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.

Department of Defense (DoD) Goal Security Architecture (DGSA)

Center for Information System Security
Defense Information Systems Security Program

Version 1.0

1 August 1993

PREFACE

The Defense Information Systems Security Program is a joint undertaking of the Defense Information Systems Agency (DISA) and the National Security Agency (NSA). This document was prepared by the Architecture and Engineering Directorate of the DISA Center for Information System Security. The Architecture and Engineering Directorate is located at the NSA and comments on this document may be delivered to:

DIRECTOR, NATIONAL SECURITY AGENCY
Fort George G. Meade, MD 20755-6000
Attention: CISS/A&E/V37
Facsimile: 410-859-6813
Internet: rmcallister@dockmaster.ncsc.mil

TABLE OF CONTENTS

SECTION	PAGE
1.0 INTRODUCTION	1-1
1.1 PURPOSE	1-1
1.2 SCOPE	1-2
1.3 ARCHITECTURAL TYPES	1-2
1.3.1 Abstract Architecture	1-2
1.3.2 Generic Architecture	1-3
1.3.3 Logical Architecture	1-3
1.3.4 Specific Architecture	1-3
1.4 PROCESS	1-3
1.5 DOCUMENT ORGANIZATION	1-4
2.0 SECURITY REQUIREMENTS	2-1
2.1 DGSA SECURITY POLICY	2-3
2.2 DGSA SECURITY REQUIREMENTS	2-4
2.2.1 Multiple Information Security Policy Support	2-4
2.2.2 Open Systems Employment	2-4
2.2.3 Appropriate Security Protection	2-5
2.2.4 Common Security Management	2-5
2.3 DGSA DERIVED SECURITY REQUIREMENTS	2-6
2.3.1 Security Requirements Refinement	2-6
2.3.2 Interaction Between Mission-Related Operational Objectives and Security Requirements	2-8
3.0 DIS TARGET ARCHITECTURE PERSPECTIVES AND EVOLUTION	3-1
3.1 TAFIM GUIDANCE	3-3
3.1.1 DoD Technical Reference Model	3-3
3.1.2 TAFIM Information Management Integration Model	3-5
3.2 C4IFTW GUIDANCE	3-7
3.3 DIS TARGET ARCHITECTURE EVOLUTION	3-8
3.3.1 Voice-Video/Imagery-Data	3-8
3.3.2 Portable-Mobile-Fixed	3-9
3.3.3 Secure-Nonsecure	3-10
3.3.4 Workstations-Servers-Hosts	3-10
3.3.5 Programmable-Nonprogrammable	3-10
4.0 SECURITY VIEWS AND CONCEPTS	4-1

4.1 INFORMATION SYSTEM ARCHITECTURE SECURITY VIEWS 4-1
4.1.1 Abstract Information System Architecture Security View 4-1
4.1.2 LSE Generic Security View 4-2
4.1.3 TAFIM IM Integration Model Security View 4-5
4.2 SECURITY SERVICE ALLOCATIONS 4-6
4.2.1 CN Security Service Allocation 4-6
4.2.2 LSE Security Service Allocation 4-7
4.3 SECURITY CONCEPTS 4-10
4.3.1 Information Domains 4-10
4.3.2 Strict Isolation 4-12
4.3.3 Interdomain Information Sharing and Transfer 4-12
4.3.5 Absolute Protection 4-15
4.3.6 Uniform Accreditation 4-16
4.3.7 Security Management 4-17

5.0 END SYSTEMS AND RELAY SYSTEMS 5-1
5.1 END SYSTEM SECURITY ARCHITECTURE OVERVIEW 5-2
5.1.1 The LSE Protects the Hardware 5-2
5.1.2 The Hardware Protects the Software 5-2
5.1.3 The Software Protects Information 5-3
5.2 END SYSTEM SECURITY ARCHITECTURE DESCRIPTION 5-3
5.2.1 Separation Kernel 5-6
5.2.2 Security Contexts 5-8
5.2.3 Security-Critical Functions 5-11
5.2.4 Security-Relevant Functions 5-15
5.3 END SYSTEM SECURITY ARCHITECTURE TECHNOLOGIES 5-17
5.3.1 LSE 5-18
5.3.2 Hardware 5-18
5.3.3 Software 5-18

6.0 SECURITY MANAGEMENT 6-1
6.1 SECURITY MANAGEMENT RELATIONSHIPS TO DGSA CONCEPTS 6-1
6.2 ISO 7498-2 AND DGSA SECURITY MANAGEMENT CONCEPTS 6-4
6.2.1 Information Domains 6-4
6.2.2 Security Management Information Bases 6-4
6.2.3 Communication of Security Management Information 6-6
6.2.4 Distributed Security Management Administration 6-7
6.2.5 Security Management Application Protocols 6-7
6.2.6 End Security Management Functions 6-8
6.2.7 Security Service Management 6-9
6.2.8 Security Mechanism Management 6-13

6.3 SECURITY MANAGEMENT TOOLS 6-18
6.3.1 Security Policy Rule Specification 6-18
6.3.2 Security Mechanisms Catalog 6-18
6.3.3 Maintenance Applications for Security Administrators 6-19
6.4 AREAS FOR SECURITY MANAGEMENT STANDARDIZATION 6-20

7.0 TRANSFER SYSTEM 7-1
7.1 DISTRIBUTED SECURITY CONTEXTS 7-2
7.1.1 Distributed Security Contexts for Interactive Communications 7-2
7.1.2 Staged Delivery Distributed Security Contexts 7-6
7.1.3 Other Aspects of Distributed Security Contexts 7-7
7.2 TRANSFER SYSTEM SUPPORT 7-8
7.2.1 Security Management Application Process 7-8
7.2.2 Security Management Information Base 7-9
7.2.3 Security Protocols 7-10
7.2.4 Cryptographic Support 7-11
7.2.5 Distributed Management Systems 7-12
7.3 DGSA TRANSFER SYSTEM ISSUES 7-13
7.3.1 Traffic Flow Security in Open System Communication Environments 7-14
7.3.2 Limitations on Distributed Processing 7-14

8.0 SECURITY DOCTRINE 8-1
8.1 SECURITY SERVICES 8-1
8.2 DOCTRINAL SECURITY MECHANISMS TO PROVIDE SECURITY SERVICES 8-2
8.2.1 Mechanisms for Identification and Authentication 8-2
8.2.2 Mechanisms for Access Control 8-2
8.2.3 Mechanisms for Confidentiality 8-3
8.2.4 Mechanisms for Integrity 8-4
8.2.5 Mechanisms for Non-Repudiation 8-4
8.2.6 Mechanisms for Availability 8-4
8.3 COTS PRODUCT CONSIDERATIONS 8-5
8.4 SECURITY MANAGEMENT 8-5

9.0 EXAMPLE LOGICAL ARCHITECTURE DEVELOPMENT PROCESS 9-1
9.1 MISSION DESCRIPTIONS 9-2
9.1.1 Intelligence Mission Area Function Description 9-2
9.1.2 Tactical Mission Area Function Description 9-5
9.1.3 Drug Interdiction Mission Area Function Description 9-6
9.2 MISSION AREA INFORMATION DOMAINS 9-8

9.3 INFORMATION DOMAIN SECURITY POLICIES 9-12
9.3.1 Intelligence Mission Function Information Domains 9-14
9.3.2 Tactical Mission Function Information Domains 9-17
9.3.3 Drug Interdiction Mission Function Information Domains 9-21
9.4 LOCAL SUBSCRIBER ENVIRONMENTS 9-25
9.4.1 Intelligence Site LSE Architecture (LSE-A) 9-28
9.4.2 Tactical Site LSE Architecture (LSE-B) 9-31
9.4.3 DEA Site LSE Architecture (LSE-C) 9-34
9.4.4 DOJ Site LSE Architecture (LSE-D) 9-35
9.4.5 Drug Interdiction User Site LSE Architecture (LSE-E) 9-35
9.5 SYSTEM SECURITY POLICIES 9-36
9.5.1 Intelligence Site System Security Policy (LSE-A) 9-36
9.5.2 Tactical Site System Security Policy (LSE-B) 9-36
9.5.3 DEA Site System Security Policy (LSE-C) 9-37
9.5.4 DOJ Site System Security Policy (LSE-D) 9-37
9.5.5 Drug Interdiction User Site System Security Policy (LSE-E) 9-37
9.6 SYSTEM SECURITY ARCHITECTURES 9-39
9.6.1 Intelligence Site System Security Architecture (LSE-A) 9-39
9.6.2 Tactical Site System Security Architecture (LSE-B) 9-46
9.6.3 DEA Site System Security Architecture (LSE-C) 9-47
9.6.4 DOJ Site System Security Architecture (LSE-D) 9-47
9.6.5 Drug Interdiction System Security Architecture (LSE-E) 9-47
9.7 INTERDEPENDENCY ANALYSIS 9-47
9.7.1 Interoperability Analysis 9-48
9.7.2 Security Service and Mechanism Interdependencies 9-53
9.7.3 Vulnerabilities, Residual Risk, And Recommended Countermeasures 9-57

