

D. CASE STUDY—COMMAND-LEVEL SITUATIONAL AWARENESS AND COLLABORATION SERVICES

Key Takeaways:

- Powerful collaboration capabilities, when deployed in an operational setting, will stimulate rapid innovative development of new functionality and federation with other capabilities, particularly those with high value content.
- Urgent Operational Needs (UONs) can come directly from and through collaboration of senior officers who want more information integrated into their collaboration space. Combinative behaviors like this can be greatly facilitated by an SOE.
- Once experimental capabilities are recognized as important to current operations, incorporating them into an existing PoR has many benefits. These include enabling IOT&E shortcuts, piggybacking on mature training and logistic activities, and focusing already programmed resources on federations with other important capabilities.
- Exigencies of the real-world operational environment can stimulate adoption of service-oriented technologies such as collaboration that would not necessarily occur in less-threatening environments

CPOF, a DARPA-sponsored capability based on collaboration technology, migrated into Army PoR status in 2006. It was incorporated into mainline Army Battle Command Systems (ABCS) capability MCS to improve command-level situational

UNCLASSIFIED

awareness and collaboration. CPOF applications communicate with ABCS through GCCS-A. With a user-friendly map and easy access to GCCS data, users can confer via chat or voice and annotate and share graphics. The tool enables commanders to visualize segments of the battlefield, obtain and reflect recent information about operations, and express their ideas in annotations. This common picture sharing and manipulation together with effective conferencing capabilities has proven to yield far more efficient decisions. Graphical features, used in the system, provide a better and more accurate view of the battlefield using real terrain data and GPS. Capability enhancements have been proposed by other service development efforts that capitalize on CPOF's collaboration infrastructure and mapping functionality. The CPOF received the 2009 Network Centric Warfare Award for Outstanding U.S. Government Program.

CPOF provides, in essence, a "sand box" service that offers numerous command and staff officers the opportunity to share their ideas around a common electronic representation of the battlespace in which their operations are being or will be conducted.

CPOF began as an investigation into improving C2 through networked information visualization systems, with the goal of doubling the speed and quality of command decisions. A virtual workspace is the main human interface, in which all CPOF content is a shared piece of data in a networked repository. Shared visual elements in CPOF include iconic representations of hard data, such as units, events, and tasks; visualization frameworks such as maps or schedule charts on which icons appear; and brush-marks, ink-strokes, highlighting, notes, and other whiteboard-like hand annotation.

All visual elements in CPOF are interactive via drag and drop. Users can drag structured data and annotations from one visualization framework into any other (i.e., from a chart to a table), which highlights different data-attributes in context depending on the visualization used. Most data-elements can be grouped and nested via drag-and-drop to form associations that remain with the data in all of its views. Drag-and-drop composition on live visualizations is CPOF's primary mechanism for editing data values, such as locations on a map or tasks on a schedule (for example, moving an event-icon on a map changes the latitude/longitude values of that event in the shared repository; moving a task icon on a schedule changes its time-based values in the shared repository). The results of editing are conveyed in real time to all participants in a visualization session.

UNCLASSIFIED

When one user moves an event on a map, for example, that event icon moves on all maps and shared views, such that all users see its new location immediately. Data inputs from warfighters are conveyed to all collaborators as the "natural" result of a drop-gesture in situ, requiring no explicit publishing mechanism.

CPOF gained traction as a live data alternative to PowerPoint briefings, which are nonetheless still used extensively within and among SW Asia command facilities. During a CPOF briefing, commanders can drill into any data element in a high-level view to see details on demand, and view outliers or other elements of interest in different visual contexts without switching applications. Annotations and editing-gestures made during briefings become part of the shared repository. With CPOF, the commander's SA is based on ground-truth as observed during the collaboration timeframe; and the commander can then share his intentions live as they evolve.

CPOF's tool-and-appliance capabilities are designed to let users create quick, throw-away mini-applications to meet their needs in situ, supporting on-the-fly uses of the software that no developer or designer could have anticipated.

CPOF users at any level can assemble workspaces out of smaller tool-and-appliance capabilities, allowing members of a collaborating group to organize their workflows according to their needs, without affecting or disrupting the views of other users. This capability empowers edge innovation. In addition, a multitude of collaborative interactions has made CPOF repositories an extremely rich source of empirical data on the nature and specific content of C2 business processes. Databases populated in the field in the course of ongoing operations are being mined for the empirical record they provide of human behavior in the context of C2 collaborations. This is a powerful requirements definition and validation mechanism that can augment anecdotal evidence.

After CPOF was transitioned to a PoR, the program manger initiated efforts to port the capability onto the same open technologies that the latest mainline Army systems use. Finding ways to allow rapid federation of CPOF with currently deployed and well-supported IT has proved to be a major challenge.

UNCLASSIFIED

1. CPOF Operational Details

CPOF can receive real-time or near-real-time data from a variety of sources such as GCCS-A, C2PC, and ABCS, and it can display it using MIL-STD 2525B symbols on maps and charts. Plans, schedules, notes, briefings, and other battle-related information can be composed and shared among warfighters. A Voice over IP solution is included, although it can also integrate with a pre-existing voice conference solution.

2. CPOF Deployment

The 1st Cavalry initially used CPOF in 2004 at a handful of locations in Baghdad. The 3rd Infantry Division was the first unit to receive CPOF with enhancements from in-theater experience, and it deployed with another 140 machines the following year. Since 2006, PM Battle Command at Fort Monmouth has directed and managed deployment, sustainment, and feature improvements for CPOF. It is currently in use throughout Iraq and Afghanistan, becoming the primary battalion and above battle command platform in the SW Asia theater of operations, with approximately 1,000 systems in use by Army, Marine Corps headquarters, and Air Force liaison elements.

CPOF was first deployed operationally in a handful of locations in Baghdad by the 1st Cavalry Division in 2004. Two years later, it became a program managed by PM Battle Command at Fort Monmouth, which directs deployment, sustainment, and continuing improvement of the system. It is currently fielded throughout Iraq and Afghanistan with more than 1,000 instances operational.

The system requires a special workstation with three screens that provide a user-friendly, shared environment capable of displaying and manipulating current operational information about friends, foes, and features. Information, including images and data, is seen in two and three dimensions across the distributed workspace accessible by scores of participants. CPOF has a built-in collaboration infrastructure and interactive technology that permits users in different physical locations to operate various tools simultaneously, tools such as collaborative sketching and text or image information sharing. New capability applications can become icons on CPOF displays.

3. CPOF Development

CPOF development actually started pre-9/11 as a C2 technology demonstration sponsored by DARPA and was designed with the assistance of two retired Marine Corps generals. DARPA later expanded the system by adding advanced visualization tools (a multi-screen video wall, video and audio conferencing, online collaboration tools, etc.), which allowed brigade commanders to communicate, share information, and collaborate.

CPOF was equipped with a potent combination of shared data and voice communications that allows its users to rapidly process, prioritize, and respond effectively to new information. Developers paired chat with Voice over IP, which enables the user not only to make and receive telephone calls using a broadband Internet connection, but also to text in support of setup or to back up oral discussions. Fault tolerance for low bandwidth, high latency, and/or error-prone TCP/IP networks is supported by CPOF's multi-tiered client-server architecture, which is specially designed for this purpose. It can be deployed on systems from a two-hop geosynchronous satellite link to a radio network such as JNN while remaining collaborative. The software is largely Java-based but is currently deployed on a Microsoft Windows platform. This is another example of C2-specific infrastructure.

"The ability to give planners immediate situational awareness of activities occurring in the battlespace, regardless of geographic location, is a very powerful tool," said Lt. Col. Richard Hornstein, Product Manager for Tactical Battle Command. "When a Significant Activity (SigAct), such as an IED occurs in theater, a patrol can send the information through an FM radio to a division operations center where it can be posted onto CPOF's shared operational picture. Instantly, that information is available to each individual in the battle space viewing the same digital map display. With near real-time awareness of SigActs, units in the vicinity can either move in to provide support or they can steer away to avoid danger."

A main force driving CPOF development has been the desire to add new data object representations and formalisms into the collaboration space. This has drawn together 14 different software threads from the ABCS. In addition to being federated for

UNCLASSIFIED

information sharing with ABCS, data from TIGR and CIDNE can be loaded into CPOF via thumb drive or external hard drive. Now commanders can do “quick back-of-the-envelope” analyses, to collaboratively work out and decide on courses of action and to share that information with subordinates—not just as end products but as they are developing along with conclusions as they are reached. This results in better information being shared in a greatly accelerated OODA cycle.

The CPOF case suggests that widely adopted collaboration capabilities, which are technically open to federation with new data sources, may prove to be one of the most powerful implementation motivators in furthering the Department’s emerging SOE.

E. CASE STUDY—AIR OPERATIONS TASKING AND CONTROL SERVICES

Key Takeaways:

- Large, complex C2 mission processes generally require composite services. These complex C2 processes must be supported by many “atomic” services, some created by exposing legacy systems, which are orchestrated to assemble and share required information products.
- PoR fielded systems can improve their value through service-oriented transformations, which facilitate repurposing, recombination, and federation to reduce gaps in IT support for business processes.
- Collaborative capabilities crucial to C2 mission processes, and a more efficient planning cycle, are often provided for major system users by applying non-PoR resources in the local command center setting.

This case study features C2 common services delivered through a large Major Defense Acquisition Program (MDAP) PoR by leveraging Web technology to integrate new data sources, improve collaboration, and expand its user base. IT supporting air operations management is fielded within globally distributed AOCs, principally TBMCS in the Combined AOCs (CAOCs), which has a decades-long developmental and operational history. TBMCS is the backbone of joint/combined force automated and integrated capability to plan and execute air battle plans. TBMCS applications, a number of which were derived from previously nonintegrated capabilities, include:

UNCLASSIFIED

- Air Campaign Planning (ACP)
- Airspace Deconfliction (AD)
- Theater Air Planning (TAP)
- Joint Defense Planning (JDP)
- Weather
- Air Tasking Order/Air Control Order Tools (AAT)
- Execution Management: Replanning (EM-R) & Control (EM-C), Close Air Support Tool (CAST), Scramble and Time Critical Targeting (TCT)
- Intelligence: Tactical Command, Control, Communications, Computers, and Intelligence (TC4I) and Target/Weaponneering Module.

Each one of these “applications” is clearly a candidate for development as a service within an overall composite air operations management capability. Although far from fielding at this juncture, Web service development efforts are underway to federate TBMCS capabilities and other air operations planning tools using workflow engines. This Web service approach is expected to significantly speed up the business process as strike plans are worked and shared among headquarters, wings, and squadrons. Transparency will be improved through monitors that relay process status information on wing and squadron Web pages. This enhanced visibility into the process of planning, scheduling, and executing missions is expected to streamline operations, realizing significant reductions in the time it takes to plan, evaluate, and execute decisions.

This case study addresses a prominent TBMCS process and its product from the service-oriented perspective; viz., users who need sorties to accomplish specific air support missions input their timing, routing, and other requirements for coordination/deconfliction and organization into a schedule, which is then promulgated as individual assignments in an ATO. The ATO, as a key information product in every theater of operations, is taken as the focal point for this case study analysis.

At present, the process to create an ATO is extremely complex and ponderous with up to five orders concurrently under development. It is also widely acknowledged that any given ATO from initial conceptualization to execution takes more than 2 days, with an “official” ATO release occurring every 24 hours. Less well understood but known to be vital is the lateral interaction among processes within one ATO cycle and the vertical interaction among concurrent ATOs under development. Identifying these

UNCLASSIFIED

interactions for possible service evolution could reveal specific business process improvement opportunities in this complex command and control environment.

