

Chapter 6

Developing Interoperability Metrics

By
Dr. John A. Hamilton, Jr., Auburn University
Captain Jerome D. Rosen, US Air Force Reserve
Major Paul A. Summers, US Air Force

The views expressed in this chapter are the opinions of the authors, and do not reflect the official opinions of any U.S. government agency.

6.1 INTRODUCTION

Modern military operations require interoperability. The DoD has made tremendous interoperability gains over the last few years. Unfortunately, without a way to assess the status of interoperability throughout the department, it is difficult to quantify this progress. While interoperability issues are persistent and visible, the number of interoperability successes is easily overlooked. Most systems that are developed today meet the interoperability requirements that were specified in their Operational Requirements Document (ORD). The application of a set of metrics addressing this domain would shed more light on the situation and highlight the successes of the many agencies that have labored to produce interoperable systems. Effective metrics would enable the services and agencies to make informed decisions about the allocation of scarce resources to solve interoperability in already fielded systems.

This chapter begins by discussing some of the challenges involved in developing effective interoperability metrics. Next, it proposes a set of metrics that, the authors believe, would help realize the objectives described above. Finally, institutional and programmatic issues surrounding metrics are discussed.

6.2 BACKGROUND

The pejorative use of the term “legacy system” often occurs when describing communications and computer systems. This is unfortunate since many fielded systems are performing well and meeting or exceeding their original specifications. C4ISR refers to systems that are part of the Command, Control, Communications, Computer, Intelligence, Surveillance and Reconnaissance domain. The C4ISR domain is one of four domains for which the Joint Technical Architecture (JTA) specifies a domain annex. C4ISR is defined in the JTA [JTA 99] as those systems which:

- Support properly designated commanders in the exercise of authority and direction over assigned and attached forces across the range of military operations.
- Collect, process, integrate, analyze, evaluate, or interpret available information concerning foreign countries or areas.
- Systematically observe aerospace, surface or subsurface areas, places, persons, or things by visual, aural, electronic, photographic, or other means.
- Obtain, by visual observation or other detection methods, information about the activities and resources of an enemy or potential enemy, or secure data concerning the meteorological, hydrographic, or geographic characteristics of a particular area.

The JTA specifies a minimal subset of interoperability requirements. It is becoming trite to say that the Joint Technical Architecture is necessary but not sufficient to achieve interoperability. Interoperability remains a critical issue and it certainly should be. However, as the Commander of DISA’s Joint Interoperability Test Command, Colonel Thomas Andrew, USAF, has observed, “There is a lot of interoperability out there. Many C4ISR systems do interoperate quite well together.” Colonel Andrew is in a strong position to speak with authority, since he has visibility over the DoD’s sole certifier of joint interoperability for systems.

The continued accelerated advancement of information technology ensures that fielded systems do not have the latest and greatest capabilities. The Revolution in Military Affairs is rapidly accelerating the rate at which requirements change, but the essential question should be, “Does the fielded system meet mission requirements?”

The Revolution in Military Affairs is built on software. The rapid linking of disparate weapons and command systems is done via software. Therefore a significant number of interoperability issues are software-based. Laymen commonly think of software in terms of application software. More often than

not, interoperability issues dealing with passing targeting data from a sensor platform to a weapons platform (sensor-to-shooter) involve low-level software to include firmware. Firmware is essentially software-reprogrammable chipsets.

Rapid technological advances have also fueled the Revolution in Business Affairs. Innovative solutions are sought to accelerate the fielding of new technology. Commercial off-the-shelf (COTS) software is widely touted as the silver bullet for speeding the delivery of updated software to the field. Unfortunately, there are no COTS products for purely military applications such as embedded weapons systems. Even with application software, some commercial products produce interoperability problems because they are designed to be proprietary, closed systems.

