To: Architecture Methodology Working Group
Subject: DOD TAFIM Volume 6: Draft DOD Goal Security Architecture, Version 1.0

The Draft DOD Goal Security Architecture (DGSA), dated 1 August 1993, is hereby submitted to the Architecture Methodology Working Group as volume 6 of the DOD Technical Architecture Framework for Information Management (TAFIM). Comments on the DGSA should be submitted according to the same guidelines as the other volumes of the TAFIM.

The DGSA authors have been working with the other TAFIM authors to integrate security within the current (29 October 1993) releases of volumes 2 and 3, but additional effort is needed to complete this integration. One aspect of this integration that will be considered in preparing the next version of volume 6 is the revision, or even deletion, of section 3 and appropriate changes to other parts of volume 6. One consequence of the current incomplete integration is that volume 6 references the previous versions of volumes 2 and 3.

The DGSA was circulated for review within the Defense Information Systems Agency (DISA) and the National Security Agency (NSA) prior to its present submission to the Architecture Methodology Working Group. Less than three weeks were available to consider the comments received from the DISA and NSA reviewers. However, several issues raised, questions asked, and errors discovered were deemed important enough that they should be presented to current recipients of the DGSA. Simple errors are dealt with in the attached *Errata*. Technical questions for which succinct answers could be given and which would enhance current recipients' reading of the DGSA have been recorded in the attached *Questions and Answers*. Several significant technical issues were raised. A list of these issues is attached. In those cases where brief responses were feasible, they are given with the issue statements. Others, which would require significant explanation or for which an immediate response is not available, are simply listed. In most cases, the issue statements are a composite of multiple comments received. The open issues will be considered for the next version of the DGSA. In addition, many excellent editorial, presentation, and stylistic suggestions were received which will be reflected in the next version of the DGSA.

Several questions were raised about the relationship of the DGSA to other efforts within the Center for Information Systems Security (CISS) and how CISS plans to cause the DGSA principles to be adopted by government and embraced by industry. First, it is planned that a future DGSA Executive Summary will enable program managers to understand how they can use and can respond to the DGSA, what the benefits of the DGSA are, and how they can get help in applying the DGSA. Second, there is already underway within CISS the creation of the DGSA Overall Transition Strategy (DOTS). DOTS will create the means for DOD information system planners and managers to incorporate DGSA principles into their specific information system architectures. Transition to the DGSA (via DOTS) is defined as the incorporation of DGSA security concepts into current and new DOD information system architectures. The DOTS is organized around several areas, called segments, that are critical to creating tools, products, and support for transition to the DGSA. The segments are: standards, product development, research and technology, security management, local subscriber environments, communications systems, certification and accreditation, policy and doctrine, and education and training. DOTS is a vehicle not only for program transition, but will involve the commercial community to provide off-the-shelf products that will allow specific information systems to achieve the DGSA vision. The DOD community will be invited to join the CISS on-going DOTS effort. For further information on DOTS, please address inquiries to Carl Deutsch via e-mail at deutsch@dockmaster.ncsc.mil or at Defense Information Systems Agency, Center for Information System Security (TGF), Suite 400, 5113 Leesburg Pike, Falls Church, VA 22041-3230.

ERRATA
29 October 1993

1. 2.3.2.1, first sentence: replace "Center for" with "Corporate".

2. Table 2-1, first entry under Multiple Security Policy Support: replace "policy" with "policies".

3. 4.1.3, title, first sentence, and figure 4-5: replace "IM Integration Model" or "Information Management Integration Model" with "TRM".

4. Figure 4-5: replace "Communications" with "External Environment".

5. 4.3.7, paragraph 4, third sentence: replace "*processes*" with "*process*".

6. 5.1.2, paragraph 2, third sentence (last word): delete "security".

7. 5.2.4, title: replace "Relevant" with "Related".

8. 8.0, third sentence: delete ", which is part of the security policy".
   8.0, fourth sentence: replace "policy" with "doctrine".

9. Page ACR-1, CIM: replace "Center for" with "Corporate".

10. Page ACR-2, JCALS: replace entry with "Joint Continuous and Life-Cycle Support System".

11. Page ACR-3, PABX: entry should be "Private Automated Branch Exchange".

QUESTIONS AND ANSWERS
29 October 1993

I. Local Subscriber Environment (LSE)

1. Are there levels of LSEs (i.e., can an LSE be made up of several other LSEs)?

   ANSWER: No, there is no notion of levels of LSEs.