The TBMCS Program Management Office is the Air Force Electronic Systems Center at Hanscom Air Force Base. An operational instance of TBMCS implemented in a theater AOC executes a complex business process requiring many potential services, including C2 common, other domain, and enterprise. The air tasking processes are analogous to machines on the factory floor except that the working medium consists of information flows not only within a developing ATO but also among ATOs. Bottlenecking of information and information inventory can be shown to back up through document completion delays. An Information Age attack on problems like this would be to provide a visibility service that enhances self-synchronization as a continuing activity. Time/space clusters of assignments, labeled as provisional, might then be assembled and pulled by users as they emerge.

Services can be used to provide visibility into bottlenecked information processes. That information can then be used to streamline those processes.

Within a CAOC, the mission of the Combat Plans Division is to provide detailed event descriptions, sequencing, and organizational assignments for upcoming air operations. This division follows up the JFACC's vision and the Joint Force Commander's Campaign Plan by building the air campaign plan and expressing the specific air tasking it contains via the ATO product. ATOs are the orders issued to all JFACC controlled aircrews that assign them offensive missions, defensive missions, and support air missions. TBMCS is the PoR that provides much of the IT to support this process.

In addition to CAOC operations, TBMCS-equipped Air Support Operations Centers provide a tactical extension of the capability for air elements co-located with, and in support of, Army units to provide for the coordination of Army target requests. Allied coalition forces also have access to TBMCS capabilities, with access depending on the particular coalition formed and air war situation.

New technology insertion blossomed under a "horizontal integration" policy and yielded significant results. The most notable result so far is the advent of a far more open architecture to allow these air operations-related applications to share data without having

to significantly reconfigure systems. Recently added capabilities enhance the AOCs' ability to plan with accuracy GPS navigation in support of precision-guided munitions, giving air battle planners the benefit of vastly improved GPS accuracy across the battlefield. The program continues to pursue rapid implementation of capabilities in response to emergent user priorities.

AOC planners have access to space flyover data through the TBMCS-based AOC Web portal, while space vehicle operators will view the ATO and target nomination data through a Space Battle Management Core Systems (SBMCS) capability. The air-space integration effort is synchronized with the TBMCS and SBMCS spiral development schedules and will roll out new capabilities with each new spiral. The next system iteration will support coalition partner access through a Web browser interface.

F. CASE STUDY—FORCE DEPLOYMENT PLANNING AND EXECUTION MANAGEMENT SERVICES

Key Takeaways:

- Some large capabilities delivered by PoRs are simply one contributor among many in a larger, more complex business process.
- When transformed into a service, systems can be more easily federated with other capabilities in support of their “native” business process. Some may be amenable to repurposing in support of new composite services.
- Widespread use of PoR-sponsored systems to support real-world C2 in Web-technology-equipped centers can lead to service-oriented enhancements, particularly federation strategies to enhance information product sharing.
- Capabilities delivered via traditional acquisition approaches can require more than a decade to be fielded, during which time their technological advantage erodes and they accrue very high costs per user.

JOPES is a global C2 common capability, JTF and above, that enables collaborative deployment plan development, execution monitoring, and information sharing within and among theaters in a COCOM's AoR and with national command.

UNCLASSIFIED

JOPES procedures and system capabilities are the mechanisms for submitting movement requirements to TRANSCOM in support of joint operations.

JOPES is the Department's principal tool for designing, monitoring the progress of, and managing force deployments, generally large-scale deployments. It is part of the GCCS federation of capabilities. This system is networked (via SIPRNet) to execute a highly complex business process, and it requires accurate data entry from multiple, widely distributed human and machine sources. Under CENTCOM staff (G-4) auspices, it is actively engaged in support of logistic evolutions for both Iraq and Afghanistan. JOPES functionality is intended to support planning, routing, scheduling, controlling, coordination, and in-transit visibility of personnel, equipment, and consumables. Primary organizations utilizing these services would typically be a Theater Movement Control Agency (TMCC), Corps Movement Control Centers (CMCCs) and smaller unit organizations headed up by Movement Control Officers (MCOs).

Like TBMCS, JOPES is an "industrial strength" capability built on a Relational Data Base Management System (RDBMS) with multiple instances distributed around the world; however, JOPES instances are synchronized to form a centralized virtual knowledge base. Commanders assemble, maintain, and share a situational awareness picture of ongoing materiel distribution by combining JOPES deployment data with resource consumption and flow models plus information on actual or near-term scheduled unit movements and re-supply actions. This picture is key to supporting C2 decision making.

JOPES procedures and supporting IT are the mechanisms for submitting movement requirements to TRANSCOM and PACOM for joint operations in support of OIF and OEF.

In PACOM, where JOPES has come to be used in conjunction with Global Force Management (GFM) prototypes and the Adaptive Planning and Execution (APEX) module, a standard federation architecture for these capabilities is under development in close collaboration with staff users. Lessons learned from operations in SW Asia are also being incorporated. This process appears to be providing far more efficient and user-friendly logistics solutions, methods, and tools than high-level requirements definition.

UNCLASSIFIED

Although JOPES procedures and supporting IT are the mechanisms for submitting movement requirements to TRANSCOM for joint operations in support of OIF and OEF, most supply data are not shown in JOPES. The automated system used by the Distribution Process Owner (TRANSCOM) to track supplies moving through the Defense Transportation System is the Global Transportation Network (GTN). To view supply data in GTN, users must know the transportation control numbers (TCNs) of the cargo or the DoDAAC (activity address code) of the receiving unit. The users cannot simply select a unit identification code (UIC), unit line number (ULN), force module, or Time-Phased Force Deployment Data (TPFDD) and query on all re-supply cargo en route.

A finite number (fewer than 100) of globally distributed staff logistics specialists who understand the processes involved with JOPES are essential to filling, managing, and retrieving useful information from its database. These specialists are also well represented in efforts to enhance the efficiency of the system itself.

The system deals in five major types of movement information: the deploying units, the dates associated with the movement, the locations involved with the movement, the number of personnel and the type and quantity of cargo to be moved, and the type of transportation that will be required to move the forces. JOPES is used to address force movements that range in size from an 18,000-soldier Army division down to a brigade, a battalion, a company, a platoon, or even an individual service member.

JOPES organizes the information obtained from four globally distributed databases, along with scenario-specific information, into TPFDD for movement plans known by a Plan Identification Number (PID). A PID directly corresponds to an operational plan (OPLAN) or concept plan (CONPLAN) and contains all of the unit line numbers and force modules associated with that plan's movement of forces. JOPES Functional Managers grant permissions, restrict access to operation plans on the database, and perform periodic reviews of user IDs and the content of the JOPES database to ensure outdated plans and accounts are removed when no longer required.

A sample JOPES scenario relevant to contemporary operations would proceed as follows: CENTCOM's J-3 staff is tasked to assist in planning for an upcoming rotation of joint forces operating in Iraq. Hundreds of units will be involved in the deployment and redeployment, and the logistician's boss wants him to ensure that this rotation will have an increased fuel storage capacity of 60,000 gallons in case the local fuel pipelines

UNCLASSIFIED

continue to suffer periodic interdiction. The JOPES contribution is to support development of a plan that addresses these needs. Obviously, vital information required to do this task originates in capabilities outside JOPES, and the system's value-added service includes marshalling and relating a great deal of disparate data.

Information-sharing capability that lies outside JOPES program boundaries to support collaboration is exemplified by the U.S. Army Central Command (USARCENT) G-4 JOPES Newsgroup validation message and its daily transmission to USCENTCOM. Information for this Newsgroup product is gathered from Forces Command (FORSCOM), USARCENT staff officers, and Coalition Forces Land Component Command (CFLCC) Forward concerning validation of units that are ready for deployment, units with issues in meeting their movement schedules, and/or units requiring adjustment of their TPFDD data.

To be effectively federated with C2, especially operations order development and dissemination, JOPES and other capabilities would benefit from a service-oriented evolution. Even in its present state, JOPES handles a number of arcane identifiers to relate various older systems and processes. The Net-Enabled Command Capability (NECC) program analyzed how various planning capabilities might be decomposed and rearranged as a composite service, including the functionality that JOPES provides.

1. JOPES and Adaptive Change Process

PoR-sponsored efforts are underway to evolve this joint C2 capability in accordance with DoD's service-oriented vision. In FY09, GCCS-J was scheduled to complete development, testing, and fielding of spiral releases that address currently unanswered operational requirements and net-centric implementation requirements. This involved core infrastructure upgrades to the GCCS operating system, database, and security capabilities, completing the implementation of unified account management via Public Key Infrastructure (PKI) and single sign on. New functionality was to include (inter alia) Web-based access to force planning and readiness data, a capability to aggregate readiness data, implementation of dynamic and deployment force structure modules, and Web enablement of the JOPES Rapid Query Tool (RQT). Using the RQT, commanders can monitor arrival of forces in their AoRs using GTN data combined with the Scheduling and Movement (S&M) application in JOPES. Utilizing data from movement control agencies and intratheater lift providers, this application allows

UNCLASSIFIED

supported commanders to receive, stage, onward move, and integrate forces arriving from points of departure in CONUS. This enhancement amounts to fulfillment of the “track shipment” requirement combined with an ability to manage further force dispositions. The RQT requirement was first identified and defined at the turn of the century; so a decade was required for it to near IOC. Architectural enhancements were to include the migration of Adaptive Course of Action (ACoA) from a local to an enterprise-level capability and eliminating the need for local replication of readiness data.

2. JOPES Web Interface Example

The Consolidated Air Mobility Planning System (CAMPS) is an Air Mobility Command (AMC) sponsored suite of airlift and air refueling planning, scheduling, and analysis tools. CAMPS has demonstrated the capability to move into a global environment where the system is used not only at HQ AMC’s Tanker Airlift Control Center (TACC) at Scott AFB, but also is used in the CENTCOM theater for planning and scheduling wartime missions. Oak Ridge (ORNL) has continued to improve capabilities to do automated scheduling of missions based on a host of criteria that the planner can model with CAMPS. ORNL has also developed a set of Web-service-based interfaces with JOPES whereby a CAMPS user can request an OPLAN from JOPES, and CAMPS will automatically retrieve the appropriate data using Web services, and then load the data into CAMPS for subsequent scheduling.

UNCLASSIFIED

(This page is intentionally blank.)

UNCLASSIFIED

UNCLASSIFIED

Appendix B
C2 INFORMATION SHARING FRAMEWORK (C2ISF)

UNCLASSIFIED

Appendix B

C2 INFORMATION SHARING FRAMEWORK (C2ISF)

This appendix expands on the technical aspects of two critical artifacts of the C2ISF design-time infrastructure: (1) the Service Description Templates, and (2) the Run-Time Infrastructure Rules and Protocols. Refer to Section III for an introductory-level discussion of both these subjects in the context of the overall C2ISF.

A. SERVICE DESCRIPTION TEMPLATE USE

1. Motivation

As mentioned in Section III.B, creating templates for the descriptions of C2 services improves the ability of C2 Service Users to find the information and services for which they are looking by providing a regular structure and vocabulary, which enables effective registration, search, and discovery. It greatly reduces the possibility of a user missing a service to fill his needs because he uses the wrong search word or because a description poorly denotes a service's attributes. Generally, using templates for describing services results in improved:

- Information, document, and service discovery
- Information and data mediation (including transformer creation)
- Information sharing.

Knowledge Organization Systems (KOS) [which can comprise any number of Controlled Vocabularies (CVs) and taxonomies] are designed to provide a standard way of describing a group of information. Structured data conforming to them (i.e. described schematically by them) may be readily composed, parsed, or semantically interpreted. Rich, formal models with standardized syntax for describing URIs and links exist for KOSs, making the KOS useful for describing information and services available over the Web. Thus, service templates conforming to a well-defined KOS can assure the discovery of the services on the Web to which they refer (the services' presence on the Web is embodied by their having a URI or, especially, a URL).

UNCLASSIFIED

Two information-scientific measures of search performance valid for any type of IT-assisted resource searching exist: precision and recall. Precision is defined as the number of returned documents that are relevant to the user's search criteria divided by the total number of documents returned. Accordingly, a high value of precision implies that more of the returned search hits are useful to the searcher. Recall is defined as the number of relevant documents returned by the search divided by the total number of relevant documents that exist. A high value of recall implies that the search results include most of the relevant resources that are available.