APPENDIX A A-1

APPENDIX B B-1

REFERENCES REF-1

LIST OF ACRONYMS ACR-1

LIST OF FIGURES

FIGURE	PAGE
2-1. Security Policy and Requirements	2-2
2-2. Mission-Specific Security Architecture Development	2-3
3-1. Building Blocks of the DIS	3-2
3-2. DoD Technical Reference Model	3-4
3-3. TAFIM IM Integration Model	3-6
3-4. DIS Target Architecture	3-9
4-1. Abstract Security Perspective	4-1
4-2. Security View of LSEs	4-2
4-3. Example Variations of Network Elements	4-3
4-4. Example LSEs and Network Elements	4-4
4-5. TAFIM IM Integration Model and DGSA End System Mapping	4-6
4-6. Secure-to-Nonsecure LSE Communications	4-7
4-7. Secure LSE Communications	4-8
4-8. Security Service Allocation Summary	4-10
5-1. End System Security Architecture Abstract View	5-5
5-2. Security Context Software Abstract Relationships	5-16
7-1. Component Relationships for Creating a Security Association	7-6
9-1. Example Mission Functions and Their Relationships	9-3
9-2. Example Mission Area Information Domain Relationships	9-13
9-3. LSE Locations, Mission Applications, and Communications	9-27
9-4. Intelligence Site LSE Functional Architecture (LSE-A)	9-29
9-5. Intelligence Site Transfer System Architecture (LSE-A)	9-30
9-6. Intelligence Site Security Protocols (LSE-A)	9-31
9-7. Tactical Site Functional Architecture (LSE-B)	9-32
9-8. Tactical Site Transfer System Architecture (LSE-B)	9-33
9-9. DEA Site Functional Architecture (LSE-C)	9-34
9-10. DEA Site Transfer System Architecture (LSE-C)	9-35
9-11. Drug Interdiction Site Functional Architecture (LSE-E)	9-36
9-12. LSE-A End System Information Domain Allocations	9-43

LIST OF TABLES

TABLE	PAGE
2-1. Summary of Security Requirements Refinements	2-9
7-1. DGSA Security Protocols and Security Services	7-11
9-1. Hierarchical Policy/Mission Relationships	9-4
9-2. Intelligence Mission Area Positions and Responsibilities	9-5
9-3. Tactical Mission Area Positions and Responsibilities	9-6
9-4 Drug Interdiction Mission Area Positions and Responsibilities	9-8
9-5. Summary of Mission Area Information Domains	9-11
9-6. Intelligence Information Domain Security Policy	9-18
9-7. Intelligence Security Management Information Domain Security Policy	9-19
9-8. Tactical Information Domain Security Policy	9-20
9-9. DEA-Only Information Domain Security Policy	9-22
9-10. DOJ-Only Information Domain Security Policy	9-23
9-11. Drug Interdiction Information Domain Security Policy	9-23
9-12. DEA/DI Security Management Information Domain Security Policy	9-25
9-13. Domain Allocations by LSE	9-28
9-14. Intelligence Site System Security Policy (LSE-A)	9-38
9-15. DEA Site System Security Policy (LSE-C)	9-39
9-16. DOJ Site System Security Policy (LSE-D)	9-40
9-17. Drug Interdiction User Site System Security Policy (LSE-E)	9-41
9-18. LSE-A End System Security Architecture Elements	9-44
9-19. Intradomain Communications Interoperability Summary	9-51
9-20. Interservice Relation to System Requirements	9-53
9-21. Hypothetical Mechanism Strength Scale	9-55
9-22. Weighted Mechanisms by Domain and LSE	9-56
9-23. Weighted Mechanism Totals, by Category and LSE-Information Domain	9-57

1.0 INTRODUCTION

The Defense Information Systems Security Program (DISSP) was initiated at the request of the Assistant Secretary of Defense (Command, Control, Communications and Intelligence). The Defense Information Systems Agency (DISA) and the National Security Agency (NSA) have agreed to cooperate in achieving eight security objectives. The objectives involve security policy, architecture, standards and protocols, accreditation procedures, technology, transition planning, organizational improvement, and products and services availability. Accordingly, a DISSP Office was established which included an Architecture and Engineering Division responsible for the development of the Department of Defense (DoD) Goal Security Architecture (DGSA). The DISSP has since become a part of the Center for Information System Security (CISS) in DISA and the Architecture and Engineering Division is now a Directorate within CISS. The Directorate also will assist DoD organizations in the transition of existing systems, and in the development of new systems for the Defense Information System (DIS) in accordance with the DGSA.

Within the DISA, several efforts had been under way to define information system architectures for the DIS. All such architectural efforts have now been focused under the Technical Architecture Framework for Information Management (CIM, 1992).

1.1 PURPOSE

This DGSA, although described and defined separately in this document, is an integral part of the TAFIM being developed for the DIS. The DGSA is being developed in conjunction and harmony with the total requirements for automated services. Consideration is being given to the protection of information and system assets as part of the total view of the missions, threats, performance, interoperability, extensibility, usability, and cost of implementations. The TAFIM, including its security architecture, is intended to be generic and sufficiently flexible in its definition so that specific systems may be developed or modified to satisfy their specific missions. It includes the requirements of the joint tactical mission of command, control, communications, computers, and intelligence (C4I) for the Warrior (C4IFTW), while meeting the broader and longer range objectives of the DIS. The TAFIM is thereby a "goal information system architecture" and its security architecture is a "goal security architecture."

This document describes the DGSA, but it does not provide a specification for any particular information system or component. Rather, it specifies security principles and target security capabilities that will guide system security architects in creating specific security architectures that are consistent with the DGSA. Further, there is no

particular date by which the DGSA will be completely achieved. The DIS architecture and the DGSA are under constant development and will continue to evolve with time and with technological advances. The DGSA will also provide the basis for the development of security products and mechanisms that may be chosen by security systems engineers and integrators.

1.2 SCOPE

The DISSP was instituted to draw together various information system applications, information transport systems, programs, and architectural activities to bring about consistency, efficiency, and interoperability in the security designs for the DIS. Several programs and systems were identified, such as the Defense Message System (DMS), the Defense Information Systems Network (DISN), the Integrated Tactical/Strategic Data Network (ITSDN), and the DoD Multilevel Security (MLS) Program as candidates from which DISSP personnel could gather a complete set of security requirements. These programs cover the bulk of the DIS and are reasonable representatives of DoD information processing needs, but the DGSA must be applicable to all individual DIS programs and systems if the DIS objective of interoperability is to be met.

1.3 ARCHITECTURAL TYPES

Information system architectures range in definition and occur in sequence from abstract views to specific views of what is to be developed. Experience shows that four types are frequently used. The types are abstract, generic, logical, and specific, each of which is described below. The TAFIM presented in section 3 of this document is considered to be an abstract and generic architecture. Thus, the DGSA, as part of the TAFIM, also is abstract and generic. This document also includes an example logical architecture based on the DGSA in section 9.

1.3.1 Abstract Architecture

An abstract architecture begins with knowledge of the requirements and defines corresponding functions to be performed. It defines principles and fundamental concepts that guide the selection and organization of functions. Abstract security architectures cite principles, fundamental concepts, and functions that satisfy the typical security requirements. These concepts and functions are allocated to elements of an abstract definition of the information system architecture.

1.3.2 Generic Architecture

The development of a generic architecture is based upon the abstract architectural decisions. It defines the general types of components and allowable standards to be used, and identifies any necessary guidelines for their application. A generic security architecture proceeds from an initial allocation of security services and functions and begins to define the types of components and security mechanisms that are available to implement the security services with particular strengths. Any limitations in combining components and mechanisms because of incompatibility or security degradation must be cited in the guidelines for application.

1.3.3 Logical Architecture

A logical architecture is a design that meets a hypothetical set of requirements. It serves as a detailed example that illustrates the results of applying a generic architecture to specific circumstances. The only differences between a logical and a specific architecture are that the specific requirements are real, not hypothetical, and since the logical architecture is not intended to be implemented there is no need to perform a cost analysis. In logical security architectures, the logical design is accompanied by an illustration of the security analysis to be performed in specific architectures.

1.3.4 Specific Architecture

The objective of any system architect is to accomplish a level of design specification such that components may be acquired to implement the system. The specific architecture addresses components, interfaces, standards, performance, and cost. Specific security architectures show how all the selected information security components and mechanisms, including doctrine and supporting security management components, combine to meet the security requirements of the specific system under consideration.