2. How do we decide what is an LSE vs. several LSEs? If multiple tenants share an LCS, is it a single LSE or multiple LSEs? Is the difference between an LCS and a CN merely a matter of who owns it? Why is part of the RS outside the dashed line in figure 4-2 representing the transfer system?

   ANSWER: LSE identification is determined by examining the policy authority that controls of the environment and the resources in it. If multiple tenant organizations share a building, they may or may not have independently controlled resources within that building. Each group of independently controlled resources forms an LSE.

   In figure 4-2, a part of the RS is indicated as being outside of the transfer system because there may be functions within the RS that are not associated with the transfer of information.

II. Information Domains

1. Can a user operate in more than one information domain at once? What exactly is meant by these restrictions? [Question refers to 4.3.3, paragraph 5]

   ANSWER: A user can be a member of more than one information domain simultaneously. If an end system(ES) that a user is currently working on supports two or more of those information domains of which the user is a member, then the user can have one or more security contexts established (representing different activities) for those information domains at the same time. The burden is on the ES to maintain separation of the user’s activities and other ES functions.

Regarding the second part of the question, the phrase "these restrictions" refers to those stated at the beginning of the identified paragraph.

2. Considering the fact that users will operate in more than one information domain, is there a concept of a super-user for the system such that a workstation "owner" will have access to all system files?

   ANSWER: While this is not precluded, it is not recommended. The notion of a super-user in the DGSA context would require the user to have all priviliges for all information domains implemented on the workstation, including security management.

3. Will membership in an information domain group be explicit, implicit or some combination based on the sensitivity of the information in a particular domain?

   ANSWER: Membership must in all cases be explicit in the security policy as implemented in the SMIB, except for the "public domains" (those collections of information objects to which anyone may have access).

4. When an information object is copied to another information domain does it need to be updated when the original is updated?

   ANSWER: The DGSA does not require that such transfers be automatically updated. On the other hand some applications may require this and a suitable implementation should be designed. That is, if there is a requirement to update copies of an information object, the implementation will be no more difficult (and no less difficult) than updating copies in any distributed system.

III. Security Management

1. Is a MIB a managed object?

   ANSWER: A MIB (and by implication a SMIB) is a logical construct which in and of itself is not a managed object. Its component parts are managed objects.

IV. Metrics for Security Mechanisms and Values for Information

1. The scale [reference is to section 9.7.2, paragraph 10] is allegedly hypothetical, what about the scheme for combining the individual results?

   ANSWER: It is hypothetical also and we continue to investigate how to assign values to mechanisms. This investigation will include two notions of combining, first, where more than one security mechanism is required to implement a given security service, and, second, where a collection of security mechanisms supports different security services.

V. Absolute Protection

1. Apparently, looking at the example [reference is to section 9.7.2, paragraph 12], it’s OK to have two LSEs, with very different "ratings" protecting the same information. What is the deal?

   ANSWER: Protection of information within an information domain may be provided by different security mechanisms from LSE to LSE. Therefore, a security mechanism in one LSE may have a different "rating" than another security mechanism in another LSE in which the same information is being protected. What absolute protection states is that the collection of security mechanisms used to protect information must provide at least the minimum protection required for that information domain in any of the end systems that support a particular information domain and its associated information

VI. Miscellaneous

1. Would the requirement for strict isolation exclude systems that utilize parallel processing capabilities?

   ANSWER: No, it would not. However, one would have to examine the features of the specific parallel processor architecture to determine if it supports the DGSA concepts. This question may result in an example in section 9 in a future version.

2. Why should the external interfaces be consistent with existing

standards? This may introduce additional vulnerabilities especially since some of the existing standards do not address security issues.

ANSWER: While it is true that many of today's standards do not address security, two situations arise. First, some of these standards will not be relied upon to provide any security service(what is required is correct implementations) and are therefore very appropriate for the future. Second, those that must address security for the *goal*architecture represent research and development activities that will be identified. As systems transition to the DGSA, decisions about existing capabilities and the tradeoffs associated with their use will be judged.

3. Has a distribution plan/process been developed to handle subsequent iterations of the DGSA?

   ANSWER: Yes, since the DGSA will [and with this release has] become a volume of the TAFIM, we will follow the TAFIM distribution procedures in future releases.