Implementing or extending the component parts of a KOS—CVs and taxonomies—can aid in improving the values of these search performance measures for a given search system. CVs help reduce ambiguity in what the best search terms may be for a concept, which promotes more accurate registration of C2 services by the C2 Service Owners. Recall is improved when a suitably comprehensive set of CVs is employed by the searcher. If the CVs and the taxonomy are complete, orthogonal, and sufficiently granular (differentially descriptive) for the domain in which they are employed, precision is also improved, as the resources themselves (e.g., services) may be arbitrarily finely distinguished such that only hits on relevant classes of results be returned. Given that it is always in the best interests of the enterprise to reuse services that are expensive to build rather than duplicate them, having high levels of performance according to these measures is clearly desirable. It therefore behooves the DoD to consider establishing a KOS around the discovery activity. Later sections describe important aspects of implementing this KOS.

A rich KOS model includes relationships between elements, terms, and categories of terms. Knowing these relationships facilitates information mediation because they can be used to infer the meaning of an unfamiliar element, term, or term category in terms of related elements, terms, or term categories. Broader terms may be immediately entailed by a narrower term, thereby increasing the chance that a user unfamiliar with the narrower term may still be able to interpret resources associated to it in light of one of its related, broader terms. In an analogous fashion, a rich KOS model may also facilitate transformation of information with one schema into information compliant with a different schema.

Structured data conforming to the KOS may be readily shared between different parties, since the formal semantic interpretation of the information contained in those

UNCLASSIFIED

Web resources is completely described by the KOS model. KOSs are usually published (in our case, discoverable on the network) to allow the largest audience to exploit the host of structured data conforming to them.

2. Example of Discovery Infrastructure

Searching for C2 services should be performed with constraints from a KOS, and the KOS must be able to handle descriptions of Web-visible resources. This suggests the need for a Web-based User Interface (UI) that allows the user to navigate through the contents of the CVs and the taxonomies, as well as the associated services descriptions. Moreover, this discovery service must be able to adequately relate information on Web-based resources to the user, and it and the underlying KOS must, as a practical necessity, be able to handle and process URIs and URLs automatically (i.e., not merely as strings, but as indicators to existing and managed Web-resources (managed at least in the providing of the resources' metadata to the C2ISF discovery service)). This should be accomplished in conformance with the Architecture of the Web [<http://www.w3.org/TR/webarch.>]

It is economical for new DoD services to leverage existing DoD enterprise discovery service capabilities, namely the Metadata Registry (MDR) and the Universal Description, Discovery, and Integration (UDDI)-based Discovery Service of NCES. The former capability uses as its primary registry artifact (resource discovery metadata) the DoD Discovery Metadata Specification (DDMS), and description documents conforming to it are stored in the DDMS in the form of XML documents. UDDI is an industry-standard means for discovering Web services, such as those that are part of a SOA. Its scope is too narrow to afford the needed visibility into all types of services needed for C2. Nonetheless, its object model gives the necessary framework for publication of, finding, and binding to Web services, and it is an important technology that enables use of Web services (such as in a SOA).

It is recommended that normalization occur between the MDR and the NCES Discovery Service such that a system interoperating with, but more general than either MDR or the NCES Discovery Service, can be implemented to handle documents in the manner depicted in Figure B-1. New descriptions of services should be entered into the service registries roughly in the manner depicted in Figure B-2.




Figure B-1. Generic Registration and Discovery Infrastructure Use Case Diagram




Figure B-2. C2 Service Metadata Creation Sequence Diagram

UNCLASSIFIED

These documents describing services are conceived to be KOS instance records and are naturally serialized (e.g. RDF/XML) resource descriptions conforming to the KOS's CVs and taxonomies. One candidate KOS framework is SKOS (Simple Knowledge Organization System), which may be normalized to work with other, semantic metadata initiatives such as the ontological modeling of the U-Core found in U-Core-SL.

The interoperability of the KOS system with the legacy MDR and UDDI-based systems will be an easy matter, given that both of these systems are based on open standards. Especially for UDDI, all updating and federation of the registry may be performed through an Organization for the Advancement of Structured Information Standards (OASIS) Application Programming Interface (API). Because the MDR conforms to an Model-View-Controller (MVC) architecture and its resources are XML-based, it is suspected that full process interoperability is also achievable.

Storage of service descriptions should be based around the resource document instance as a basic unit of information. Provision of these descriptions to the users is most naturally enabled by Hypertext Transfer Protocol (HTTP) over a Transmission Control Protocol/Internet Protocol (TCP/IP) stack, and should conform to the Atomicity, Consistency, Isolation, Durability (ACID) management paradigm. Ideally, these description documents would also be exposed (for automated access) by one or more Web service endpoints [which conform, e.g., to a Simple Object Access Protocol (SOAP)-based Web service access scheme or the RESTful paradigm].

B. SERVICE DESCRIPTION TEMPLATE MANAGEMENT

To keep current the templates for describing C2 services, so they can effectively differentiate between existing and new services, the organizing system for the types of data in the templates must be updated to include new terms and new relationships between terms. To this end, there are many lessons to be learned from the Library of Congress' (LoC) management of its fully operational KOS (in operation for over a century), which includes 5 million CV terms (e.g., names) and greater than 0.3 million taxonomically structured topics in the LoC Subject Heading (LCSH) taxonomy.

1. KOS Management Example (Library of Congress)

The KOS's editorial board meets weekly to align newly suggested terms to the existing names and "classification manual" of the LCSH taxonomy. The board has experts who normalize the suggested terms both with respect to that taxonomy and to "authority work," which is a collection of efforts to understand the current usage of English language and specialized jargon. Authority work includes, for example, tracking the usage of English language terms to identify synonyms and how they are used. This information can be then used to find the most appropriate candidate for a taxonomical term among synonyms when finding a substitute for existing terms may be too narrow or specialized or may have drifted away from the vernacular. After the term suggestions are normalized, they then enter the comment phase, where subject-matter experts, stakeholders, and identified interested parties comment on the suggestions. Finally, accepted terms are added to the KOS, and necessary substitutions and prudent deletions are made concurrently to maintain the orthogonality and currency of the KOS term collection. This latter work is, perhaps, the most ad hoc and specialized of the KOS maintenance activity, requiring experienced human capital. The entire process is performed continuously, and it takes 3 weeks for a term to percolate through and enter the KOS. An editorial control group governs this process, ensuring conformance to industry-standard "best practices" (c.f. ISO 2788-1986 (E) and ANSI/NISO Z39.19-2005) and performs change management of the artifacts [e.g., linking versions of the artifacts, final decision authority on contentious issues, and oversight of dissemination (to customer libraries)].

2. Additional C2ISF Template Management Considerations

Tasks analogous to those in the LoC's KOS maintenance activity should be performed for the C2ISF template maintenance, and it would greatly improve the process to additionally make it and its products net-ready. This would involve both standard employment of Web technologies (including Web services) and use of a document resource base (set of artifacts) that is Web-oriented itself. In particular, the interchange and archival serialization forms of the KOS instance records (e.g., C2 service descriptions) should themselves conform to the Architecture of the Web, not only in their description of the resources but also in the way they are published. This means they must be describable and discoverable according to standard Web-based and library-science-information-technology means.

SKOS and other semantic technologies [e.g., those using Web Ontology Language (OWL)] have an abstract syntax given by the Resource Description Framework (RDF). This syntax has subject, predicate, object statements (“triples”) with a limited amount of (type) metadata associated to each triple.¹ This abstract syntax is conveniently serialized using an XML encoding for RDF (RDF/XML), so the semantic and syntactic information set of a KOS item (artifact or instance record) is stored as an XML document.

Change-request handling, one of the major activities in KOS management, must be performed by a set of skilled librarians under the supervision of a governance board. The governance board must have sufficient authority to propagate changes to the KOS within the discovery service domain as well as to related, interested parties (e.g., customers with associated and automated information processing technologies).

KOS change management may take place similarly to the LOC’s KOS management overview above. It should be noted that the primary knowledge base (set of artifacts) that describe the KOS conceptually includes information on all the competencies of the personnel involved in the service life cycle. Creation of the templates must be performed with the intention of modeling and inter-relating fields relevant to personnel both using and building the KOS, as indicated in the next section.

C. TEMPLATE FEATURES

1. General Requirements

Conceptually, the templates model the view of Web resources of a specific category of users—other services. Each template is, formally, a taxonomical structure imposed onto the domain of services. It should ideally satisfy the principles of completeness (within the domain) and orthogonality (to reduce confusion during registration and search). Nonetheless, it is likely that a much more relational structure is desired within the abstract syntax for describing taxonomies themselves, which could be used to improve the ability to search the template’s taxonomy, for example. One way to accomplish this is to form a thesaurus from the terms of the taxonomy. Searching within a particular view for a specific user group is likely best aided by a thesaurus holding the

¹ Full-resource metadata-referencing quads are out of the scope of the specification, but have been added to numerous RDF-processing APIs.

UNCLASSIFIED

taxonomical terms (and having essentially the same structure as the taxonomy view). Relating terms across views and versions of the template set or across taxonomies, such as for normalization of different communities' templates to one another, requires the increased (semantic) expressiveness inherent to the structure of an ontology.

In Figure B-2, we showed the templates' highest-level contents. While more modeling must be done to make these concept classes unambiguous, not much conceptual modeling is needed initially since all these concept classes have been described at length in other DoD publications. The initial modeling work needed is essentially a translation of the existing models from these domain-specific modeling languages (e.g., JCIDS authoritative publications) to the abstract syntax appropriate to a thesaurus/taxonomy, which allows for true semantic interchange and discovery.

Supposing the templates' description of resources is Web-based (i.e., uses URIs), the templates may naturally store information about external documents or artifacts (for example, through XLink, RDF itself, or even the venerable `@ref/src` of (X)HTML). Accordingly, not all of the information needed to describe the resource must be stored directly in the resource description (a "filled-out template," or template instance document); in keeping with the Architecture of the Web, the resource description may include references to remote documents that describe the resource (and that also allow the user to call up those remote documents' data). As a basic example, an XML Schema Definition (XSD) may be referenced in a resource description instead of including a copy in the description of every single service that uses it. The schema could contain very detailed, and possibly classified, information about a fundamental information exchange that the service can participate in [according to some Web Services Description Language (WSDL), for example], but the resource description may merely list the unclassified metadata exchange endpoint URL of the schema. That endpoint may be protected (through proxying or NAT, for instance), and can be navigated to according to proper (WS-)Policy to learn the details of the schema. The template instance, regardless of the contents of the schema, is most naturally an unclassified document.

2. Semantics

Although more formal models of semantics may be developed, one that is sufficiently simple to implement with Web-based IT is that of the OWL. To illustrate what semantic models are, we will describe the much less restrictive, more basic RDF.

(The RDF is less powerful from the perspective of description logic frameworks and rule engines, both of which are outside of the scope of this study.) Informally, the RDF essentially models semantic statements (meaningful statements) as a triple of resource parts: a subject, a predicate (i.e., the relationship) and the object. Each resource part is identified by its URI, which may or may not correspond to a Web-based resource (this is part of the power of the formalism). The RDF triple is an example of an “abstract syntax,” which is a way of stating the relationship of the resource parts to one another. These abstract syntactic structures of semantics are easily adapted to use with standard library science knowledge organization constructs, which are readily applicable to the C2ISF template specifications. As a specific example, the SKOS is a candidate KOS framework that is an RDF model and also OWL compatible. Not surprisingly, these semantic models naturally play a role in formalizing Web-oriented taxonomical structures and thesauri (e.g., for the LCSH).