1.4 PROCESS

Development of the DGSA is achieved through an iterative process, cycling through the following steps until a satisfactory generic security architecture (or architectures) is accomplished:

- Perform requirements analysis
- Create structure for the architecture
- Make security service allocations

- Select security components and mechanisms
- Perform interdependency analysis (evaluation)

Each iteration of this document will be a new version of the DGSA. Each version will present the abstract and generic views of the DIS target architecture and its security.

1.5 DOCUMENT ORGANIZATION

This document begins by stating a broad set of requirements in section 2 and then explains the relationship of the DGSA to the TAFIM and C4IFTW in section 3. In section 4, an abstract information system architecture is presented that includes the identification of major components of a generic information system, an abstract information model is discussed, and security responsibilities are allocated to the major components based upon realistic expectations of the protections that can be achieved. Section 4 also presents several key security concepts used through the remainder of this document. The major components identified in section 4 are then considered in greater detail, specifically end systems and relay systems in section 5, security management in section 6, transfer systems in section 7, and doctrine in section 8. Section 9 presents logical architecture examples of the application of the DGSA.

2.0 SECURITY REQUIREMENTS

This section addresses the security requirements applicable to the DGSA and the process by which organizations can identify the specific security requirements of their missions. Section 2.1 summarizes the DGSA security policy. Section 2.2 describes the DGSA security requirements. Section 2.3 discusses the DGSA derived security requirements needed to support multiple security policies.

An *information system* is a collection of information processing and communications components, and the environment in which they operate, that is used to support the operations of one or more missions. A *security policy* pertains to a mission and is based upon the threats to the means by which that mission is accomplished. A security policy (or, in a more general sense, a collection of security polices) documents the security requirements to be placed upon resources used by an organization. These security requirements express, for the information system personnel, the user organization's desired protection for its information and other system resources.

A security architecture designed to meet a specific mission's security requirements defines the security services and mechanisms and allocates them to components of the mission's information system architecture. Since the DGSA is intended to address the needs of all DoD organizations, it is a more general statement about the common collection of services and mechanisms any information system might offer and allocates the security services and mechanisms to generic components of information systems.

The DoD organizations that will employ the DGSA have many different missions. The security policy addressed by the DGSA is a general expression of the security requirements commonly found among the mission requirements of DoD organizations. Figure 2-1 shows that security policy and requirements are derived as a result of examining the threats to the mission and are therefore a subset of the mission's requirements. It also indicates the strong relationship among mission, users, information, and policy.

The establishment of security requirements follows the same process whether it is for the DGSA or a specific mission. The process is composed of the following mandatory steps: the information to be managed is identified; the operational requirements for the use of the information are stated; the value of the information is determined; and the potential threats to the information are identified. The security policy for either the generic case (the DGSA) or a specific mission can next be stated in terms of the requirements for:

- Protection for the information based on the potential threats
- Security services that afford the appropriate protection of the information based upon the value of the information and the threats to it

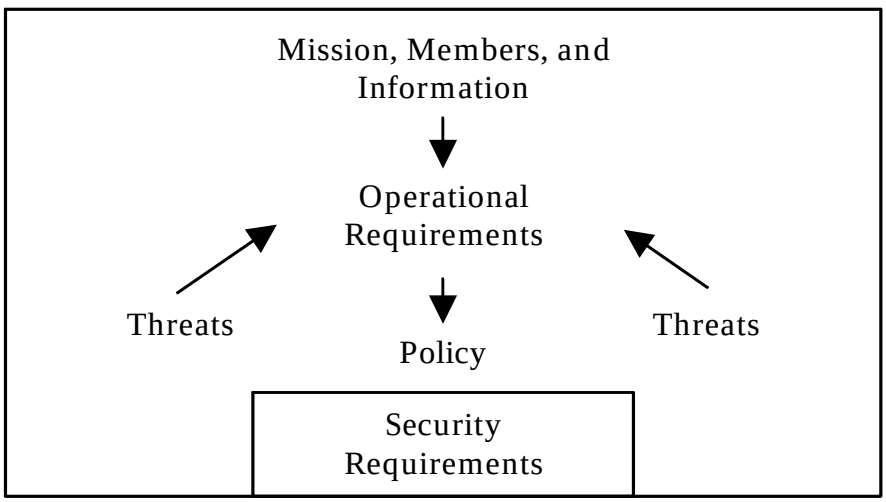

Figure 2-1. Security Policy and Requirements

As stated previously, the DGSA is a security architecture covering the full range of DoD missions and related information system security services and security mechanisms. The development of a mission-specific security architecture, as shown in figure 2-2, begins by applying the DoD security policy to the specific mission requirements to develop a mission-specific security policy, which includes identifying the appropriate security services and mechanisms an information system should offer to satisfy those requirements. The mission-specific information system security architecture is derived from this set of requirements and security services. This mission-specific architecture is stated as the set of mechanisms appropriate for providing the level of protection required. Guidance documents such as the TAFIM, and particularly the DGSA, should be applied to a specific information system architecture to ensure that the necessary security protections are appropriately allocated to specific information system components. Specific security architectures also need to address any applicable policy, public laws, and executive orders. Information system security architects should understand the complete methodology and the way other aspects of the DGSA are taken into account as demonstrated in the examples in section 9.

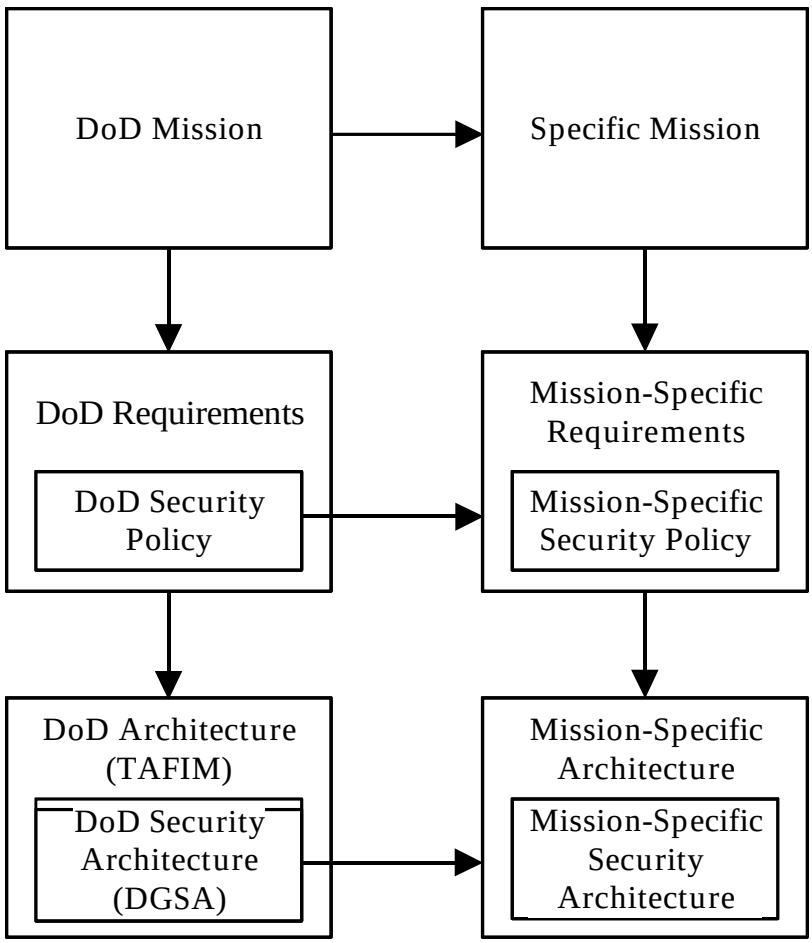

Figure 2-2. Mission-Specific Security Architecture Development

2.1 DGSA SECURITY POLICY

The DGSA security policy is based on the security requirements cited in the *DoD Information Systems Security Policy* (NSA, 1993), which is summarized as follows:

1. DoD information systems must support information processing under multiple security policies of any complexity or type, including those for sensitive unclassified information and multiple categories of classified information.

2. DoD information systems must be sufficiently protected to allow distributed information processing (including distributed information

system management) among multiple hosts on multiple networks in accordance with open systems architectures.

3. DoD information systems must support information processing among users with different security attributes employing resources with varying degrees of security protection, including users of nonsecure resources if a particular mission so dictates.

4. DoD information systems must be sufficiently protected to allow connectivity via common carrier (public) communications systems.

2.2 DGSA SECURITY REQUIREMENTS

This section describes the DGSA security requirements based on the security policy stated in section 2.1. The four security requirements discussed are Multiple Information Security Policy Support, Open Systems Employment, Appropriate Security Protection, and Common Security Management.