6.3 DEFINING THE PROBLEM SPACE

Given the enormous number of C4ISR systems in use in today's Armed Forces, it is critical that we understand clearly which systems are being addressed by our approach. This section establishes the scope of the problem with which this chapter will concern itself.

When systems are fielded from outside of the DoD acquisition process, interoperability responsibility for these systems are also outside the DoD acquisition commands. For example, CINC initiative funds have, in the past, been used to develop homegrown systems that are used in some theatres in place of the systems of record. Although we recognize the existence and importance of these systems, we will restrict our attention in this paper to those that have been fielded through PMs and PEOs.

Our general approach is to narrow the field to *C4ISR* systems; ensure that these systems have *interoperability* requirements; and ensure that we focus on *combat requirements*. We discuss each of these elements in greater detail, and suggest a methodology for arriving at the "right" list of systems.

What is a C4ISR system? Our approach is concerned with C4ISR systems and may not be applicable to other types of interoperability. The JTA's definition of C4ISR systems is presented in the introduction. Understanding this definition in the DoD, however, is complicated by the overlapping set of definitions found in DoDI 5000.2, Enclosure 2 [DOD 01]. It is often unclear how these definitions relate to one another. For example, a *National Security System* is "Any telecommunications or information system operated by the U.S. Government, the function, operation, or use of which: involves intelligence activities; involves cryptologic activities related to national security; involves command and control of military forces; involves equipment that is an integral part of a weapon or

weapons system; or, subject to the limitation below, is critical to the direct fulfillment of military or intelligence missions. This does not include a system that is to be used for routine administrative and business applications (including payroll, finance, logistics, and personnel management applications).” The relationship between this definition, which is from the Clinger-Cohen Act [USC 96], and the definition of C4ISR systems in the introduction is anything but obvious. Add to this the definitions of *Automated Information System (AIS)*, *Information Technology*, *Mission Critical Information Systems*, and *Mission Essential Information Systems* also found in DoDI 5000.2, and the difficulty involved becomes even greater. The situation is so confused that CJCSI 3170.01B currently requires that all AIS-related requirements documents be reviewed for possible JROC interest, simply to ensure that none of the “important” (i.e., joint) systems slip through the definitional cracks. [JCS 01]

Fundamentally, C4ISR systems move data that is critical to the conduct of military operations. Information systems that do not move this type of data have no operational requirement to be interoperable. (There may be a functional requirement for reasons of organizational efficiency, but this is distinctly different from an operational requirement.) For the purposes of this article, either the JTA’s definition of C4ISR systems, or DoD 5000.2’s definition of National Security Systems will suffice; for the remainder of this article, we will use the JTA’s terminology.

The task of identifying all C4ISR systems seems daunting. C4ISR systems are developed by a dizzying array of organizations within the DoD. Fortunately, the preparations for the Year 2000 (Y2K) crisis have resulted in extensive efforts throughout the DoD to inventory all computer and communications systems throughout the department. These databases can be leveraged to provide an initial, all-inclusive list of systems. Applying the definition of C4ISR systems provided in the JTA can narrow this list of computer and information systems to those systems that perform C4ISR functions.

The second step is to eliminate problems that are not interoperability problems. Interoperability as defined by the Institute for Electrical and Electronic Engineers (IEEE) as “The ability of two or more systems or components to exchange data and use information.” [IEEE 90]. This roadmap does not apply to deficient capability that is isolated to a particular system. Rather, it focuses on situations where the ability for multiple systems to communicate, cooperate, or co-exist is lacking. At the same time, it is important that we do not construe the potential for interoperability too narrowly – most systems will have interaction with other systems at some level. A system should be eliminated at this point only if the system meets none of the following criteria:

- It generates data that is used by another system.

- It processes or consumes data that is generated by another system.
- It relies on another system for delivery of data.
- It is software that operates on the same platform as another system.