Possible semantic requirements that may be made for the C2ISF semantic model are for the creation of a synonymy (preferred terms to allow for redundancy of terms within the thesauri) or a hierarchy (which allows different levels of specificity in a user’s specification of a term), but they are not necessary for a functioning C2ISF. Standard KOS considerations should be adhered to when applicable, especially those for monolingual thesauri indicated in ANSI/NISO Z39.19-2005 and ISO 2788-1986. Also, the versioning regimen in the KOS must be robust, a feature readily provided by OWL, which has sufficient metadata classes and relationships for this activity. It is further necessary that the semantic model used for the C2ISF be able to represent the evolution of a service through the governance regimen described in the CONOPS (Appendix A).

3. Candidate Ontologies

U-Core SL (Universal Core Semantic Layer) is a transcription and alignment of the U-Core XML serialization (XSD data model) to the OWL’s abstract syntax. (The semantics in the U-Core SL are not automatically apparent in U-Core’s XSD). U-Core SL makes the semantics of the U-Core explicit, thus providing a general-use ontology whose concepts users of U-Core will find familiar. U-Core SL additionally allows for extensions and data model management through the OWL syntax (the need for which was mentioned previously), making it straightforward to include U-Core SL in an OWL-based KOSs. Moreover, other relationship types may be similarly modeled within the OWL framework as extensions off of U-Core SL.

UNCLASSIFIED

The LCSH taxonomy (a thesaurus-structured subject heading taxonomy) is another general-use taxonomy that has an extremely wide user base. Most academic libraries employ this taxonomy for their collections, and, due to its consequent and continuous management and updating, it is a KOS component of extremely high pedigree. Although the subject matter of the LCSH is considerably broader than needed for the services supporting C2, the term structure, generality, and wide user acceptance of this taxonomy all argue for its inclusion in the C2ISF artifact set (at least as a reference taxonomy). The LoC office managing the KOS to which it belongs has deliberately designed the taxonomy to be general-use, intending that more-specific, more-tailored taxonomies or associated vocabularies will be derived from it. The obvious advantages of this are that the developer of the derived taxonomy or vocabulary will immediately have a very broad-based, semantically rich, and expressive alignment to a functional KOS.

The Basic Formal Ontology (BFO) is a formal upper ontology, meaning it is a general concept model used for (model-theoretic) alignment of other ontologies in a KOS. As it is a formal and abstract ontology, its direct exposure to an uninitiated end-user may have detrimental effects on that user's discovery effort. An upper ontology such as BFO is, however, useful in the C2-services KOS, as a complex set of end-user-focused taxonomies and thesauri is likely to develop. The upper ontology will help assure semantic interoperability of these artifacts, which aids user comprehension, improves the effectiveness of the KOS-based search or resource description, and hastens the development of the tools used to employ the KOS. Using an upper ontology like BFO will decrease ambiguity in the overall data model and the potential for formal or unintentional semantic error.

All three of these candidates can be seen as complementary parts of an upper-level or general-use ontology. BFO is, most likely, the ontology with the most immediate applications of organizing the KOS itself and assuring alignment of its constituent parts (from the point of view of generalization). U-Core SL is a cross-cutting, generalized, but still generally useable higher-level data model, and it also has end-user (operational) applicability as its particular terms should be immediately comprehensible to most DoD/IC/Government users.

4. Serialization

As stated earlier, the choice of the precise serialization mechanism is of a less critical nature than the construction of the template model set and semantics it expresses. For instance, through a predefined SAWSDL (Semantic Annotations for WSDL) transformation, a complete RDF semantic model can be recovered from a semantically imprecise XML schema (such as DDMS) that was constructed using an underlying semantic model. (DDMS, incidentally, was partly derived from the Dublin Core, which already has a semantic underpinning. Producing the SAWSDL transformation is largely an exercise in stylesheet composition and not one of data modeling.) Other XML serializations of RDF-based abstract schemas such as SKOS are readily achieved using standard, available APIs, and are generally as straightforward as object serialization in modern object-oriented programming frameworks. Thus, there exists a “Web service middle layer”—the serialization step, which is simply and formally easy to accomplish. This activity enables data interchange, as well as publication/Web presence of the instance data—in the form of an XML document.

D. C2 RUN-TIME INFRASTRUCTURE RULES AND PROTOCOLS

1. Motivation

As mentioned in the C2ISF description in Section III.B, the C2ISF Run-Time Infrastructure Rules and Protocols are the engineering-level description of how the C2ISF operates. They are the critical component of the design of the C2ISF that explains how the various software, information, and standards [e.g., World Wide Web Consortium (W3C) standards, C2ISF artifacts] will be coordinated to enable C2 information sharing. The C2ISF Run-Time Infrastructure Rules and Protocols must be designed to provide the information needs of the C2 services in the C2 portfolio while accounting for the constrained network capabilities and characteristics. In addition to a general discussion of example C2ISF rules and protocols, this section includes a description of current technologies, methods, and standards that can motivate the design of the C2ISF in influencing the rules and protocols. This description is not intended to be a specific, comprehensive, or authoritative set of rules and protocols for the eventual C2ISF.

2. Notional Registration Needs and Methods

For services in different C2 service tiers (see CONOPS, Section IV.C), there will naturally be different data required when filling out the Service Description Templates. For users to effectively discover the services they need, there is a minimum amount of metadata required for each C2 service:

- URI (and URL where possible)
- Service Owner
- POC
- High-Level Description (of its functionality)
- Virtual Coverage (what network domains it can be called from)
- Access Procedure
- Operational Status
- Classification Level.

The URI is the name for the service, and a URL should ideally be included. (The URL is the location of the service in a network—the service’s address.) This is clearly needed to allow C2 Service Users to know that they are referring to the same service and to allow them to call it reliably. The part of the Service Owner that must be identified is the entity that sets requirements and the entities that implement changes. These parts of the Service Owner must provide their identities and digital signatures, in addition to the listing of the overall Point of Contact (PoC) that C2 Service Users should call for assistance. A high-level description of what the C2 service provides is important for the C2 Service Users to identify whether the C2 service is even close to what they need, in the absence of a more detailed service description. The virtual coverage that the C2 service has is the set of network locations where it may be called (per DDMS). This information should ideally identify which (Web) domains or subnets have access, if possible. The access procedure for the service should include, at minimum, a description of what sorts of users automatically have access and what sorts may be granted access if they apply (the default for this field should be no networkable access, and contain default information in the vein of “call POC for individual access”). The C2 service’s operational status might be set manually but must indicate whether the service is working, in general. Possible values could be “in development,” “available,” “under repair,” “deprecated,” or “retired.” The classification levels of the service should

UNCLASSIFIED

conform to the basic security metadata standards in IC-ISM and should describe not only the minimum level of clearance the C2 Service User must have to use the service, but also the levels at which the algorithms and output are classified.

Most of this information is included because it is so fundamental to the operations of the C2 Service Users that many of them could conceivably rule out using the service almost immediately with just this information. For example, if a service is not currently operating, there is no use in calling it from a code that is already in use.

One fundamental C2ISF rule is that all C2 Service Owners must provide descriptions of their services to one of the C2ISF service registries (i.e., register their services). This requirement is necessary to allow unanticipated users to find C2 services they could benefit from, and for the general visibility of all C2 services. Notional examples of the minimum metadata requirements for C2 services according to their category in the governance taxonomy are presented in Table B-1.

Table B-1. Notional Data Required by C2ISF Service Description Templates versus Service Type (in governance taxonomy)

Category	Data	C2 Local	C2 Common	C2 Infr.	Enterprise Infr.
General Data	Service Name (URI)	X	X	X	X
	Contact Info (POCs) (Service Owner)	X	X	X	X
	Version		X	X	X
	High-Level Description (including motivation)	X	X	X	X
	Location (URL)	X	X	X	X
	Operational Status	X	X	X	X
	Classification	X	X	X	X
	Other Security Data		X	X	X
	Other Joint IC/DoD Enterprise Services Registry Taxonomy (Appendix D) required data: Namespace, Creator, Publisher, Creation Date, Effective Date, Validation Date, End of Life Date, Geographic Coverage, URI of Related Data Resources, Rights to Data (copyright, etc.), Classification Data (incl. dissemination and access controls)		X	X	X

UNCLASSIFIED

Category	Data	C2 Local	C2 Common	C2 Infr.	Enterprise Infr.
Governance Data	Current JCIDS Milestone Achieved		X	X	
	Link to Source of JCIDS Documents				
	JCIDS Schedule and POCs			X	
	Funding Sources		X	X	
	Requirements		X	X	
	Which Authorities Ensure It Is Useful (Process Owner)		X		
	What Commands Are Using It		X		
	Access Control Policies	X	X	X	
Technical Data	Technical Service Specifications (Input Format, Output Format, Call Procedure, Standards and Technologies Used) (for a Web service, this is a WSDL)		X	X	X
	Dependencies on Other Services (e.g., list of services and data required for full operation)		X	X	X
	Description of Algorithms (a detailed explanation of how the service works)				X
Operational/ Functional	Performance and Usage Metric Measurements (usage statistics)		X	X	X
	Maintenance and Provisioning Data: When and Where It Can Be Expected To Work, How Much Traffic It Can Handle (calls/hour)		X	X	X
	Dependencies: Needed Hardware, Software, Bandwidth, Data and Data Sources, Communication Infrastructure, Schemas		X	X	X
	Which C2-Objects, Core Taxonomy, Joint Common System Function List/JCA-Derived Objects It Uses		X	X	
	Categorization According to Joint IC/DoD Enterprise Services Registry Taxonomy (Appendix C)		X		

In the notional example, C2 common services and C2 infrastructure services have more data in their descriptions than C2 local services or Enterprise Infrastructure Services. As suggested by the description of the governance taxonomy, C2 local services descriptions must be lean so as to burden innovators creating new services as little as possible, while still assuring compliance with C2ISF rules. It may also prove difficult (if not impossible) to describe such a rapidly changing set of capabilities. C2 common services will be more widely used than C2 local, and they must have more detailed descriptions to match the likely greater reliance on them. Also, there is more oversight of C2 common services, a greater participation in formal acquisition for them, and more

UNCLASSIFIED

management need for usage monitoring data on them. Since C2 common services are designed primarily for the C2 community, more detail is needed to distinguish the services' functionality from one another than is needed to distinguish a generic Enterprise Infrastructure Service from any community service or other enterprise service. C2 Infrastructure Service Descriptions must be general, since the services must be useable across the C2 community, but they are still subject to extensive changes and complicated access control, so they require more detail than Enterprise Infrastructure Services' Descriptions. The Enterprise Infrastructure Services are designed to be generic enough and fundamental enough that they are useful across the entire enterprise, so their descriptions may focus on how to use them and how to obtain help on their use.

3. Notional Federation Method

One example of a protocol would be a detailed description of how the service registries are federated. This would include the query form that a service search engine should use to automatically search a registry with which it is federated, how long the search engine should wait for search results, how frequently the registry should indicate it is still searching, and whether the service search engine or the registry is responsible for translating search terms into ones compatible with the registry (which of those two services calls the mediation service before the information is delivered to the User).

In general, federation of services refers to the facilitated use of processing and data services that are distributed across a network of largely autonomous nodes with a set of rules but without central ownership. This means that there is an agreement among participants on how to exchange or invoke distributed processing and data services to perform an operational or technical task or create a composite service. Federation participants (e.g., processing nodes on a network) are often called federates and execute under local resource control and management. In addition, the degree of federation of services can vary from loose to tight in terms of both governance and the technicalities of implementation and as a function of both user operational requirements and the maturity of the SOE.

Federation components, rules, and protocols have two basic sets of requirements, regardless of the information flow topology (e.g., synchronized or some form of

UNCLASSIFIED

hierarchical² updating). First, a small set of common exchange, serialization message models are necessary for exchanging the data on which the federation operates. Second, remote operations (transactional method invocations) that conform to the ACID paradigm must be used for making changes to documents. The second requirement is made so that the fidelity of the states of all documents, records, and operations thereon are maintained at all times at all points in the federation.