2.2.1 Multiple Information Security Policy Support

Although most current information systems support only one information security policy at a time, there has long been a desire by users to operate simultaneously at multiple sensitivity levels or under multiple security policies (e.g., by using multilevel secure systems) on a single device (e.g., workstation, outboard protocol device). Policy statement 1 above recognizes that support for multiple security policy operation must become more common. The successful implementation of policy statements 1, 3, and 4 largely depends on the ability of information systems to separate users and information at different sensitivity levels or to separate information subject to different security policies. That is, implementations must provide users with confidence that there will not be any security policy violations because shared information systems and communications systems are used that support users operating under differing security policies.

2.2.2 Open Systems Employment

DoD information systems must be open in the sense that potential connectivity among them always is supported, even if a particular request for communication is denied because of a security policy decision. Although the use of open systems as a high-level operational requirement (and the policy statements from which this requirement is derived) may seem to be focused on operational issues, it is equally critical to the DGSA in that it promotes a particular approach to providing

information security among cooperating DoD information systems. In the past, isolated systems were created and information was over-classified to satisfy security requirements. Given that users operating under different security policies may need to share components, and that complex policies for sharing and transferring information among users operating under different security policies must be supported, it is critical that truly open systems (both information processing systems and communications systems) be employed. Not only is this requirement directly derived from policy statement 2, but it supports policy statements 3 and 4 as well.

2.2.3 Appropriate Security Protection

Policy statements 2, 3, and 4 refer to information systems being "sufficiently protected" or supporting users by employing varying degrees of security protection. The combination of automated, procedural, and physical methods, from the complete set offered by a particular information system, appropriate for protecting a set of users and information can only be determined by those persons responsible for the particular information and who are able to assess its value and the threats to it. The corresponding generic DGSA requirement is that specific means must be available to users to invoke security mechanisms appropriate to the task at hand.

What constitutes appropriate security protection, in part, is affected by the security protection provided by the communications system that is used among distributed systems. Policy statement 4 requires that when common carrier communications must be used, the information systems must be prepared to provide all of the appropriate security protection. The only service that should be assumed from a common carrier communications system is availability.

2.2.4 Common Security Management

Like the open systems requirement, security management appears to be concerned with operational issues, but it actually provides the foundation for many of the security mechanisms that implement the security services chosen to satisfy the other security requirements. To ensure that distributed information processing is properly supported, the DGSA must address common security management. This commonality will allow security administrators to manage, in a uniform manner, systems that operate under multiple security policies in accordance with policy statements 1 and 2.

2.3 DGSA DERIVED SECURITY REQUIREMENTS

This section first discusses the refinement of the security requirements (section 2.3.1) and then the interaction between mission-related operational objectives and security requirements (section 2.3.2). The process of security requirements derivation is shown by example; it is not intended to identify every possible security requirement. The expectation is that developers will perform similar, but complete, analyses for specific systems.

2.3.1 Security Requirements Refinement

The refinement of the security requirements is stated as a set of security services, functions, or activities that will be allocated among users, administrators (acting on behalf of the users), information systems, and communications systems for a particular distributed information system architecture.

2.3.1.1 Multiple Information Security Policy Support

Several derived requirements are consequences of the need to support multiple information security policies. The most basic of these is the ability to support each security policy independently of other security policies supported in shared information systems or communications systems. Security policy enforcement is dependent on the ability of supporting information systems to maintain reliably the identities of users and the identification of information objects under each security policy. The traditional expression of policy enforcement is that all references by users (or processes representing them) to information objects must be mediated by a reference monitor. The DGSA adopts the reference monitor concept. (Note that any number of reference monitor implementations may be possible.)

When information processing operations are supported by distributed information processing systems, the security policy enforcement for information in transit is commonly supported by mutual authentication, access control, data integrity, data confidentiality, and non-repudiation communications security services. For local (e.g., within a workstation) information processing, a similar set of security services can be applied.

2.3.1.2 Open Systems Employment

When a user seeks to perform functions in a distributed environment, the user must be able to convey information to another user (or a process) that will become the basis for decisions about what (if any) kinds of interaction will be allowed. The

DGSA presumes that international standard protocols (or at least national or DoD standards, not industry proprietary schemes), information, and mechanisms will enable users to determine the capabilities and environment of other users or system processes with which they attempt to communicate. The determination may be made on the basis of information available before any communication is attempted (e.g., from a directory service), or where the determination is made as part of the initial communications service negotiation, or a combination. The result of such a determination might be that (within the information security policies shared by the users) the only common capability is to share only non-sensitive information or that no further communication is possible.

Beyond the normal means to begin distributed processing, standards for the representation and exchange of security information are needed. Some of this information is made available as part of the communications exchanges and some is provided through security management-related exchanges. Taken together, this information is used in the provision of various security services.

2.3.1.3 Appropriate Security Protection

The requirement for appropriate information systems security protection dictates that security mechanisms must be identified that implement security services at the level of protection required in security policies. Since some security mechanisms may be used to provide (parts of) multiple security services and some security services may be implemented by multiple mechanisms, a determination must be made that the mechanisms are appropriate individually and in combination. Initially, this is a technical activity, but the final determination involves deciding whether shortfalls in the collected security mechanisms can be accepted or whether additional measures must be put in place. This determination must be made by the users of mission information, or as is most common, the accreditor who represents the users.

2.3.1.4 Common Security Management

The basic elements that must be managed within the DGSA are users, security polices, information, information processing systems that support one or more security policies, and the security functions that support the security mechanisms (automated, physical, personnel, or procedural) used to implement security services. For each of these elements, the managed objects that constitute them must be identified and maintained. For example, users must be known and registered, the security policies must be represented and maintained, and information objects must be identified and maintained. The format for presenting the information in managed objects and operations on them must be standardized. Section 6 presents a detailed

discussion of these managed objects and an architecture for security management within the DGSA.

2.3.1.5 Summary of Security Requirements Refinements

The requirements refinements discussed above are summarized in table 2-1.

2.3.2 Interaction Between Mission-Related Operational Objectives and Security Requirements

This section describes the interaction between mission-related operational objectives and security requirements. This presentation is designed to promote a thought or investigative process that should be applied to specific missions. Not all of the operational objectives discussed here pertain to every mission.

2.3.2.1 Prevalence of Enterprise Initiatives

DoD-wide enterprise initiatives, such as Center for Information Management (CIM) and C4IFTW, impose operational objectives that have an impact on security. CIM promotes information centralization, information access, and interoperability. All three of these activities eliminate the idea of isolated or stand-alone implementations as a means of providing security. Their effect on security requirements is the need to consider both the coexistence of varying sensitivities of information on the same information system and the provision of proper separation, authentication, labeling, and access control. C4IFTW is designed to provide the war-fighting soldier with access to any information needed to do the job, regardless of sensitivity, media, or branch of Service. Such operational objectives also provide security challenges and considerations. System interfaces are quite different for war-fighting equipment, thus presenting new authentication issues. Access to the information in a pull-from (information-on-demand) mode emphasizes both interoperability and availability requirements. The integration of voice, imagery, and data requires data correlation and a general secure display (windows) implementation. The implications of CIM and C4IFTW and any other relevant initiatives should be considered for their effects on specific missions.

Table 2-1. Summary of Security Requirements Refinements

Multiple Security Policy Support	Open Systems Employment	Appropriate Protection	Common Security Management
– Enforce security policy – Maintain user identities – Maintain information identification – Provide data integrity service – Provide data confidentiality service – Provide non-repudiation service	– Provide common security capability identification – Use standard security information exchanges – Use standard security information representations – Provide authentication service – Provide access control service – Provide availability service	– Identify appropriate security mechanisms that provide required level of protection for each security service (individually and in combination)	– Identify and maintain user information managed objects – Identify and maintain information system managed objects – Identify and maintain supporting security function managed objects – Use standard managed object representations

2.3.2.2 Use of Off-The-Shelf Equipment

Economics have always been a driver in decisions to employ security solutions for information systems. Implementation of automated security measures has raised systems costs with questionable returns on investment. One of the reasons that costs of security measures have remained high compared to their value is that security measures have been implemented in specialized, often retrofitted, components. Particularly in the face of current budgetary constraints, it is highly desirable that security features become standard elements of commercial off-the-shelf (COTS) or government off-the-shelf (GOTS) equipment so that security has minimal impact on price. For this to happen, vendors must be persuaded to create products with security features that are integral parts of those products. Vendors will need to be convinced that a real market for such products exists. Evaluation, certification, and

accreditation must become streamlined and conclusive processes so that the vendors can be assured of reasonable investment and return. Creation of a viable security product market will depend on use of standards for commercial, international, and DoD use. Availability of COTS and GOTS products with integral security features will affect the ability to achieve mission security requirements.