For legacy systems, these criteria may narrow the list of systems considerably. However, in today's changing environment, this filter may quickly become less effective. Because of the movement to network-centric systems required by Joint Visions 2010 and 2020 [JV2010/JV2020], there will be increasing interaction between systems that were previously disjoint. [JCS 00] In addition, the proliferation of applications and the need for lighter, more agile components will drive more and more applications away from dedicated computing platforms and onto common ones. This will drive the need for shared service between applications that were previously disjoint.

Finally, we wish to focus on systems that are *critical to the unified commands*. Problems that do not affect the sharp end of the spear clearly deserve less attention. We are interested in an approach to interoperability that increases the combat capability of the U.S. Armed Forces. Systems that directly impact our ability to accomplish our mission should be the first on our list. We can accomplish this by considering the Joint Mission Essential Task Lists (JMETLs) of the unified commands. If a particular system does not help to accomplish a JMET for a unified commander, it can be safely eliminated from consideration.

Thus, we have recommended a systematic, three-step methodology. First, each command responsible for C4ISR procurements should identify the universe of all C4ISR systems. Next, eliminate those systems for which interoperability is not an issue. Finally, compare the list to the Joint Mission Essential Task List (JMETL) of the unified commands, using this to eliminate systems that are either not joint, or not mission essential.

6.4 CONSIDERING EXISTING METRICS

Interoperability is notoriously difficult to measure. While one might at first think that interoperability is an all-or-nothing proposition, this is an oversimplification. For example, systems that can exchange all of the required data elements might be interoperable – but if the speed of the exchange is too slow to support the operational requirements, then the so-called interoperability will not be of operational value. Sometimes interoperability is realized through labor-intensive workarounds. Are two systems interoperable in this case? It depends on a number of factors – the required frequency, the availability of personnel to operate the workarounds, and, in general, the ability of the procedure to meet the operational requirements (both in terms of effectiveness *and* suitability).

There has been at least one attempt to measure the interoperability of two systems, through an effort called the Levels of Information System Interoperability (LISI). Rather than a single measure, LISI is actually a collection of related models, a tool for use in applying these models, a set of metrics and techniques for applying the models, and an initiative (or process) aimed at using these models to address a wide set of interoperability objectives. At its core, LISI is based around classifying levels of interoperability by the “richness” of the communication that a particular system or group of systems allows. [CAWG 98] The LISI reference model was adopted as Appendix D to the C4ISR Architecture Framework Document Version 2.0 [CAWG 97], and has gained acceptance in some circles in the DoD. While a full discussion of LISI is beyond the scope of this paper, we believe that the model is, at root, too complicated for use in aggregating the status of systems at the level discussed in this paper.

A recent policy memorandum signed by the Undersecretary of Defense for Acquisition, Technology, & Logistics (USD(AT&L)); the Assistant Secretary of Defense for Command, Control, Communications, and Intelligence (ASD(C3I)); the Director of Operational Test and Evaluation (DOT&E), and the Director of the Joint Staff established an “Interoperability Watch List” and an “Interoperability Review List” to raise the level of attention applied to systems with interoperability problems. [DOD 00] The authors argue that the existence of this process makes the need for an accepted set of metrics even more critical.

6.5 A BASIS FOR MEASUREMENT

We propose a simplified model, wherein each system will be labeled with a color code based on two factors. The first factor is whether or not the system has any known interoperability problems. A problem exists if and only if some operational requirement cannot be met because of the deficiency. Some capabilities might be “desirable,” but if they are not *required* by the unified commands, then there is no interoperability problem for the purposes of our measure. This first factor, then measures whether or not the system meets *operational* requirements. The second factor is whether the system meets its interoperability requirement set. By this we mean that the system has implemented all documented interoperability functionality and has received a joint certification from the Joint Interoperability Test Command (JITC). This factor, then, focuses on *acquisition* requirements. We use this to distinguish between problems of *requirements definition* and problems of *requirements implementation*.

Using these two methods, we arrive at a four-colored system, as described in the two tables below.