As described formally within the Web services stack of protocols (including SOAP, WSDL, and even ULex), common exchange, serialization data models are necessary for exchanging the data on which the federation operates. Essentially, above the “application layer” of HTTP, XML wrappers of XML documents are used to convey expected, atomic message-handling activity to the recipient Web service. In fact, one of the underlying motivations for the development of the WSDL, SOAP, and other even higher-layer protocols (e.g., within the WS-* protocol suite and the OASIS ebXML set of business automation protocols) is to allow the description of all aspects of individual XML document handling and associated remote method invocation (such as authentication, fall-back, and any other operation sequencing) in the document itself. Expected WSDL protocol components that must have standards for federation include requiring the use of a SOAP envelope around well-described (schema-compliant) XML instance documents or the use of a REST set of HTTP operations with an XML payload. In either case, these clearly ride over HTTP.

When using these interchange protocols, remote transactional method invocations that conform to the ACID paradigm must be used for making changes to documents. Although true synchronization across the federated set of (discovery) services is not necessary, or even feasible (e.g., in a DIL environment), the state of the each of the records (i.e., the descriptions of services that conform to the templates) must be assessable locally so they can be compared to any other document within the federation. Also, during the time when different instances of the same document are being compared (e.g., for updating or other replacement) or otherwise operated on, other modifications to the document must be locked out (delayed) so there is no more than one action occurring at any point in time on the same document. This ensures the existence of an authoritative version of the data and allows the institution of high-level rules regarding the document

² More generally, acyclic.

UNCLASSIFIED

handling so that local or authoritative update, possible synchronization, and other federation features can effectively occur.

In addition, the authority of a user to add data to an existing document must also be verified before the changes are made. The owner of the data is usually allowed to modify it, but others are only allowed to read and possibly extend the data (e.g., with “annotations”). The typical basic set of remote invocations (“commands”) used to enable federation normally includes the Create, Read, Update, Delete (CRUD) method set, as seen in Structured Query Language (SQL) database distributed topologies and in the Universal Description, Discovery, and Integration (UDDI) object model that supports federation. It should be noted that a built-in, systematic tolerance to latency to accommodate the federation’s need to operate at Web speeds may be required of the federated set of systems. Standard asynchronous access paradigms should be sufficient for the near-term service visibility use case, as synchronization is often not a must for such Web-based applications [c.f., Domain Name System (DNS)].

In addition to interchange protocol standards’ effects on federation requirements, there is also the issue of the schematic and semantic interpretation of the payload itself. These payloads only need to employ a standard serialization of template instance documents—basically every user and service will use HTTP, the same templates as one another, and all the other same protocols in between, which ensures the semantic interoperability of the participants in the interchange. For example, one may use SOAP invocation of standard CRUD actions over HTTP, authenticated via a WS-Policy-containing set of elements and containing an RDF/XML serialization of a template-conformant instance document.

Another more complete example of federation is the CRUD method set implemented in UDDI. Here, the protocol stack is set up, and remote method invocation and ACID discipline is enforced in such a way that UDDI instance services can formally federate and interoperate according to some paradigm (top-down refresh, for example), all the while maintaining local repository integrity, if not cross-federation synchronicity. Cross-federation synchronicity is difficult to achieve with Web-speed-like latencies but is likely not required for service visibility of the type envisioned in the C2ISF. A similar tolerance to latency works in DNS whois updates, where domain name information-set changes (additions, deletions, metadata modifications) are propagated out from centralized servers roughly hourly across the global Internet, through a hierarchical

UNCLASSIFIED

bubble-up-and-disseminate-down topological flow paradigm. The existing whois data records, used for registering internet resources, have a subset of the type of information that is needed for registering C2ISF services.

UNCLASSIFIED

UNCLASSIFIED

Appendix C
DEFINITIONS AND GLOSSARY

UNCLASSIFIED

UNCLASSIFIED

Appendix C DEFINITIONS AND GLOSSARY

PART 1—DEFINITIONS

ACID: Atomicity, Consistency, Isolation, Durability. “In computer science, ACID (atomicity, consistency, isolation, durability) is a set of properties that guarantee that database transactions are processed reliably.” (<http://en.wikipedia.org/wiki/ACID>)

Agility: The ability to respond effectively and in a timely manner to changing circumstances against a thinking and adaptive enemy, from anywhere in the battlespace, at any time, even when the networks and command structure are degraded. Agility includes both “flexibility” and “responsiveness.” Agility enables organizations, systems, or processes to react and adapt to changing situations and conditions, such as performing C2 during operational transition and reorganization/reconstitution; while airborne, afloat, or “on the move”; or in response to enemy actions. (C2 Joint Integrating Concept, v. 1.0, 1 Sep 2005)

Assured (service): (see “Information Assurance”)

Attribute: A quantitative or qualitative characteristic of an element or its actions. (CJCSI 3010.02B; CJCSI 3170.01E). In this case, a piece of information describing the element that is interpreted according to the data model associated with that element.

CRUD: Create, Read, Update, Delete. “Create, read, update and delete (CRUD) are the four basic functions of persistent storage.” (http://en.wikipedia.org/wiki/Create,_read,_update_and_delete)

Architecture: (1.) Structure. (2.) A framework or structure that portrays relationships among all the elements of the subject force, system, or activity. (JP 3-05)

Architecture of the Web: A standard describing how the parts of a functioning network and included Web-ready applications should interact. (<http://www.w3.org/TR/webarch/>).

Artifact: A part of a data model. For example, ontologies, knowledge organization systems, taxonomies, controlled vocabularies, and schema are all artifacts.

UNCLASSIFIED

Bandwidth: The difference between the limiting frequencies of a continuous frequency band expressed in hertz (cycles per second). The term bandwidth is also loosely used to refer to the rate at which data can be transmitted over a given communications circuit. In the latter usage, bandwidth is usually expressed in either kilobits per second or megabits per second. (JP 1-02)

Bottleneck: A feature of the system that particularly constrains the flow of information or material through the system.

C2 Core: Part of the design of the C2ISF that includes the Joint C2 Conceptual Model and Vocabulary and C2-specific extensions to U-Core.

C2ISF (Command & Control Information Sharing Framework): A collection of services and information used for creating, finding, using, and managing C2 services and information. It includes the C2ISF Run-Time Infrastructure (services), the C2ISF Run-Time Infrastructure Rules and Protocols that describe its operation, the C2ISF Service Description Templates used to describe services, and data model artifacts (such as the C2 Core) that give the language used in the service descriptions.

C2ISF Design-Time Infrastructure: The information containing the design of the C2ISF, including the C2ISF Run-Time Infrastructure Rules and Protocols, the C2ISF Service Description Templates, and data model artifacts.

C2ISF Run-Time Infrastructure: The services used for creating, finding, using, and managing C2 services and information. It must include critical services such as service registries, artifact registries, a URI/URL management service, and cryptographic services and may include other useful services such as monitoring, role management, schema transformers, tagging engines, and feedback services.

C2ISF Run-Time Infrastructure Rules and Protocols: Rules and protocols that are designed to regulate the operation of the C2ISF Run-Time Infrastructure.

C2 Information Support Service (see Service): A means of facilitating a C2-related mission outcome through the provision of information or data processing support without the customer owning the mechanism for providing that support. Generally, information services will provide access to data, computational or transactional functions, or management or orchestration functions.

UNCLASSIFIED

C2 Common Services: Capabilities that fulfill data and/or functionality requirements inherent in multiple C2 missions but that are not expressly tailored to or necessarily useful for supporting other mission areas. C2 common services will typically be available for use by multiple commands in one or more AoRs (i.e., regionally) or globally.

C2 CONOPS: An organizing construct that ties together C2 service categories, key implementation roles and responsibilities throughout the complete service lifecycle, and a high-level governance construct. (see “CONOPS”)

C2 Infrastructure Services: Mission-specific, general-purpose capabilities, configured expressly to address C2-community-specific performance requirements, business processes, and/or behavior characteristics that provide for basic communications, collaboration, publication, discovery, security, and information and service management.

C2 Local Services: Mission-oriented information services that are tailored to meet the needs of a limited group of users, e.g., specific organizations or entities, usually within a single organization and/or AoR.

C2 Service (see Service, C2 Information Support Service): An information service that is designed especially to support a C2 operational process or to facilitate the operation of one that does (rather than being designed to be simply generally useful).

C2 Service Description Templates: A structure that indicates which metadata about an Information Service should be provided in descriptions of an Information Service that will be stored in a particular registry. The provided metadata must use the terms and definitions in the design information (artifacts) for that registry.

C2-Specific Extensions from the U-Core: Schema components and vocabulary added to the U-Core as required, providing an ability to share more detailed data within the C2 community. C2-Specific Extensions from U-Core are under configuration management of the C2 CPM in cooperation with the C2 community.

Capability: The ability to achieve a desired effect under specified standards and conditions through a combination of means and ways across doctrine, organization, training, materiel, leadership and education, personnel, and facilities (DOTMLPF) to perform a set of tasks to execute a specified course of action. (DoDD 7045.20) (CJCSI 3170.1E); the ability to achieve an effect to a standard under specified conditions through multiple combinations of means and ways to perform a set of tasks. (CJCSI 3010.02B)

UNCLASSIFIED

Capability Portfolio: A collection of grouped capabilities as defined by JCAs and the associated DOTMLPF programs, initiatives, and activities. (DoDD 7045.20)

Capability Portfolio Management: The process of integrating, synchronizing, and coordinating Department of Defense capabilities needs with current and planned DOTMLPF investments within a capability portfolio to better inform decision making and optimize defense resources. (DoDD 7045.20)

Capability Portfolio Manager (CPM): The civilian and military co-leads accountable for the execution of capability portfolio management activities for a defined portfolio. (DoDD 7045.20)

Client-Server (architecture): A model for computing where a server performs all the processing of data according to data requests sent from client software installed at the user.

Cloud Computing: An arrangement where a service that is provided to a user over the network is run (hosted) on a third-party server that is not owned or operated by the provider of the service.

Command and Control (C2): The exercise of authority and direction by a properly designated commander over assigned and attached forces and resources in the accomplishment of the mission. (JP 1-02, modified to reflect current JROC approved/DAWG endorsed JCA language.)

Community: A group of users that cooperates to accomplish a mission.

Community of Interest: A collaborative group of users that routinely shares information in pursuit of its shared goals, interests, missions, or business processes and therefore must have shared vocabulary for the information it exchanges. (C2 Core IAT report, December 2008)

Community Service: A service used by a particular community of users.

CONOPS (Concept of Operations): A verbal or graphic statement that clearly and concisely expresses what the joint force commander intends to accomplish and how it will be done using available resources. The concept is designed to give an overall picture of the operation. Also called commander's concept or CONOPS. (JP 5-0)

UNCLASSIFIED

Concept: A unit of thought. The semantic content of a concept can be re-expressed by a combination of other and different concepts, which may vary from one language or culture to another. Concepts exist in the mind as abstract entities, which are independent of the terms used to label them. (<http://www.willpowerinfo.co.uk/glossary.htm#concept>)

Conceptual Model: A map of concepts and their relationships. These models describe the semantics of an organization and represent a series of assertions about its nature. Specifically, they describe the things of significance to an organization (entity classes) (about which it is inclined to collect information), and characteristics of (attributes) and associations between pairs of those things of significance (relationships).

Controlled Vocabulary (CV): prescribed list of terms or headings each one having an assigned meaning; (<http://www.willpowerinfo.co.uk/glossary.htm#controlledvocabulary>) A complete set of allowed terms that can be used for describing something (e.g., in a registry). There may be many possible values to choose from for each attribute of the described object, but the user may not add new values.

Crawler: Software that recursively follows URLs in documents to find new documents.

Customer: A CONOPS role that includes providing requirements and (potentially) resources for local instantiations of C2 services, negotiating Service Level Agreements (SLAs) with the Service Owner (when appropriate), and advocating for capability needs.

Data: A collection of characters that has meaning in some context (e.g., a news article), even if only according to an unregistered schema.