2.3.2.3 Need for Increased Connectivity

A common and significant operational objective is to take advantage of computer and communications technology to accomplish the mission at hand. This objective can be partially achieved by increasing the potential for connectivity, making additional resources available. Other operational objectives demand that such increased connectivity cannot increase cost significantly. One approach to increased connectivity is to employ commercially available, common carrier networks. However, this approach introduces significant potential risks. There is always the possibility that a hostile entity, with access to the network, will use any means affordable to mount attacks on information systems using the network. The resulting security requirement is that the security mechanisms chosen to protect information must be adequate to deter such a hostile entity.

Increased connectivity and use of common carrier systems present a perfect environment for DoD-wide interoperability. The connectivity to common carriers will dictate lower layer (International Organization for Standardization (ISO) Open Systems Interconnection (OSI) Reference Model (RM), ISO 7498 (ISO, 1984) Layer 3 and below) standard protocols, while the DoD missions will have to address upper layer (ISO Layer 4 and above) standards for interoperability between local environments. This standardization will include authentication information, security protocols, key management and distribution, and security management information. Equivalent standards use for voice communications should be used. Additionally, the potential threat of a hostile entity will require standard methods of evaluating the protections afforded to information and other resources to assure that remote user environments are providing equivalent protection.

2.3.2.4 Access to Information and Resource

A common operational objective is to provide users with access to any information and resources needed to complete a task. This objective includes operational concepts such as information pull, distributed processing, and information sharing. Some missions will require support by non-DoD personnel and resources.

Security considerations cause enclaves to arise based on mission criteria that require separation of users and information, while operational objectives create the need to traverse enclave boundaries. For example, pull-from may mean information will come from another enclave. This requires interoperability of communications and security services. In dealing with access to and sharing of information and resources, the following security implications must be considered: establishment and separation of enclaves, interpretation and exchange of information in standard forms, and management of information.

Transparency in distributed processing is an often stated objective, that is, users wish to behave as if all resources are locally available. Users wish to be able to be authenticated once to the local system and then transparently interact with the other systems to access resources. The effect of this objective on security is that information systems must have adequate local authentication schemes and security management mechanisms that free the user from the burdens of procedures such as multiple logins.

2.3.2.5 Certification and Accreditation

Certification is the process of determining the effectiveness of all security mechanisms. *Accreditation* is the process by which an organization (or an individual on behalf of the organization) accepts or rejects operational responsibility for an information system's performance, including security, in supporting their enclaves.

Certification and accreditation are complementary procedures that need to be consistent, uniform, and applicable across DoD systems and products. Certification procedures have lacked uniformity and a clear path to completion. This deficiency has caused tremendous frustration on the part of both users and developers of systems. In many cases, accreditation procedures are subjective and ad hoc. The results of these procedures applied to particular products and systems should be of value to evaluators and accreditors of products and systems that have common elements. The challenge is to develop a set of uniform procedures that will limit and reduce the time to achieve product and system acceptance and that will eliminate disparities in the accreditation process. Uniform procedures will ensure consistent and interoperable security support for an enclave throughout a distributed environment.

Certain specific information is needed in the certification and accreditation processes that generally is not available today. Knowing the effectiveness of security mechanisms is an important part of determining how well required security services are supported. Knowing how a collection of security mechanisms interact and

support one another is important in assessing whether mission requirements have been met while minimizing security risks.

2.3.2.6 Need for Separation

Most missions will require the creation of several groups or enclaves joined together to achieve some specific purpose. It is also likely that the individuals involved will be members of more than one of these enclaves and will need to operate in two or more simultaneously. Organizations can no longer afford to build separate systems to support each of these enclaves, nor is it effective to require the user to change interface components (such as a workstation) every time the need arises to operate in a different enclave. The resulting security requirement is the establishment of criteria for mechanisms that allow multiple enclaves to share systems and information while guaranteeing the separation of information and users as necessary.

2.3.2.7 Maximizing Return on Investment

Operations today must exist in an environment in which major trends tend to be at odds with one another. Technology advancement has provided an opportunity to create an operational vision barely imaginable a few years ago. However, the high cost of transitions and diminishing budgets act against employing the new technologies. Intelligent strategies which may not reduce up-front costs but show valuable long-term benefits and reductions in costs will win favor. These strategies must support the long-term operational objectives. Such strategies include portability of applications and other software, continuous upgrades of hardware and software, ensuring scalability of applications and communications resources, reuse of software components, and reuse of certification and accreditation results. Each strategy has the post-transition value of providing low-cost growth paths if supported properly. Each strategy has an effect on security. Recertification of systems and products after change may be the most important of the strategies in its long-term payoff.

3.0 DIS TARGET ARCHITECTURE PERSPECTIVES AND EVOLUTION

This section summarizes the DIS target architecture perspectives and evolution. The DIS target architecture provides a consistent structure for distributed information management across the entire DoD. The objectives of the DIS target architecture are cost reduction, greater efficiency, greater mission effectiveness, interoperability, and consistent security.

The TAFIM is being promulgated by the Office of the Secretary of Defense (OSD). It mandates the use of open systems, software reuse, and data management, communication, and security standards to achieve significant DoD cost reduction, interoperability, and security for all defense mission areas. *The TAFIM is the DIS target architecture.*

The DGSA describes the security perspective of the DIS. The DGSA will become a separate volume of the TAFIM, and will later become integrated with portions of the TAFIM volumes. *The DGSA is the DIS target security architecture.*

The DIS target architecture is composed of mission information systems and the DISN. The completed DIS target architecture is the TAFIM with the integrated DGSA and DISN architecture guidance documents, which will be applied to and support the specific architectures of mission information systems.

The requirements drivers for the TAFIM, the DGSA, mission information systems, and the DISN are set by elements of the DoD hierarchy. These requirements drivers form the basis for all programs and joint initiatives within the DoD that develop specific applications for information management. The TAFIM Information Management (IM) Integration Model organizes the DoD hierarchy into five levels (enterprise, mission, function, application, and personal), which are described in section 3.1.2. The Integration Model provides a basis for assigning IM integration responsibilities by making explicit the intersections of functions that can be categorized and assigned management responsibilities. Thus, requirements at all levels of the TAFIM IM Integration Model are considered in the DIS target architecture and are combined to drive major application development programs and joint initiatives. Major program and joint initiative application requirements, which are used to develop specific architectures, are therefore consistent when specific architecture development is guided by the generic DIS target architecture. These requirements interrelationships are not accidental. A coordinated, consistent, orderly, and effective DIS evolution cannot be realized without this interrelated means of expressing requirements.

Open system standards are a major element of the DIS target architecture. Many of the security requirements, discussed in section 2, are directly attributable to the open systems mandate. Requirements for compliance with standards are pertinent to all levels of the TAFIM IM Integration Model. Only in special cases will exceptions be made to specific standards applicable to a mission area (e.g., in tactical systems, which in limited cases may have unique standards).

Time constraints, the state of technology, budget limitations, and backward compatibility act as limitations to achieving the DIS target architecture. Therefore, *transition strategies* are necessary. Some requirements that drive the DIS target architecture cannot be achieved by any specific date. Such requirements are categorized as *far term*. Similarly, many mission, function, or application requirements will be in this same category. Many transition strategies (categorized by estimated timeframes) will be developed to achieve the realization of the DIS target architecture in evolutionary phases. Specific mission, function, or application programs will be expected to develop transition strategies that establish a planned evolution consistent with DIS target architecture transition phases. The DIS target architecture evolution is discussed in section 3.3.

Figure 3-1 illustrates the objectives, guidance, relevance, and some example implementation and integration programs that are building blocks in the evolving DIS target architecture.