Table 1
Acquisition/Requirement Stoplight Model

		Meets Acquisition Requirements	
		Yes	No
Meets Operational Requirements?	Yes	Green	Yellow
	No	Orange	Red

Table 2
Color Code Table Explanation

GREEN	The system meets its interoperability requirement set and has no known interoperability problems	Fielded system without known issues that meets all documented requirements
YELLOW	The system does not meet its interoperability requirement set, but has no known interoperability problems	Documented requirements do not reflect operational use of the system.
ORANGE	The system meets its interoperability requirement set, but has known interoperability problems	Revisit requirements and determine if requirements are adequate.
RED	The system does not meet its interoperability requirement set, and has known interoperability problems	Improvement, migration and/or action plans needs to be put in place.

Note that use of the stoplight model in table 1 involves drawing hard lines between meeting and not meeting requirements. This may not be an entirely straightforward proposition. Nevertheless, the value of drawing a line in the sand between “acceptable” and “unacceptable” cannot be overstated. Only by providing unambiguous judgments on the status of our systems can we move the debate from one of educated guessing (i.e., rumor and innuendo) to one with a degree of rigor and reproducibility.

There is one important point that needs to be made about this grading system. The operational requirements considered should be those that are in force at a particular time. If we are grading a system today, it should be graded with respect to today’s operational requirements – if we are projecting a grade for one year

from now, it should be based on our best information about next year's operational requirements.

For the acquisition requirements, however, we should use those acquisition requirements that were to have been implemented by a particular point in time. For example, if we have a radio system with a scheduled upgrade, and the requirements process was geared for release of the upgrade in one year, we should not include the upgraded requirements in an evaluation of the system's readiness today. We *should*, however, include the upgraded requirements in an evaluation of the system's readiness next year – and we should do so using our best information about whether/when the system will meet these requirements. This is important (and fair) because we must recognize that the acquisition community is constrained by technical and fiscal realities, and does not always have the ability to deliver improvements in time. Its schedule is dictated by the ever-present trades with cost and performance, and may even be influenced by the timeliness of the identification of the requirement. This part of the measurement is designed to measure the acquisition community's ability to respond to the requirements process – and schedule is a key part of that process.

Of course, in both cases, the farther in the future we project these requirements, the more cautiously we must consider their results. Nevertheless, short- and medium-term forecasts are likely to have tremendous planning value. Consider the following notional example as shown in Figure 1.

System XYZ was fielded one year ago. Its original release had some problems with one of its interfaces, and as a result, was unable to pass JITC testing. Nevertheless, System XYZ was fielded because it provided significant capability to the Unified Commands. Unfortunately, its failures have limited its operational employment. Therefore, its current interoperability readiness status is red – it has both operational and acquisition problems.

A software fix, just submitted for JITC testing, has undergone favorable initial review and is expected to resolve both the operational and acquisition problems. JITC testing will be completed in three months, so at that time we project that its interoperability readiness status will be orange. Three months later, the software fix will be incorporated in all fielded units, so in six months, we project the interoperability readiness status will be green.

Meanwhile, however, system XYZ was fielded with a Preplanned Product Improvement (P3I) strategy, since the operational requirements were expected to increase dramatically over the course of the next few years. The upgraded system is due in two years, but since the current system implemented a few features early, it will actually continue to function effectively for the next thirty months. As a result of a funding cut, the PM is already warning that the upgrade will be three

months late. Therefore, in 24 months we expect the system to enter the yellow state, where it is behind the acquisition cycle, but still meets operational requirements. In 27 months when the next version is released, it will become green, and will stay green at 30 months because the system will still meet requirements.