Data Model: A graphical or lexical representation of data, specifying their properties, structure, and interrelationships. (C2 Core IAT report, December 2008)

DDMS: The Department of Defense Discovery Metadata Specification defines discovery metadata elements for resources posted to community and shared spaces. (<https://metadata.dod.mil/mdr/irs/DDMS/>)

Discovery: The act of locating a description of a service-related resource that may have been previously unknown and that meets certain functional criteria. It involves matching a set of functional and other criteria with a set of resource descriptions.

UNCLASSIFIED

Domain: (1) A subset of a Mission Area that represents a common collection of related, or highly dependent, information capabilities and services; (2) a logical grouping of addresses in a network.

“drill into” (something): To obtain more detailed information (about something).

Echelon (of command): (1) A subdivision of a headquarters, i.e., forward echelon, rear echelon; (2) Separate level of command. As compared to a regiment, a division is a higher echelon, a battalion is a lower echelon; (3) A fraction of a command in the direction of depth to which a principal combat mission is assigned; i.e., attack echelon, support echelon, reserve echelon; (4) A formation in which its subdivisions are placed one behind another, with a lateral and even spacing to the same side. (JP 1-02)

Endpoint: An address where an instance of a service may be found.

Enterprise Infrastructure Services: Content or mission-neutral, general-purpose capabilities—designed for use by and continuously available to all organizations within the DoD and selected mission partners—that provide for basic communications, collaboration, publication, discovery, security, and information and service management.

Expose (a service): To make a service “visible” (see definition of “visible”) by including an address where it may be found in its registered description.

Extensible Markup Language (XML): XML is a structured language for describing information being sent electronically by one entity to another. XML Schema defines the rules and constraints for the characteristics of the data, such as structure, relationships, allowable values, and data types. (C2 Core IAT report, December 2008)

Federation (of services): (1) The facilitated use of processing and data services that are distributed across a network of largely autonomous nodes with a set of rules but without central ownership. This means that there is an agreement among participants on how to exchange or invoke distributed processing and data services to perform an operational or technical task or create a composite service. Federation participants (e.g., processing nodes on a network) are often called federates and execute under local resource control and management. In addition, the degree of federation of services can vary from loose to tight in terms of both governance and the technicalities of implementation and as a function of both user operational requirements and the maturity of the SOE. (2) The evolution toward or development of a set of services that compose a federation from a set that does not.

UNCLASSIFIED

Fielding: Providing a capability to intended users (to “troops in the field”).

Fuse: Synthesize

“gain traction”: Become used or accepted.

GIG (Global Information Grid): The globally connected, end-to-end set of information capabilities, associated processes, and personnel for collecting, processing, storing, disseminating, and managing information on demand to warfighters, policy makers, and support personnel. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense) The Global Information Grid (GIG) includes all owned and leased communications and computing systems and services, software (including applications), data, security services, and other associated services necessary to achieve information superiority. It also includes National Security Systems as defined in section 5142 of the Clinger-Cohen Act of 1996. The GIG supports all Department of Defense (DoD), National Security, and related intelligence community missions and functions (strategic, operational, tactical, and business), in war and in peace. The GIG provides capabilities from all operating locations (bases, posts, camps, stations, facilities, mobile platforms, and deployed sites). The GIG provides interfaces to coalition, allied, and non-DoD users and systems. Also called GIG. (JP 3-05.1)

GIG 2.0 (Global Information Grid version 2.0): The set of all communicating hardware and software owned by the DoD, including social networking upgrades.

“Go Viral”: To become wildly popular very quickly. Implies an initially exponential increase in popularity.

Governance: Consistent management, cohesive policies, processes and decision-rights for a given area of responsibility. (<http://en.wikipedia.org/wiki/Governance>)

Granular: Specific, descriptive.

“ground truth”: Perfectly accurate information.

Icon: A symbol used to represent an entity. Usually used in this context to refer to tiny pictures that represent military units or hardware on a map.

Information: (1) Data that conforms to (is expressed according to) a known schema. In the case of Information Services, the schema must be discoverable (e.g., registered

UNCLASSIFIED

somewhere). In current technology, information is data that have an associated meaning imparted by incorporated labels (such as XML tags) that use terms defined in readily available (published) artifacts (in this case, the C2ISF Artifacts), or by being associated with published ontologies in a standard resource-description-tagging fashion. (NCSS Study) (2a) Facts, data, or instructions in any medium or form. (2b) The meaning that a human assigns to data by means of the known conventions used in their representation. (JP 3-13.1)

Information Assurance (IA): Information operations that protect and defend information and information systems by ensuring their availability, integrity, authentication, confidentiality, and non-repudiation. This includes providing for restoration of information systems by incorporating protection, detection, and reaction capabilities. Also called IA. (JP 3-13)

Information Service: See Service.

Information Sharing: The sharing of information (see “Information”) using formal, standard protocols.

Information Technology (IT): Any equipment or interconnected system or subsystem of equipment that is used in the automatic acquisition, storage, manipulation, management, movement, control, display, switching, interchange, transmission, or reception of data or information. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

Instance (of a service): A separate copy of the original service that performs the same as the original but can be independently called.

Interdependent: Only able to perform a required mission by coordinating actions with another entity that also cannot perform the required mission by itself.

Interoperability: The ability of systems, units, or forces to provide data, information, materiel, and services to and accept the same from other systems, units, or forces and to use the data, information, materiel, and services so exchanged to enable them to operate effectively together. National Security System (NSS) and Information Technology System (ITS) interoperability includes both the technical exchange of information and the operational effectiveness of that exchanged information as required for mission accomplishment. (CJCSI 6212.01E) (1) The ability to operate in synergy in the execution of assigned tasks. (2) The condition achieved among communications-electronics systems or items of communications-electronics equipment when information

UNCLASSIFIED

or services can be exchanged directly and satisfactorily between them and/or their users. The degree of interoperability should be defined when referring to specific cases. (JP 3-32)

Invoke (a service): To call a service; to send a detailed request to a service that is designed such that the service can be expected to respond with the answer to the request.

JCIDS (Joint Capabilities Integration and Development System): Three key processes in the DoD must work in concert to deliver the capabilities required by the warfighter: the requirements process; the acquisition process; and the Planning, Programming, Budget, and Execution (PPBE) process. JCIDS implements the requirements process. JCIDS supports the Chairman of the Joint Chiefs of Staff (CJCS) and the Joint Requirements Oversight Council (JROC) in identifying, assessing, and prioritizing joint military capability needs as required by law. The capabilities are identified by analyzing what is required across all joint capability areas to accomplish the mission. (CJCSI 3170.01G and JCIDS Manual)

JFCOM J3: United States Joint Forces Command director for operations, plans, logistics and engineering.

Joint C2 Conceptual Model and Vocabulary: Model and vocabulary, used via services, which publish descriptions of C2 entities and their interrelationships, together with terms and definitions that express properties of those entities.

JOPES (Joint Operation Planning and Execution System): DoD's principal tool for designing, monitoring the progress of, and managing (generally large-scale) force deployments.

Knowledge Organization System (KOS): Systems that are designed to provide a standard way of describing a group of information and that can comprise any numbers of Controlled Vocabularies (CVs) and Taxonomies.

“lat/lon”: A latitude and longitude pair (a position on the surface of the earth).

Leverage: To make use of existing assets or capabilities to aid in the performance of a new task.

Loose Coupling: A situation where the set of factors that a system has to comply with in order to consume the features or services provided by other systems has been reduced to

UNCLASSIFIED

a minimum. (from <http://www.w3.org/2003/glossary/keyword/All/?keywords=loose+coupling>)

Mash-up: A web page or application that combines data or functionality from two or more external sources to create a new service. The term “mash-up” implies easy, fast integration, frequently using open APIs and data sources to produce results that were not the original reason for producing the raw source data. An example of a mash-up is the use of cartographic data to add location information to real estate data, thereby creating a new and distinct Web API that was not originally provided by either source. ([http://en.wikipedia.org/wiki/Mashup_\(web_application_hybrid\)](http://en.wikipedia.org/wiki/Mashup_(web_application_hybrid)))

Mediation: Expressing data conformant to one schema according to another.

Mediators: Information services that perform mediation.

Metadata: Information describing the characteristics of data; data or information about data; or descriptive information about an entity’s data, data activities, systems, and holdings. For example, discovery metadata is a type of metadata that allows data assets to be found using enterprise search capabilities. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

Metadata Registry: Repository of all metadata related to data structures, models, dictionaries, taxonomies, schema, and other engineering artifacts that are used to support interoperability and understanding through semantic and structural information about the data. A federated metadata registry is one in which multiple registries are joined electronically through a common interface and exchange structure, thereby effecting a common registry. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

Metric: A quantity that can be reliably computed from observations of an entity of interest.

(Database) Mining: Using sophisticated “data mining” algorithms that find patterns, trends, and correlations in large amounts of data (in a database, in this case).

Mission Area: A defined area of responsibility with functions and processes that contribute to mission accomplishment. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

UNCLASSIFIED

Mission Partners: Those entities not under the commander's direct authority that are participating in the mission. Some examples include, but are not limited to, supported/supporting commands, non-DoD Government agencies such as the Department of State, CIA, or Department of Homeland Security, coalition partners, U.S. and host nation civil authorities, international organizations, and nongovernmental organizations (NGOs). (C2IP)

(governed) Monolithically: As though the object was completely uniform, with no differences throughout.

NCSS Study: The IDA study that resulted in this document.

Near Real Time (NRT): (1) Pertaining to the timeliness of data or information delayed only by the time required for electronic communication. This implies there are no noticeable delays. (JP 1-02) (2) Timeliness of data or information delayed only by the time required for electronic communication. This implies there are no noticeable delays. Data is real time when current active tracks show current location, updates occur immediately, and the only delay is of electronic communication. (CJCSI 3151.01)

Net-Centric: Relating to or representing the attributes of net-centricity. Net-centricity is a robust, globally interconnected network environment (including infrastructure, systems, processes, and people) in which data sharing is timely and seamless among users, applications, and platforms. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

Net-Enabled: Facilitated through the use of information technology (IT) systems interconnected via a communication network or network of networks.

NetOps: (1) The activities performed to ensure the GIG is operating smoothly (with minimal interruptions to user's activities). (NCSS Study) (2) The operational framework consisting of the essential tasks, situational awareness (SA); and C2 that CDRUSSTRATCOM employs to operate and defend the GIG. The essential tasks are GIG Enterprise Management (GEM), GIG Network Defense (GND), and GIG Content Management (GCM). NetOps and its essential tasks, GEM, GND, and GCM, include IA as defined and outlined in DoDD 8500.1, Information Assurance, and CJCSI 6510.01D, Information Assurance and Computer Network Defense. (Joint CONOPS for GIG NetOps, 4 Aug 2006)

Ontology: Specification of the concepts of a domain and their relationships, structured to allow computer processing and reasoning (<http://www.willpowerinfo.co.uk/glossary>).

UNCLASSIFIED

[htm](#)); an explicit specification of how to represent objects and concepts and the relationships among them. An ontology may be used to describe the relationships between the data elements in a schema, and those relationships are used to derive the meaning of data that conforms to the schema.

Operational Process Owner (OPO): A CONOPS role that includes being the source of operational process definition and TTPs and determining how information services will be used to support operational processes, being the source of authoritative capability needs for services (and required improvements thereto) in support of C2 missions, and advocating for resources (and, in some cases, providing resources).

Orchestration: The coordinating of services to perform an overall task.

Orthogonalized: Made so the constituent parts are independent.

Other-Domain Services/Other-Community Services: Capabilities that fulfill data and/or functionality requirements for non-C2 services, but that may be useful for C2 missions.

Portfolio: See Capability Portfolio.

Portfolio Management: See Capability Portfolio Management.

“Piggy-Backing” (on something): Being wholly dependent on and taking advantage of (something that exists).

Precision (Search Precision): The number of documents returned by the search that are relevant to the user’s search criteria divided by the total number of documents returned.

Rationalize: To make compatible or consistent.