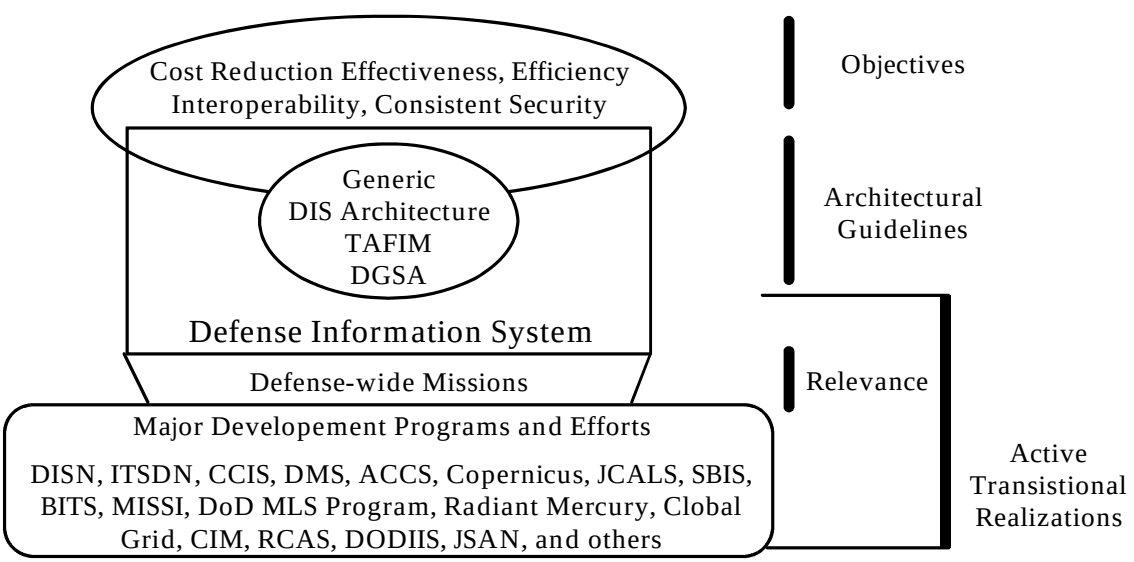

Figure 3-1. Building Blocks of the DIS

3.1 TAFIM GUIDANCE

The TAFIM is composed of several volumes. For the purpose of synopsis and correlation with the DGSA, TAFIM volume 2 has served as the reference. The TAFIM cites two concepts that apply to the DIS security architecture the DoD Technical Reference Model (section 3.1.1), and the TAFIM IM integration model (section 3.1.2).

3.1.1 DoD Technical Reference Model

The detailed DoD Technical Reference Model (TRM) is illustrated in figure 3-2. The application layers are delineated based on their mission-specific and general-purpose use. The *mission area applications layer* pertains to those applications tailored to specific mission needs (e.g., decision aids, battle management applications, imagery processing applications, message and signals processing applications). The mission area applications may make use of general support applications. The *support applications layer* provides generic services, common to multiple mission areas (e.g., word processing, spreadsheets, Graphical User Interface applications, database query language application, network services applications).

The *application platform layer* provides the common information processing and communications functions used in general-purpose and multimedia information management processing environments. A platform is composed of one or more processors, operating systems, kernel functions to support the operating system, and peripheral External Environment Interfaces (EEIs) controlled by the kernel functions for devices such as terminals, printers, and storage devices. The application platform layer also includes Open System Environment (OSE) programming services, user interface services, data management services, data interchange services, graphics services, and network services provided to mission-specific and support applications through Application Program Interfaces (APIs). APIs provide the standard interfaces between mission and support applications, and between support applications and the platform's operating system, which may extend to applications on other platforms using the network services and the kernel-controlled EEIs used to interface to communication networks. The network services provide the communication protocol stacks in the platform necessary to communicate, end-to-end, between support or mission-specific applications. The network services applications may include support for data, voice, video, and imagery (e.g., facsimile).

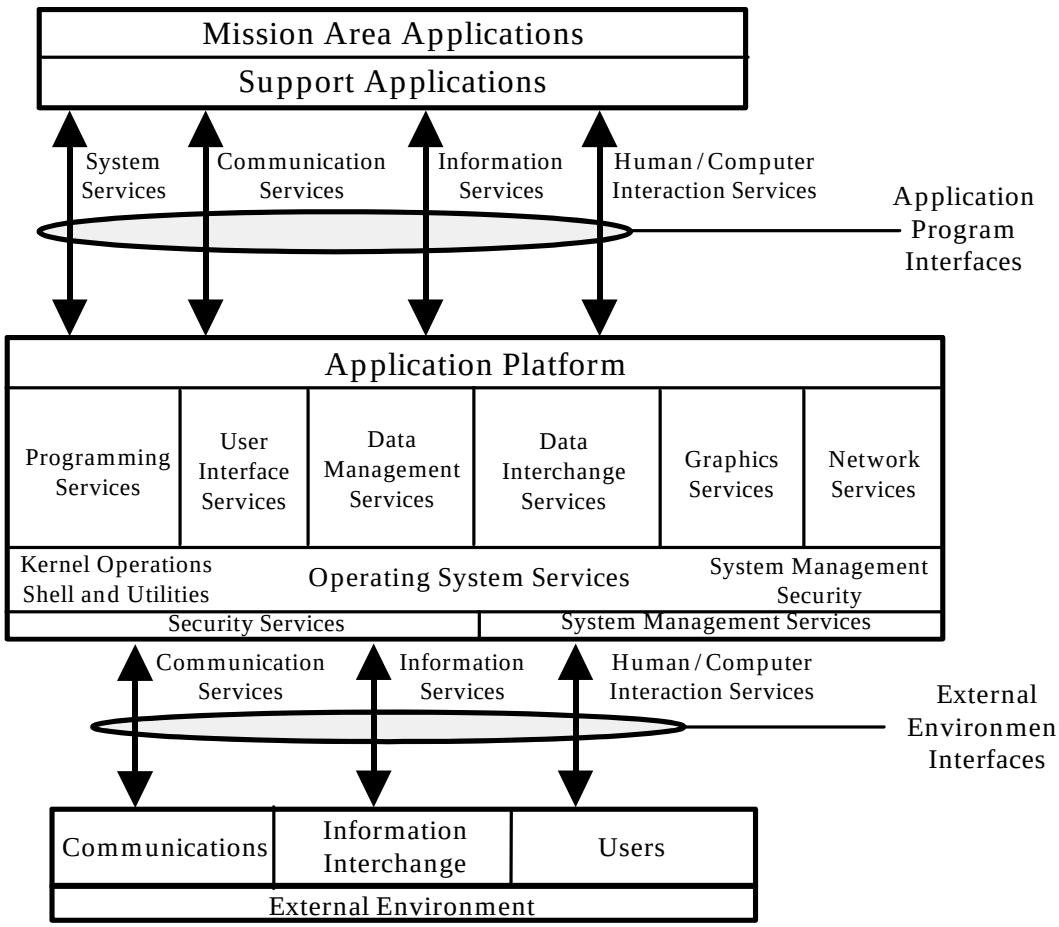

Figure 3-2. DoD Technical Reference Model

The *external environment layer* provides the external communication services necessary to provide platform-to-platform (information system-to-information system), and subsequently application-to-application services, across communication networks. External communications (the communications infrastructure) are composed of local level (e.g., base level) networks, and wide area, regional or metropolitan area (e.g., support for several bases or tactical assets), and global networks. The external environment layer also provides information services and human or computer interaction services for the platform to interact with local information interchange devices (e.g., peripherals) and with users. The external environment interfaces will be compliant with existing and future standards.

3.1.2 TAFIM Information Management Integration Model

The TAFIM IM Integration Model provides a basis for assigning IM integration responsibilities. This model makes explicit intersections of functions that can be categorized and assigned management responsibilities. This model is illustrated in figure 3-3.

Requirements established by a higher authority are inherited (must be adhered to) by the lower levels, and are expanded at each of the lower levels to meet specific operational requirements. Integration issues between levels must be addressed. Except at the enterprise and the personal levels, integration issues also must be addressed among the components within each level.

The *enterprise level* includes those elements of information management that are mandatory across the entire DoD. Elements such as policy and doctrine are established by both OSD and the Joint Staff that apply across all DoD missions. Such policy and doctrine also may be mandated by a higher authority, such as presidential executive orders and national decision directives, Office of Management and Budget (OMB) circulars, and public law. The enterprise level also promotes information technology (e.g., technical and data standards), reference models and technical architectures, methods and tools, and shared computing and telecommunications services. The TAFIM, the Defense Information Infrastructure (DII), and the DGSA are all elements of the enterprise-wide level that combine to provide IM guidance and direction in support of DoD-wide missions.

The *mission level* includes those elements of IM that are mandatory across an entire mission area, and coordinates those elements of IM that are required for inter-mission needs. The mission level is responsible for defining mission-specific functional areas (described below) necessary to fulfill the operational requirements developed to conduct and support the mission. The mission areas defined by the TAFIM are Business, Command and Control (C2), and Intelligence. The purpose of the DIS is to support the needs of all of these DoD mission areas. From this IM perspective, C4IFTW is a mission area guidance focus for C2.