4. In the referenced paragraph should the phrase "open systems" be changed to read "open systems environment"?

   ANSWER: "Open systems" does not refer to the standards-based project known as "Open Systems Environment". A future release will define more carefully what is meant by open systems in the DGSA (note the use of lower case). Roughly, what is intended is systems that are potentially open to interoperation with other open systems that adhere to a common set of communications protocols, and are flexible in support of a range of information domain security policies. Generally, this will include OSE and OSI.

VII. Security Context

1. Does the term security context refer to profiles (e.g., Federal Criteria product and system profiles)?

   ANSWER: No it does not. As described, a security context is a concept for protecting the operations of a user in an end system according to a particular information domain security policy. The Federal Criteria product and system profiles could be considered the architectural and implementation support aspects for security context assurance factors.

VIII. Registry (Cryptographic Algorithms, etc.)

1. ... there must be a registry of cryptographic algorithms and key management schemes so that specific choices can be negotiated for a particular security association. Who will develop this strategy and where will it reside?

   ANSWER: A definitive answer to this question has not been proposed at this time. This to be an infrastructure issue which is outside the scope of the DGSA. This is one of the many issues to be addressed by the DGSA Overall Transition Strategy (DOTS).

IX. Training

1. Will there be a DGSA training package available for potential DGSA users?

   ANSWER: Yes, the approach and format are being addressed in the Training Segment (working group) of the DOTS.

X. Multiple Security Policies

1. Can the security policies that govern allowable interrelationships between LSEs be made dynamic based upon external conditions?

   ANSWER: Yes, security policies can contain contingency plans that are invoked as the result of some external, authenticated event. As well, replacement of security policies, if coordinated and approved through the proper authorities, is a possibility. This latter approach obviously has an impact on the degree of speed with which this can be accomplished.

2. Does the End System have a preset policy? Does it know about all policies it supports? What about adding policies in the future?

ANSWER: The only policy that an end system must enforce is strict isolation. Beyond this, the policy for an end system is the accumulation of all policies it supports. It will be normal for policies to be added or deleted from time to time and an end system must support this requirement.

XI. Multidomain Objects

1. If creation of new information domains is relatively simple, why is it useful to create multidomain objects?

   ANSWER: The principle purpose is to create a displayed or printed instance of (parts of) information objects that are related to one another, but which belong to different information domains. If a single security policy could be constructed for the combination of related information objects (and there were enough such combined objects to make it worth while), a new information domain could be created for them. However, there may be requirements to maintain the component objects in their original information domains. One reason for such a requirement might be because it is necessary to always be able to determine the original information domain (or equivalently, the security policy). Once made part of a composite information object, it may be extremely difficult or even impossible to make such a determination.

2. How are multidomain objects marked?

   ANSWER: The components of multidomain objects are marked consistent with the security policy of the information domain to which it belongs. Either the component security policies must speak (consistently) to how a the composite displayed or printed image is to be marked, or a composite policy (known to the end system) must be applied (for example, page markings that would result from U.S. National classified information policy).

XII. Strict Isolation

1. It is not clear why there is a statement that hardware "indirectly" supports the (strict) isolation. Why isn’t it direct?

ANSWER The hardware could conceivably completely support strict isolation if an appropriate hardware architecture was available. However, today’s general purpose computers do not have such a hardware architecture, nor is it clear that such hardware support will be available in the future. Therefore, it is stated that the hardware indirectly supports this requirement by means of protecting the software which manages the strict isolation.

DGSA ISSUES LIST
29 October 1993

This issues list contains statements summarized from multiple reviewer comments. "TBS" in a response indicates either a currently open issue or one for which a short answer was not feasible at the time the list was prepared.

1. Some reviewers expressed a concern that section 2.1 is based on a security policy document not generally available. Additionally, concerns about traceability to other existing national, service or agency policies were stated.

   The "DOD Information Systems Security Policy" (NSA, 1993) [DISSP-SP.1], will be made available through CISS. This policy is a consolidation of the security objectives of DOD information systems users which takes into account national and DoD security policies. The consolidation was accomplished by the DISSP.

2. Some reviewers objected to the apparent lack of security features such as (end) system integrity, (end) system availability, and software integrity.