Real Time (RT): (1) Pertaining to the timeliness of data or information delayed only by the time required for electronic communication. This implies there are no noticeable delays. (JP 1-02) (2) Timeliness of data or information delayed only by the time required for electronic communication. This implies there are no noticeable delays. Data are real time when current active tracks show current location, updates occur immediately, and the only delay is of electronic communication. (CJCSI 3151.01)

UNCLASSIFIED

Recall (Search Recall): The number of relevant documents returned by the search divided by the total number of relevant documents that exist.

Recursion: A method of defining functions in which the function being defined is applied within its own definition; specifically it is defining an infinite statement using finite components. (<http://en.wikipedia.org/wiki/Recursion>)

Recursively: In such a manner that the defined action is performed in the course of the performing the defined action. For example, to recursively follow URLs on a Website, one must (a) follow all the URLs on the site, (b) follow all the URLs on the pages obtained in step a, (c) follow all the URLs on the pages obtained in step b, (d) and so on. (see “Recursion”)

Registry: A service that provides official descriptions of services or documents in a standard format and lexicon (i.e., that conform to the set of artifacts defining the terms in the registry).

Registration: Storing a description of a resource (e.g., a service or Website) in a registry.

Repository: An information system used to store and access information, schemas, style sheets, controlled vocabularies, dictionaries, and other work products. It would normally be discovered via a registry. (C2 Core IAT report, December 2008)

Resource: (1) Funds, equipment, or personnel that may be used for a process. (2) A service or source of information that is visible on the network.

REST: Representation State Transfer. “A RESTful Web service (also called a RESTful Web API) is a simple Web service implemented using HTTP and the principles of REST. Such a Web service can be thought about as a collection of resources.” “The REST architectural style describes the following six constraints applied to the architecture, while leaving the implementation of the individual components free to design: Client-server ... Stateless ... Cacheable ... Layered System ... Code on Demand (optional) ... Uniform Interface ...” (http://en.wikipedia.org/wiki/Representational_State_Transfer)

Role: A set of responsibilities that are performed by a single entity.

Run-time: In computer science, the duration of a computer program’s execution, from beginning to termination. In the current context, it means that if a warfighter generates a short script (program) that utilizes one or more “enterprise service(s)” to accomplish

UNCLASSIFIED

some end, when the script/program executes, the required service functions are accessible and will perform the advertised service. (DoD CIO Memorandum, March 11, 2009, Interim Implementation Guidance for the Net Centric Data Strategy (NCDS) in the Command and Control (C2) Capability Portfolio)

Scalable: Able to be readily implemented in a larger network or more populated scenario.

Schema: A format for expressing data. A schema can represent any generic model or structure that deals with the organization, format, structure, or relationship of data.

Service/Information Service: A means of facilitating a mission outcome through the provision of information or data-processing support without the customer owning the mechanism for providing that support. Generally, information services will provide access to data, computational or transactional functions, or management or orchestration functions. The more general ITIL (v.3) definition of a Service is “a means of delivering value to customers by facilitating outcomes customers want to achieve without the ownership of specific costs and risks.”

Service-Level Agreement: the SLA records a common understanding about services, priorities, responsibilities, guarantees, and warranties. Each area of service scope should have the "level of service" defined. The SLA may specify the levels of availability, serviceability, performance, operation, or other attributes of the service, such as billing. (http://en.wikipedia.org/wiki/Service_level_agreement)

Service Oriented Architecture (SOA): SOA is an architectural style whose goal is to achieve loose coupling among interacting software agents. A service is a unit of work done by a service provider to achieve desired end results for a service consumer. Both provider and consumer are roles played by software agents on behalf of their owners. SOA is characterized by on-demand services. Participants in a SOA make their resources available by publishing information in structured formats that describe their capabilities and how to access them. Other participants can discover and request those services on demand, but have no power to modify their makeup (other than by feeding back suggestions), ensuring their capabilities always remain available to other participants. This loosely coupled, on-demand assembly of resources has the advantage of being highly adaptable to change. (C2IP)

Service-Oriented Enterprise (SOE): An enterprise that combines a services-focused way of doing business with the latest technology in an operational culture where participating entities include both service providers and service consumers. This implies a broader and

UNCLASSIFIED

less technically prescriptive approach to providing and consuming services than is generally implied by usage of the term Service-Oriented Architecture (SOA).

Serialization: The encoding of a logical structure into a regular format.

Service Owner: A CONOPS role that includes being responsible for creating, acquiring, fielding, managing, supporting, and improving the information service throughout the life cycle, identifying and managing the service provider, negotiating SLAs with customers and functional requirements with OPO, and providing support to the OPO and users throughout the life cycle.

Situational Awareness: The degree of accuracy by which one's perception of his current environment mirrors reality. It is the knowledge, cognition, and anticipation of events, factors, and variables affecting the safe, expedient, and effective conduct of the mission. It is developed through the continuous integration of new observations into recurring mental assessments. (C2 Joint Integrating Concept, v. 1.0, 1 Sep 2005).

SKOS: A family of formal languages designed for representation of thesauri, classification schemes, taxonomies, subject-heading systems, or any other type of structured controlled vocabulary. SKOS is built upon RDF and RDFS, and its main objective is to enable easy publication of controlled structured vocabularies for the Semantic Web. SKOS is currently developed within the W3C framework. (<http://en.wikipedia.org/wiki/SKOS>; <http://www.w3.org/2004/skos/>; <http://www.w3.org/TR/skos-reference/>)

stage: To process, in a specified area, troops that are in transit from one locality to another. (JP 1-02)

staging: Assembling, holding, and organizing arriving personnel, equipment, and sustaining materiel in preparation for onward movement. The organizing and preparation for movement of personnel, equipment, and materiel at designated areas to incrementally build forces capable of meeting the operational commander's requirements. (JP 3-35)

Standard: Quantitative or qualitative measures for [specifying] the levels of performance of a task. (CJCSI 3010.02B)

“Stovepiped” (system): A system that was not designed to be easily made interoperable with other systems.

UNCLASSIFIED

Tagging Engine: A service that can scan through a document to find data it recognizes as being consistent with a given schema, and includes that data as derived metadata about the document.

Taxonomy: Monohierarchical classification of concepts, as used, for example, in the classification of biological organisms (<http://www.willpowerinfo.co.uk/glossary.htm>). A grouping of terms representing topics or subject categories. A taxonomy is typically structured so that its terms exhibit hierarchical relationships to one another, between broader and narrower concepts. Taxonomy structure is discussed in the NISO Z39.19 (2005) standard. (<http://www.dataharmony.com/library/taxonomyGlossary.html>)

Thesaurus: A controlled vocabulary in which concepts are represented by preferred terms, formally organized so that paradigmatic relationships between the concepts are made explicit, and the preferred terms are accompanied by lead-in entries for synonyms or quasi-synonyms. (The purpose of a thesaurus is to guide both the indexer and the searcher to select the same preferred term or combination of preferred terms to represent a given subject.) (<http://www.willpowerinfo.co.uk/glossary.htm#thesaurus>)

Track: A collection of the positions of an object, the times the object was at those positions, and information describing the object. (NCSS Study) (1) A series of related contacts displayed on a data display console or other display device. (2) To display or record the successive positions of a moving object.... (5) The actual path of an aircraft above or a ship on the surface of the Earth. The course is the path that is planned; the track is the path that is actually taken. (JP 1-02)

Translators/Schema Transformers: Information services that convert information expressed according to one schema into information expressed according to another.

TTPs (Tactics, Techniques, and Procedures): A formal set of rules and suggestions that military personnel use for the conduct of their operations that store lessons learned from a large number of previous situations.

UDDI: A platform-independent, Extensible Markup Language (XML)-based registry.... UDDI was originally proposed as a core Web service standard. It is designed to be interrogated by Simple Object Access Protocol (SOAP) messages and to provide access to Web Services Description Language (WSDL) documents describing the protocol bindings and message formats required to interact with the Web services listed in its directory. (<http://en.wikipedia.org/wiki/UDDI>)

UNCLASSIFIED

User: A CONOPS role that includes using the service in question when executing missions.

Universal Core (U-Core): A common description of entities that the DoD, IC, DoJ and DHS can use when creating new schema. It is an interagency information exchange specification and implementation profile. It provides a framework for sharing the most commonly used data concepts of “who, what, when, and where.” It serves as a starting point for data-level integration and permits the development of richer domain specific exchanges. It was created and is managed by DoD, DOJ, DHS, and the Intelligence Community. (C2IP)

Universal Core Semantic Layer (U-Core SL): An ontological model of the entities in DoD, IC, DoJ, and DHS that includes both terms describing them and the relationships between them. It is an elaboration of the high-level U-Core schema.

Vocabulary: A set of terms, headings or concept codes and their inter-relationships which may be used to support information retrieval. (http://www.willpowerinfo.co.uk/glossary.htm#structured_vocabulary) It represents agreements on the terms and definitions common to the COI, including data dictionaries. For example, one COI might define the term “tank” to mean a pressurized vessel, whereas another might define “tank” to mean a tracked vehicle. Both definitions are acceptable, but the user must understand these definitions, and their context, to properly use the data.

Visible: Able to be perceived and, to some extent, characterized by humans and/or IT systems, applications, or other processes. Visibility does not imply actual access to service-provided data or processing capabilities.

Visualization: The representation of information in an intuitive graphical format (that aids comprehension).

Web Services: A standardized way of integrating Web-based applications using open standards over an Internet Protocol backbone. Web services allow applications developed in various programming languages and running on various platforms to exchange data without intimate knowledge of each application’s underlying IT systems. (DoDD 8320.02, December 2, 2004, Data Sharing in a Net-Centric Department of Defense)

Whiteboard: The successor to the chalkboard. It is a surface that can be written on with special markers and easily erased.

UNCLASSIFIED

(This page is intentionally blank.)

UNCLASSIFIED

UNCLASSIFIED

PART 2—ACRONYMS

AAT	Air Tasking Order / Air Control Order Tools
ABCS	Army Battle Command System
ACID	Atomicity, Consistency, Isolation, Durability (see Glossary)
ACoA	Adaptive Course of Action
ACP	Air Campaign Planning
ACTD	Advanced Concept Technology Demonstrations
AD	Airspace Deconfliction
ADS	Authoritative Data Source
AFB	Air Force Base
AFRICOM	USAFRICOM
AFWA	Air Force Weather Agency
AMC	Air Mobility Command
ANSI	American National Standards Institute
AOC	Air Operations Center
AoR	Area of Responsibility
APEX	Adaptive Planning and Execution
API	Application Programming Interface
ASD(NII)/DoD CIO	Assistant Secretary of Defense for Networks and Information and Integration/DoD Chief Information Officer.
ASD/(NII)	Assistant Secretary of Defense for Networks & Information Integration
ASD/NII	Assistant Secretary of Defense for Networks & Information Integration
ATO	Air Tasking Order
BDA	Battle Damage Assessment
BFO	Basic Formal Ontology
BTI	Business Transformation Infrastructure
C&A	Certification & Accreditation
C2 CPM	Command & Control Capability Portfolio Manager
C2 SOE	Command & Control Service-Oriented Environment
C2	Command & Control
C2CS	Command & Control Core Services

UNCLASSIFIED

C2ISF	Command & Control Information Sharing Framework
C3	Command, Control, and Communications
C3S&S	Command, Control, Communications, Space and Spectrum
CA BN	Civil Affairs Battalion
CAMPS	Consolidated Air Mobility Planning System
CANES	Consolidated Afloat Networks and Enterprise Systems
CAOC	Combined Air Operations Center
CAPE	Cost Assessment and Program Evaluation Office
CAS	Close Air Support
CAST	Close Air Support Tool
CENTCOM	USCENTCOM
CFLCC	Coalition Forces Land Component Command
CIDNE	Combined Information Data Network Exchange
C-IED	Counter-IED
CIO	See DoD CIO
CJCS	Chairman of the Joint Chiefs of Staff
CMCC	Corps Movement Control Center
CO	Commanding Officer
COBIT	Control Objectives for Information and related Technology
COCOM	Combatant Command
CoI, COI	Community of Interest
CONOPS	Concept of Operations
CONPLAN	Concept Plan
CONUS	Continental United States
COTS	Commercial Off-The-Shelf
CP	Command Post
CPM	Capability Portfolio Manager
CPOF	Command Post of the Future
CRUD	Create, Read, Update, Delete (see Glossary)
CV	Controlled Vocabulary
CYBERCOM	USCYBERCOM
DARPA	Defense Advanced Research Projects Agency
DASD	Deputy Assistant Secretary of Defense
DDMS	DoD Discovery Metadata Specification