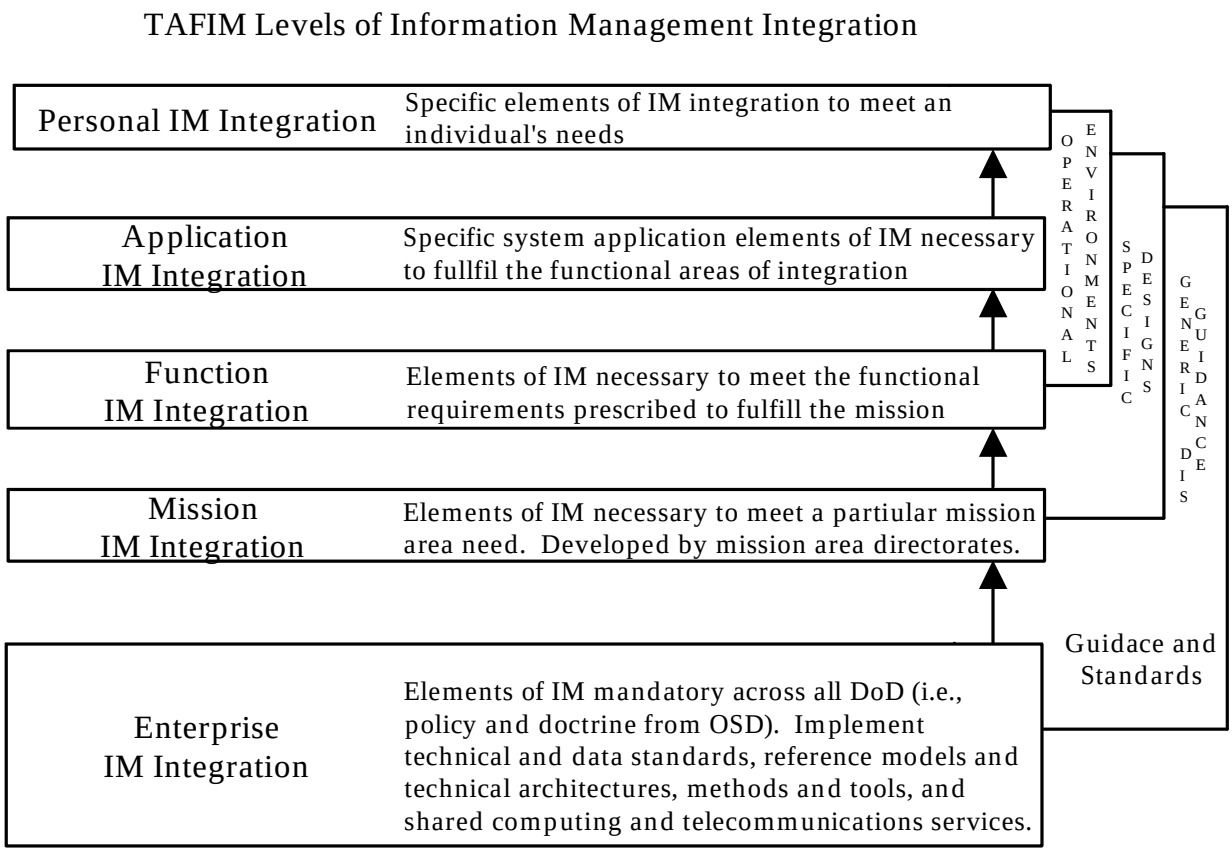

Figure 3-3. TAFIM IM Integration Model

At the *function level*, mission areas are refined into functional areas. Functional areas are generally categorized by the roles and responsibilities a particular organization or group of organizations perform to fulfill particular areas of the mission. From an information management perspective, a functional area includes those information services necessary to fulfill the specific functional areas of the mission. It is possible that functional areas of different missions may be of use in other missions. In these instances, common functional areas may be created to cross mission boundaries.

The *application level* includes the development, maintenance, and operation of the information system applications that provide required automation support to mission area functions. Integration at this level encompasses system interoperability, data sharing, and other technical issues that enable the efficient operation of information services. Integration at the boundary between the application level and the function level includes access to subject-matter databases and other system functions. The

DMS is an application level program implemented to support functional areas of all missions. The DMS Target Architecture and Implementation Strategy, as a DoD-wide application level architecture and strategy, is coordinated at the enterprise level and provided to all mission and functional areas of DoD as its common secure messaging system application service, accommodating both individual and organizational message exchange needs.

The *personal level* addresses user support requirements that involve integration with the application level, while preserving privacy, individual choice, and personal preference at the desktop and workstation. Integration factors include stability and consistency of the human-machine interface to enhance personal productivity by insulating the user from the unique characteristics of individual systems at the system application level.

3.2 C4IFTW GUIDANCE

"The C4I for the Warrior concept will give the battlefield commander access to all information needed to win in war and will provide the information when, where, and how the commander wants it." (Powell, 1992)

The C4IFTW concept is both a vision and a road map to tactical C4I modernization. It provides the structure within which a unified, joint force structure will meet its information management and communications needs to fight wars and low-intensity conflicts effectively. Although each of the Military Services currently has its own unique C4I modernization efforts under way, C4IFTW imposes mandates for interoperability and shared resources on each of the Service-unique architectures.

The C4IFTW concept identifies three distinct phases. The first is a *quick-fix phase*, which provides translators, gateways, and joint standards to achieve quick-reaction (but not necessarily efficient) interoperability. The second phase, called the *mid-term phase*, provides modular building blocks and total interoperability for new C4I systems, and a jointly shared wide area network. The third phase, called the *objective phase*, provides evolving and advanced technologies identified and assimilated for the war fighter. This phase also provides a standardized interface environment and a global C4I network of fused information for complete requirements realization.

The quick fix phase only partially maps to the TAFIM because of the existence of legacy systems. The mid-term phase maps conceptually to the TAFIM in the perspective of transitioning toward the DIS target architecture. The objective phase maps directly to the TAFIM and reflects the realized DIS target architecture. The complete range of MLS requirements postulated by C4I tactical operations only can

be achieved through the employment of the DGSA (i.e., requirements expressing the need to accommodate unclassified through Top Secret/Sensitive Compartmented Information).

3.3 DIS TARGET ARCHITECTURE EVOLUTION

The development of the DIS will be evolutionary, but is required to be consistent with the TAFIM enterprise level mandates. The most abstract representation of the DIS is that of user elements connected to one another via local communications networks, which in turn have access to wide area communications networks. The wide area communications networks are, and will continue to be, DoD-owned (e.g., DoD-owned Military Satellite Communications resources), common carrier (U.S. and foreign) leased, and non-satellite radio frequencies. Commercial carrier services will be used to a much greater degree for both the tactical and non-tactical C4I network infrastructure over time. Eventually, it will no longer be prudent or practical from a cost perspective for DoD to own vast wide area communication resources, although some will have to be retained for survivability and to meet special requirements. User elements and their local communications may be fixed or mobile. The DIS target architecture, reflecting user elements and local and wide area networks, is illustrated in figure 3-4.

It must be understood that information systems are not restricted to computers and computer networks. Although exceptions apply to what follows, the intent is to be as widely inclusive as possible in considering applications and environments, while preserving significant freedom for architects and implementors of specific information systems. In particular, the following considerations apply to information systems in the DGSA.

3.3.1 Voice-Video/ Imagery-Data

While the bandwidth, flow characteristics, error susceptibilities, and human interfaces to these media types may be different, they can be represented as bits of information and, therefore, are accommodated by digitally based processing and telecommunications.