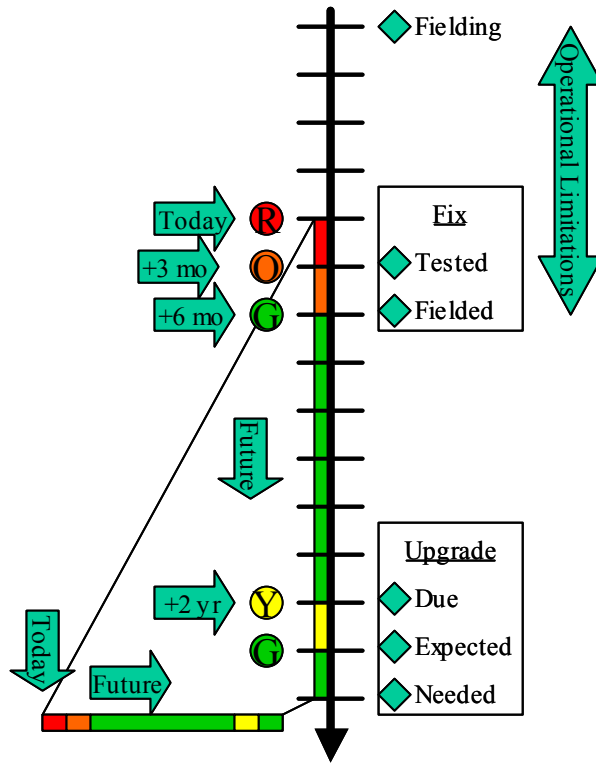


Figure 6-1 *Notional System Timeline*

6.6 APPLYING THE BASIS

The section above describes how each system can be graded using a modified stop-light style scheme. Systems graded green and yellow meet their operational requirements, although yellow systems warn of possibly over-specified requirements or potential future problems. Systems graded orange and red limit operations. Red system indicates a problem in requirements implementation,

while orange systems indicate a (potentially harder to solve) problem in requirements definition.

The problem with this scheme is that it may not provide adequate insight into the etiology (root cause) of the problem. Systems with multiple interfaces may be coded red because of problems with only one interface. Instead of grading systems, it may be necessary to unroll this rating and grade the interfaces between systems. Each pair of interacting systems could be given a color code based on the scheme described above.

The advantage to grading interfaces is a more fine-grained understanding of not only the systemic cause of the problem, but also of its operational impact. A red system may perform perfectly well in a large number of operational scenarios, if the most commonly required interfaces are not the causes of the problem. Conversely, it may have little operational value, if the most useful interfaces are the ones with the problem.

The disadvantage to grading interfaces is the dramatic increase in the magnitude of the problem space. With the increased interconnections between systems in the network-centric environment predicted by JV2010, we can expect increasing numbers of interactions between system pairs. Whereas the difficulty of measuring readiness for systems increases with the number of systems, the difficulty of measuring readiness for interfaces increases with the square of the number of systems.

One way to mitigate this problem would be to narrow the number of interfaces considered. Instead of grading all interacting system pairs, it may be useful to begin by evaluating the readiness of each system in the problem space. For red and orange systems, one might next evaluate each interaction with another system using the readiness-reporting model, and then develop remediation plans on a per-interface basis. Once this process was established, yellow systems could be similarly analyzed to determine whether the requirements were over-specified or the systems under-utilized, or whether (as in our earlier example) the yellow status was a transient situation that was not cause for great alarm.

6.7 AGGREGATE MEASURES

Clearly, these readiness measures, by themselves, can provide a useful tool in helping to address the DoD's interoperability efforts. They focus attention first on those systems that do not meet operational requirements, putting emphasis on meeting the warrior's needs. In addition, they help to identify problems in the requirements definition process. By separating problems of implementation from problems of definition, the authors believe the process will highlight the successes

of the acquisition community. More importantly, by identifying problems of definition, they will help to focus efforts in this area, hopefully helping to prevent these problems from recurring in future systems. Given the sometimes-long loop between requirements definition and operational employment, it is key that these lessons are fed back into the requirements generation system as quickly as possible so that they are not compounded.