   Response: Section 4.2 states that "Since the DGSA applies to all aspects of information security, the basic services [of IS 7498-2 plus availability] are considered to apply not only to the transfer system, but are interpreted to apply to the entire LSE." Perhaps this statement is too cryptic, but its intent is to include for access control, for example, not only access control within communications protocols, but also facility access control (a doctrinal security mechanism - see section 8), and end system access control (to the end system itself and to information within the end system). Similar extended interpretations of the remaining security services encompass the reviewers' examples.

3. Some reviewers believe that the "LSE protects the hardware, the hardware protects the software, the software protects information" paradigm of section 5.1 is too simplistic, narrow or rigid and ignores modern information security models.

   Response: The intent of the cited paradigm is to point out relationships among the environment, hardware components and systems, and software in *jointly* protecting information. For example, access control is a combined responsibility as indicated in the response to issue 2.

4. Some reviewers asked for additional justification for the decision to allow only non-hierarchical relationships among information domains. Other reviewers stated opposition to this decision.

   Response: The wording in section 4.3.1, "Information domains are not hierarchically related, ...", should say that "Information domains can be related or unrelated". They can be hierarchically or non-hierarchically by their security policies. The sets of information objects which form information domains are not related, in a security sense, except through their security policies.

5. Some reviewers questioned the claimed ability of the end system architecture to reduce covert channel concerns (section 5.2.2, last paragraph).

   Response: TBS

6. Some reviewers believe the *name* "absolute protection" (not the concept for which it is a label) to be wrong or is an obstruction to understanding the concept; several alternatives were suggested.

   We have long been seeking a satisfactory substitute for the term "absolute".

7. Some reviewers expressed doubts that the security policy decision function (SPDF) (sections 5.2, 5.2.3.1) could (or should) be made to deal with all conceivable security policies, or all aspects of security policies.

   Response: Since, as is indicated in section 5, the entire matter of SPDF implementations, and the representation of the security policies which can be interpreted by the SPDF, is still essentially a research area, these doubts cannot be answered at this time. If research efforts fail to produce completely general SPDF solutions, one fall back position might be several more specialized SPDFs, although some of the advantages of a single SPDF discussed in section 5 would be diluted.

8. Some reviewers stated that the security policy enforcement function (SPEF) (section 5.2.1) can only enforce access control policies because it is a part of the separation kernel and is subject to classical reference validation monitor (RVM) basic properties.

Response: Notwithstanding statements in section 5.2.1 that the SPEF is an *extension* of the access control enforcement function and that the separation kernel is an *extension* of the RVM, the reviewers either ignored those statements or, perhaps, did not believe them to be possible. Like the SPDF (issue 7), there are some aspects of the SPEF for which research and development must be undertaken. It should be kept in mind that not all elements of this *goal* security architecture are expected to be achieved in the short term.

9. Several related questions about the four architectural types (section 1.3) and their application to the DGSA were raised, including: to which architectural types does the DGSA correspond, and whether some of the more detailed provisions of the DGSA are appropriate to a non-specific architecture.

   Response: TBS

10. Some reviewers noted that the term "local subscriber environment" (LSE) appeared to be used in two ways, sometimes referring to a "collection of end systems, relay systems, and local communications systems", and other times referring to the "environment" in which such a collection exists.

    Response: The observation is correct. The description of LSEs in sections 4.1.1 and 4.1.2 should be explicitly made to include the environment in which the components exist. With this in mind, the intended uses of the term LSE do not appear to be in doubt when considered in the contexts where it occurs.

11. Some reviewers questioned the divisions in section 5 among security-critical, security-related, and non-security-related functions and their relationships to trusted and untrusted software and hardware.

    Response: TBS

12. Various perceived difficulties concerning information domains were noted, such as their potential large number, how they would be managed, and their apparent inability to accommodate the equivalent of discretionary access control mechanisms.

    Response: TBS

13. An apparent inconsistency was noted between section 3.3.2, where it

is stated that "... bit integrity will be inherent in the communications networks [CNs] used ..." and section 4.2.1 which allocates to CNs only the availability security service.

Response: There are many security services and mechanisms that a CN provider might employ to ensure the availability of the contracted communications service, including physical protection, configuration management and integrity of switching software, and authentication and access control for management protocols. These security services and mechanisms are *not* applied to the information being transferred between end systems, but are employed solely to guarantee a specified level of communications service availability.