UNCLASSIFIED

DHS	Department of Homeland Security
DIL	Disconnected and Intermittent Connection or Low Bandwidth.
DISA	Defense Information Systems Agency
DNS	Domain Name System
DoD CIO	Department of Defense Chief Information Officer
DoD	Department of Defense
DoDAAC	DoD Activity Address Code
DoDD	Department of Defense Directive
DoJ, DOJ	Department of Justice
DOTMLPF	Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, and Facilities
DTM	Design Tasking Memorandum
ebXML	Electronic Buisness using XML
EIS	Enterprise Information Services
EM-C	Execution Management—Control
EM-R	Execution Management—Replanning
ESC	Electronic Systems Command
EUCOM	USEUCOM
FM	Frequency Modulation
FORSCOM	USFORSCOM
GCCS	Global Command and Control System
GCCS-A	Global Command and Control System—Army
GCCS-J	Global Command and Control System—Joint
GCM	GIG Content Management
GEM	GIG Enterprise Management
GFM	Global Force Management
GIG	Global Information Grid
GND	GIG Network Defense
GOTS	Government, Off-The-Shelf
GPS	Global Positioning System
GTN	Global Transportation Network

UNCLASSIFIED

HQ	Headquarters
HTTP	Hypertext Transfer Protocol
HW	Hardware
IA	Information assurance
IAT	Independent Assessment Team
IAW	Abbreviation of “In Accordance With”
IC	Intelligence Community
IC-ISM	Intelligence Community Information Security Marking
ID	Identification
IED	Improvised Explosive Device
IMIT	International Military Information Team
IOC	Initial Operational Capability
IOT&E	Initial Operational Test & Evaluation
IPB	Intelligence Preparation of the Battlefield
ISACA	Information Systems Audit and Control Association
ISO	International Organization for Standardization
ISR	Intelligence, Surveillance, and Reconnaissance
IT	Information Technology
ITGI	IT Governance Institute
ITIL	Information Technology Infrastructure Library
ITSM	Information Technology Service Management
J3	(see JFCOM J3 in Glossary)
JCA	Joint Capability Area
JCD	Joint Capabilities Document
JCIDS	Joint Capabilities Integration and Development System
JCTD	Joint Capability Technology Demonstrations
JDP	Joint Defense Planning
JFACC	Joint Force Air Component Commander
JFCOM J3	(see Glossary)
JFCOM	USJFCOM
JNN	Joint Network Node
JOPEs	Joint Operation Planning and Execution System
JROC	Joint Requirements Oversight Council

UNCLASSIFIED

JSIC	Joint Systems Integration Center (part of USJFCOM)
JTF	Joint Task Force
JTF-GNO	Joint Task Force—Global Network Operations
KOS	Knowledge Organization System
“lat/lon”	Latitude / Longitude
LC	Life Cycle
LCSH	Library of Congress Subject Heading taxonomy
LoC	Library of Congress
MCO	Movement Control Officer
MCS	Maneuver Control System
MDAP	Major Defense Acquisition Program
MDR	Metadata Registry
METOC	Meteorological and Oceanographic
MVC	Model-View-Controller
NCDS	Net-Centric Data Strategy (May 2003)
NCES	Net-Centric Enterprise Services
NCSS	Net-Centric Services Strategy (May 2007)
NECC	Net-Enabled Command Capability
NetOps	Network Operations
NGO	Non-Governmental Organization
NII	See OASD/NII
NISO	National Information Standards Organization
NSA	National Security Agency
OASD(NII)/DoD CIO	Office of the Assistant Secretary of Defense for Networks and Information Integration/DoD Chief Information Officer
OASD/NII	Office of Assistant Secretary of Defense for Networks & Information Integration
OASIS	Organization for the Advancement of Structured Information Standards
OEF	Operation Enduring Freedom

UNCLASSIFIED

OGC	UK Office of Government Commerce
OIF	Operation Iraqi Freedom
ONS	Operational Needs Statement
OODA	Observe, Orient, Decide, Act
OPLAN	Operational Plan
OPO	Operational Process Owner
ORNL	Oak Ridge National Laboratory
OS	Operating System
OSD/NII	See OASD/NII
OWL	Web Ontology Language
PACOM	USPACOM
PEO	Program Executive Officer
PID	Plan Identification Number
PKI	Public Key Infrastructure
PM	Program Manager
PoC	Point of Contact
PoR	Program of Record
PPBE	Planning, Programming, Budgeting, and Execution
R&D	Research & Development
RACI	Responsible, Accountable/Approver, Consulted, Informed
RDBMS	Relational Database Management System
RDF	Resource Description Framework
REF	Rapid Equipping Force
REST	Representation State Transfer (see Glossary)
RoI	Return on Investment
RQT	Rapid Query Tool
S&M	Scheduling & Movement
SA	Situational Awareness
SAWSDL	Semantic Annotations for WSDL
SBMCS	Space Battle Management Core Systems
SigAct	Significant Activity
SIPRNet	Secret Internet Protocol Router Network

UNCLASSIFIED

SKOS	Simple Knowledge Organization System
SLA	Service-Level Agreement
SOA	Service-Oriented Architecture
SOAP	Simple Object Access Protocol
SOE	Services-Oriented Enterprise
SoI	Sons of Iraq
SOP	Standard Operating Procedure
SOUTHCOM	USSOUTHCOM
SQL	Structured Query Language
SW	Southwest
TACC	Tanker Airlift Control Center
TAP	Theater Air Planning
TBMCS	Theater Battle Management Core Systems
TC4I	Tactical Command, Control, Communications, Computers, and Intelligence
TCN	Transportation Control Number
TCP/IP	The Internet Protocol Suite (including the Transmission Control Protocol and the Internet Protocol)
TCT	Time-Critical Targeting
TIGR	Tactical Ground Reporting system
TMCC	Theater Movement Control Center
TPFDD	Time-Phased Force Deployment Data
TRADOC	United States Army Training and Doctrine Command
TRANSCOM	USTRANSCOM
TTP	Tactics, Techniques, and Procedures
U-Core SL	Universal Core Semantic Layer
U-Core	Universal Core
UDDI	Universal Description, Discovery and Integration
UI	User Interface
UIC	Unit Identification Code
ULN	Unit Line Number
UON	Urgent Operational Need
URI	Universal Resource Indicator

UNCLASSIFIED

URL	Universal Resource Locator
URN	Universal Resource Name
US	United States
USAFRICOM	United States Africa Command
USARCENT	United States Army Forces, United States Central Command
USCENTCOM	United States Central Command
USCYBERCOM	United States Cyber Command
USEUCOM	United States European Command
USFORSCOM	United States Forces Command
USJFCOM	United States Joint Forces Command
USPACOM	United States Pacific Command
USSOUTHCOM	United States Southern Command
USTRANSCOM	United States Transportation Command
W3C	World Wide Web Consortium
WebTAS	Web-enabled Temporal Analysis System
WSDL	Web Services Description Language
WS-Policy	Web Services Policy
XML	Extensible Markup Language
XSD	XML Schema Document

UNCLASSIFIED

Appendix D
LIST OF FIGURES AND TABLES

UNCLASSIFIED

LIST OF FIGURES AND TABLES

LIST OF FIGURES

ES-1. Major Components of Proposed C2 Services Implementation Approach.....	ES-5
1. Net-Centric Services Strategy (NCSS) Highlights	2
2. Conceptual C2 Services Framework.....	10
3. C2ISF Concept Diagram.....	21
4. Detailed C2ISF Diagram.....	22
5. Detailed C2ISF Diagram with Information Exchanges.....	29
6. Harmonizing NCDS and NCSS Implementation Activities.....	31
7. Developing the C2 Services CONOPS.....	36
8. C2 Service Tiers.....	38
9. C2 Services Tiers and Varying Characteristics.....	40
10. Variation in C2 Service Tier Characteristics Across Life Cycle	46
11. Key C2 Services CONOPS Roles.....	48
12. Example of CONOPS Roles	51
13. Summary of Principal Actors for C2 Services Across the Life Cycle.....	54
14. C2 Services Portfolios.....	55
15. C2 Services Portfolios Governance Characteristics	56
16. C2 Services Portfolios Governance Concept.....	58

UNCLASSIFIED

LIST OF TABLES

1. C2ISF Service Description Templates' Artifacts and Concepts.....	24
2. Example Data Requested by C2ISF Service Description Templates	25
3. Representative C2ISF Interface Exchanges.....	29
4. Definitions of C2 Services Tiers.....	39
5. Key Characteristics of C2 Services by Tier.....	41
6. C2 Services Life Cycle for CONOPS.....	43
7. Service Tier Characteristics by Life Cycle Phase.....	44
8. Principal Roles in C2 Services CONOPS Life Cycle.....	47
9. Typical Actors for C2 Service Roles	49
10. Analyzed Activities by Life Cycle Stage.....	52
11. Key C2 Services Roles and Responsibilities	53

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) February 2010		2. REPORT TYPE Final		3. DATES COVERED (FROM - TO) April 2009—December 2009	
4. TITLE AND SUBTITLE DoD Net-Centric Services Strategy Implementation in the C2 Domain				5A. CONTRACT NO. DASW01-04-C-0003	
				5B. GRANT NO.	
				5C. PROGRAM ELEMENT NO(S).	
6. AUTHOR(S) Walsh, P.J.; Davis, S.O.; Keifer, M.; Morrison, K.A.; Pipher, J.W.; Potrykus, H.G.				5D. PROJECT NO.	
				5E. TASK NO. BC-1-2526	
				5F. WORK UNIT NO.	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				8. PERFORMING ORGANIZATION REPORT NO. IDA Paper P-4549	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Director, Command and Control Programs and Policy Crystal Mall 3, 7 th Floor 1931 Jefferson Davis Highway Arlington, VA 22202				10. SPONSOR'S / MONITOR'S ACRONYM(S) ASD NII	
				11. SPONSOR'S / MONITOR'S REPORT NO(S).	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES Philip J. Walsh Project Leader					
14. ABSTRACT This study describes an approach for implementing C2 information support capabilities as services in compliance with the DoD Net-Centric Services Strategy (NCSS). The recommended approach includes implementation and use of the following: (1) a C2 Information Sharing Framework (C2ISF) that identifies the design-time artifacts and run-time infrastructure services needed to accomplish NCSS goals; (2) a tiered organizational structure for implementing and managing C2 information services as part of an evolving Service-Oriented Enterprise (SOE); (3) a C2 services CONOPS that ties together C2 service tiers, key implementation roles and actors; and (4) a high-level governance concept. In addition, as a result of examining how C2 information support capabilities are being implemented in the operating forces today, the recommended approach emphasizes the importance of enabling and encouraging edge-user agility and innovation in developing and improving C2 services while implementing the SOE.					
15. SUBJECT TERMS services, C2 services, C2 information services, net-centric services, net-centric services strategy, C2 capabilities, C2 information support capability, management of IT services, governance of IT services, life cycle support of IT services					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Unlimited	18. NO. OF PAGES 170	19A. NAME OF RESPONSIBLE PERSON Mr. Ronald W. Pontius
A. REPORT Unclassified	B. ABSTRACT Unclassified	C. THIS PAGE Unclassified			19B. TELEPHONE NUMBER (INCLUDE AREA CODE) 703-607-0670