However, it is also important to have an overall measure of the overall health of C4ISR interoperability in the DoD. The readiness reporting measures can be used to provide such an aggregate measure.

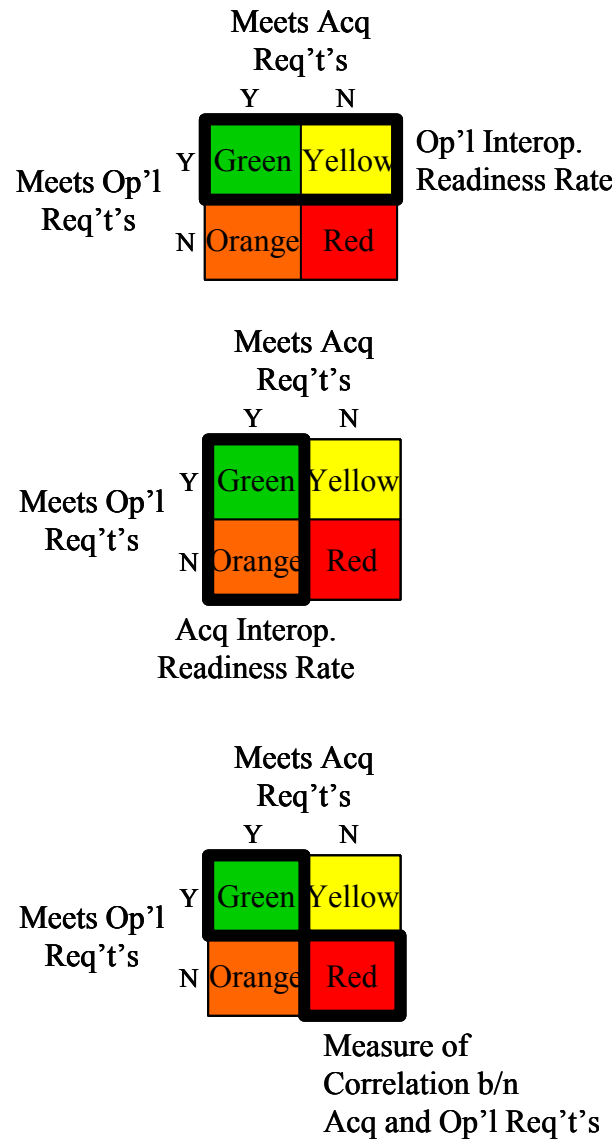
The simplest way of doing this is to measure the percentage of systems that are classified in each color. Of particular interest would be the percentage of green and yellow systems (the system operational interoperability readiness rate), and the percentage of green and orange systems (the system acquisition interoperability success rate). It might also be interesting to consider the percentage of green and red systems, which indicates the correlation between acquisition success and operational readiness – perhaps providing an indicator of the health of the acquisition process (including the requirements process) as a whole. See Figure 2 for a visual depiction of these rates.

This simple method of aggregating the scores is unambiguous. But does it provide a fair ranking of the interoperability readiness of our C4ISR systems? Because every system has equal weight, a small system filling a relatively small niche in the combat environment (for example, a medical field supply reporting system) is rated equally with a large system that interacts with many other battlefield systems (for example, an AWACS aircraft). This is probably not appropriate.

For that reason, it may be desirable to assign weights to each of the systems in our problem domain. This could potentially be done using any number of methods, taking into account factors including, but not limited to, number of interfaces, number of deployed units, or criticality of the system as rated by operational commanders.

The percentages described above could then be calculated with each system contributing a weighted score. (We refer to these measures as *weighted* system operational interoperability readiness rates, etc.) This might more accurately reflect the state of the department. On the other hand, it might tend to mask problems with relatively small, yet important systems, if the weights are not carefully constructed.

An alternative to using the system readiness indicators would be to use the interface indicators. Again, these could each be given equal weight, or could be

Figure 6-2 *Aggregate Measures*

weighted based on the importance of the particular interaction between each system-pair. We refer to these as interface interoperability readiness rates and weighted interface interoperability readiness rates, respectively.

Clearly, in order to gain value from these aggregate metrics, they must be maintained. In order for them to be maintained, the readiness scores for each system (or interface) must be maintained. It is also important to realize that as new systems are fielded, they become legacy systems and must be added to the scoring system.

Perhaps the best way to roll-up these figures would be to use them to score the ability of U.S. forces to meet various mission scenarios. By developing notional architectures for operational employment of our forces, the interoperability readiness reporting measures could be used in the same way existing readiness measures are used. In this case, typically, the lowest score would dominate and be carried up to a higher level. This would provide valuable insight into the ability of C4ISR systems to meet national strategic needs, and would help to justify funds for improvements where they are required.

6.8 DISCUSSION: SOLVING THE PROBLEM

With this collection of metrics, we can now propose a roadmap for a department wide approach to addressing the C4ISR interoperability problem. First, identify an agency with responsibility for overseeing the roadmap. Second, under the direction of the lead agency, develop an authoritative list of systems that fall into the problem space identified in this paper, together with an institutionalized process for maintaining that list. Third, develop readiness scores (and, possibly, weights) for each system on the list, together with an institutionalized process for maintaining this data. Fourth, using the first consolidated, across-the-board data set, measure the current state of our legacy C4ISR interoperability problem. Finally, set realistic goals for improving this state, allocating the resources required to realize those goals, and measuring progress along the way. This section will discuss this approach in more detail.

The issues surrounding C4ISR interoperability affect the services, defense agencies, and the unified acquisition community. In addition, success in this effort ultimately requires the ability to prioritize problems across the military community and to allocate resources in support of the priorities. The authors believe that this function is fundamentally related to Joint Forces Command's UCP-99 responsibilities as joint forces trainer, provider, and integrator, and that application of this metric is consistent with and supportive of their other objectives.

As its first task, the lead office would need to compile a master list of *mission essential*, C4ISR systems with interoperability requirements. An approach for doing this is outlined in the first section of the paper. Just as important as

compiling the initial list, the organization must put in place a process for maintaining the list, involving all of the various entities throughout the DoD that field C4ISR systems. This encompasses a large array of organizations, including CECOM, SPAWAR, ESC, NAVAIR, NAVSEA, PEO's within the AF and Army, DISA, DIA, NSA, SOCOM, and others. This process would most likely involve integration of existing processes, rather than development of new ones; nevertheless, given the number of developing organizations, this represents a potentially non-trivial effort. Even by itself, however, this task will likely result in a valuable resource by identifying all systems in this class, together with the responsible agency.

With the master-list in hand, this office could then begin to assess the readiness of each of these systems. This rating need not be done by a single office, but could potentially be delegated to subordinate agencies. The key to this evaluation, in fact, is to ensure that the operational community evaluates the operational requirements, while the acquisition community evaluates the acquisition requirements. The office, of course, would be responsible for providing sufficient oversight to ensure honest responses were provided. In addition to understanding "today's" situation, data should be collected on the projected status over the medium term, perhaps 3-5 years. As described in the example earlier, changes in status can sometimes be predicted in advance, and knowing when these changes will take place is important data for decision makers.

Once again, determining the status of these systems is not enough. The office must put in place procedures for maintaining this information. The challenge here comes from the large number of players. Although much of this information exists within the department, making it accessible in normalized form in a centralized location would represent a key contribution to dealing with the interoperability problem.

Once the list of systems and the data on each system is available, and an institutional process for maintaining it is in place, the office's work becomes at once simpler and more significant. In addition to keeping the process running (which is likely to remain non-trivial in light of the propensity for reorganization within the federal government), the office must compile aggregate statistics (including projections of the aggregate statistics). More importantly, the office must drill down into the interfaces of problem systems, determining the true sources of problems, and engage the necessary players to effect solutions. This could be done both to solve present problems, and to prevent or mitigate problems that may be anticipated.

Of course, in addition to the department wide statistics, the data could be cut along different lines. One could consider the interoperability readiness of all sensors, all shooters, all systems within a particular Area Of Interest, or all systems produced

by a particular agency. Using interface data, one could even construct a picture of the “interoperability” between two agencies, by considering the status of interfaces for systems produced by the two agencies. This might identify instances where institutional obstacles played a role in the observed interoperability problems. Equally important, it would likely highlight the tremendous levels of cooperation between many of the agencies acquiring C4ISR systems.

Armed with hard data, “ground truth” can be provided to the interoperability debate. Realistic goals could be set, and the resources to achieve these goals could be allocated to the organizations in the best positions to do so. Most importantly, progress toward these goals could be objectively measured. The process would separate requirements issues from acquisition issues, offering opportunities for improving both of these systems in cases where systemic factors are found to have contributed to problems. At the same time, it would ideally protect all parties from recrimination by focusing the entire community on solving the problem – delivering C4ISR systems that meet the interoperability demands of our warriors.

6.9 CONCLUSION

One can easily argue that America’s unrivalled dominance on the battlefields of the late 20th century is due largely to the success of our acquisition system, even in light of the declining defense budgets of the last decade.

Our ability to attain full-spectrum dominance in the 21st century will rely heavily on our C4ISR infrastructure. [JCS 00] Interoperability of our C4ISR systems is essential to achieving this goal. We are doing well in this area, but we can do better.

That which is measured improves. We will never fully eliminate interoperability problems, because evolving operational requirements will continue to challenge the developers of C4ISR systems. By its very nature, our enterprise will always be pushing the limits of technology, generating new problems even as the old ones are solved. It is important to show progress, analyze the drivers behind our interoperability problems, and apply maximum effort at the point of greatest leverage to solve them. Only in this way can we provide the greatest possible utility to the soldier on the battlefield.

We have developed a simple, readiness-reporting style method of measuring interoperability. We have also proposed a method for aggregating the data in a manner that will facilitate tracking progress on a DoD-wide basis. Finally, we have outlined a mechanism for applying the metrics within the DoD in order to

facilitate solution of the problem. It is our sincere hope that this roadmap will generate open discussion within the department that will ultimately lead to a more rigorous approach to interoperability.

The views expressed in this paper are the opinions of the authors, and do not reflect the official opinions of any U.S. government agency.

6.10 ENDNOTES

[JTA 99] Joint Technical Architecture Version 3.0 Draft 1, Defense Information Systems Agency, Arlington, Virginia, 26 February 1999, p C4ISR-1.

[DOD 01] Department of Defense Instruction 5000.2, "Operation of the Defense Acquisition System," Change 1 (4 January 2001) to 23 October 2000.

[USC 96] Section 1401 et seq. of title 40, United States Code, "Clinger-Cohen Act of 1996"

[JCS 01] Chairman of the Joint Chiefs of Staff Instruction 3170.01B, "Requirements Generation System," 15 April 2001.

[IEEE 90] IEEE STD 610.12-1990, IEEE Standard Glossary of Software Engineering Terminology

[JCS 00] Chairman of the Joint Chiefs of Staff "Joint Vision 2020." US Government Printing Office, Washington, DC, June 2000.

[CAWG 98] C4ISR Architecture Working Group, "Levels of Information System Interoperability (LISI)," 30 March 1998.

[CAWG 97] C4ISR Architecture Working Group, "C4ISR Architecture Framework Version 2.0," 18 December 1997.

[DOD 00] Memorandum entitled "Promulgation of DoD Policy for Assessment, Test, and Evaluation of Information Technology Systems Interoperability, 4 December 2000," jointly signed by USD(AT&L), ASD(C3I), DOT&E, and the Director, Joint Staff.