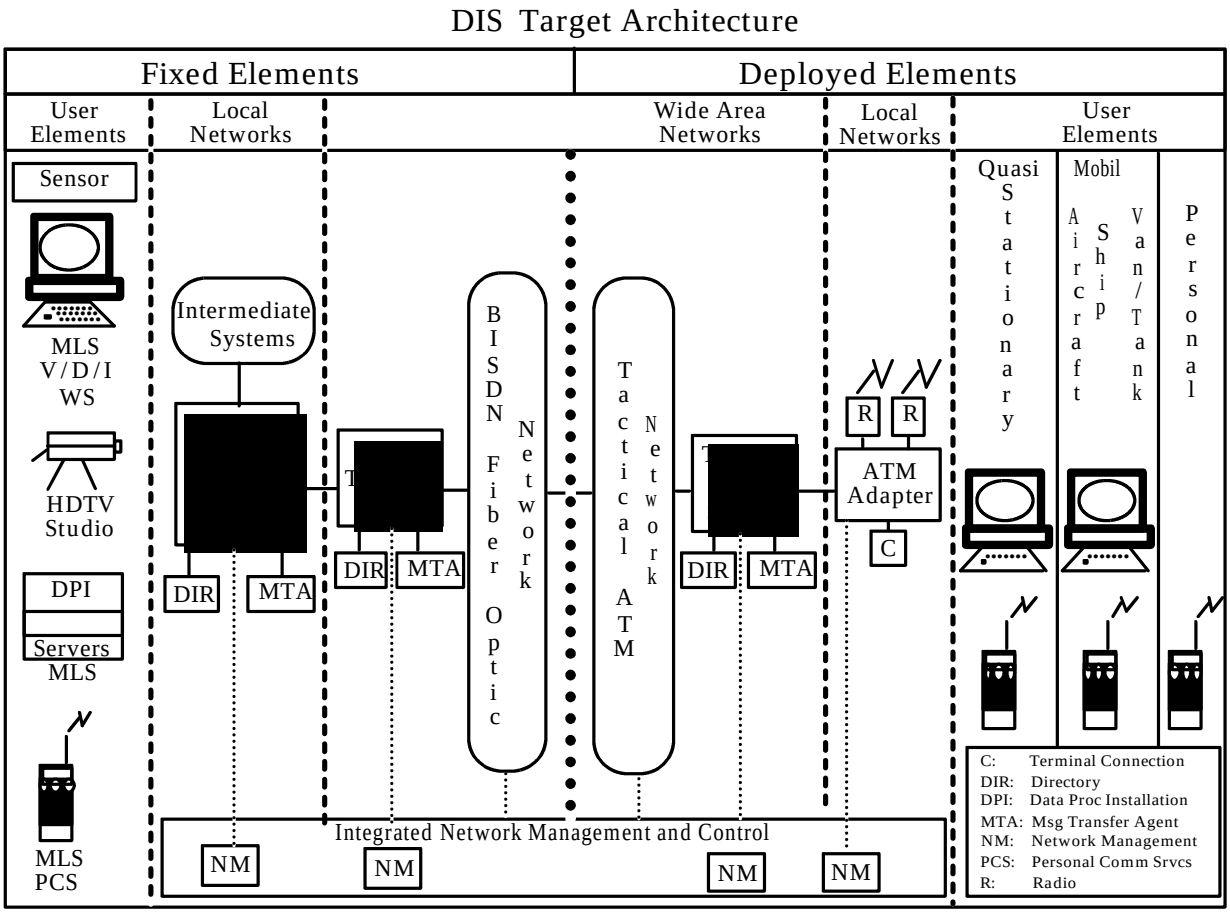

Figure 3-4. DIS Target Architecture

3.3.2 Portable-Mobile-Fixed

Physical portability and mobility are achieved through efficient packaging, low-power design, and transmission media independence. It is assumed that bit integrity will be inherent in the communications networks used and that residual errors will be the responsibility of user systems components (as part of availability services). Likewise, naming and addressing issues are assumed to be solved in the context of open systems communications. Security mechanisms such as authentication and cryptography are required to support portable and mobile information systems.

3.3.3 Secure-Nonsecure

The requirement for open systems demands the compatibility of secure and nonsecure information systems. The DGSA is based on the expectation that the integration of secure and nonsecure components will be necessary to achieve the desired degree of protection, and that products will be developed, with varying degrees of protection, that will operate in open systems.

3.3.4 Workstations-Servers-Hosts

Automated information systems will require relatively minor hardware changes, but significant operating system, support software, and application software changes will be needed to achieve security as envisioned in the DGSA.

3.3.5 Programmable-Nonprogrammable

Weapons systems, sensors, and other remotely controlled devices will have custom interfaces and high demands for communications availability, but they can be made to fit well into information systems architectures. The security, as well as the operation, of unmanned objects that may be lost or destroyed needs careful consideration in implementation. Nonprogrammable devices need careful consideration in terms of their security life cycle but are otherwise not a factor in the DGSA.

4.0 SECURITY VIEWS AND CONCEPTS

This section, which is based on the information system architectures identified in section 3, first describes the security views of those architectures. The DGSA builds on elements from those architectures (section 4.1) and begins to address the security service allocations for those elements (section 4.2). To accomplish the security service allocations, several views and concepts are presented that support the DGSA. This section provides the first exposure to the primary security concepts (section 4.3) that will be used throughout the remainder of the DGSA.

4.1 INFORMATION SYSTEM ARCHITECTURE SECURITY VIEWS

The DIS target architecture, illustrated in figure 3-4, first divides the DoD world into fixed elements and deployed elements. These elements are subdivided into user elements and network elements (e.g., local area, wide area). From a security perspective, the subdivision into user and network elements is particularly useful for establishing architectural views to which security services can be allocated.

4.1.1 Abstract Information System Architecture Security View

For security purposes, the view of the DIS target architecture shown in figure 3-4 is first simplified into an abstract perspective, defined as *local subscriber environments* (LSEs) that are connected to one another by a *communications network* (CN). Figure 4-1 illustrates this first security view of the information system.

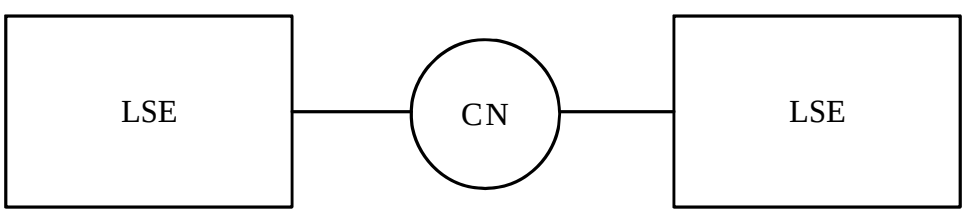

Figure 4-1. Abstract Security Perspective

The LSEs include all devices and communication systems under user (organization) control. The CN provides communication capabilities that allow LSEs to share information.

4.1.2 LSE Generic Security View

The first generic information system architecture security view concentrates on LSEs, which may be composed of open systems and *relay systems* (RSs), as described in ISO 7498, and *local communications systems* (LCSs) within a given operating environment. Open systems accessed directly by users are conventionally referred to as *end systems* (ESs) and will be throughout the rest of the DGSA. The *transfer system* includes the communication protocols integrated into end systems and relay systems, and the interconnecting LCSs and communications networks (CNs). LCSs are controlled within LSEs, but CNs are not controlled within LSEs. Figure 4-2 illustrates this generic LSE security view.

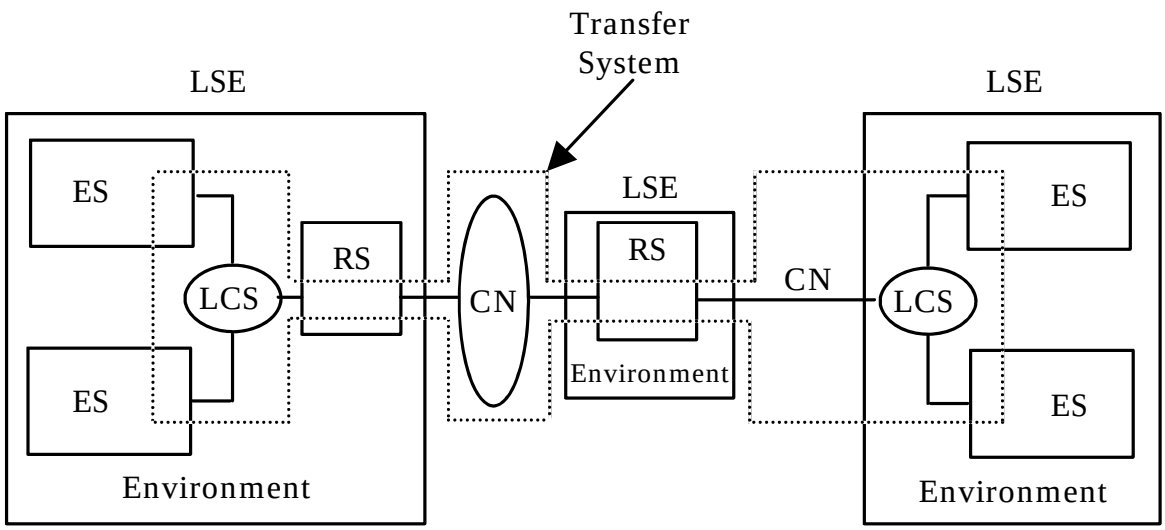

Figure 4-2. Security View of LSEs

An LSE may contain a single end system such as a workstation, a single relay system such as a router, or a complex interconnection of end systems and relay systems through LCSs. All components of the information system architecture are either part of an LSE or are CNs. This security view does not imply that LSEs are only connected to one CN or that they are connected only in pairs.

4.1.2.1 CN Description

Figure 4-3 shows two types of network elements. In one type of network element (e.g., CN_b), both ends of the (private) link are within LSEs and can be protected with

complete traffic flow security (TFS). In the other type of network element (e.g. CN_a), one end of each link terminates in the commercial control zone, and address information is needed by the switches in the commercial control zone, which limits the extent of TFS possible. One consequence of providing complete TFS between two LSEs is that the link cannot be used for any other purpose and, thus, creates a closed system. Given the importance of requirements for open systems, apparent requirements for complete TFS between LSEs must be examined very carefully. Figure 4-3 provides examples of the network element variations that are consistent with the generic security view.

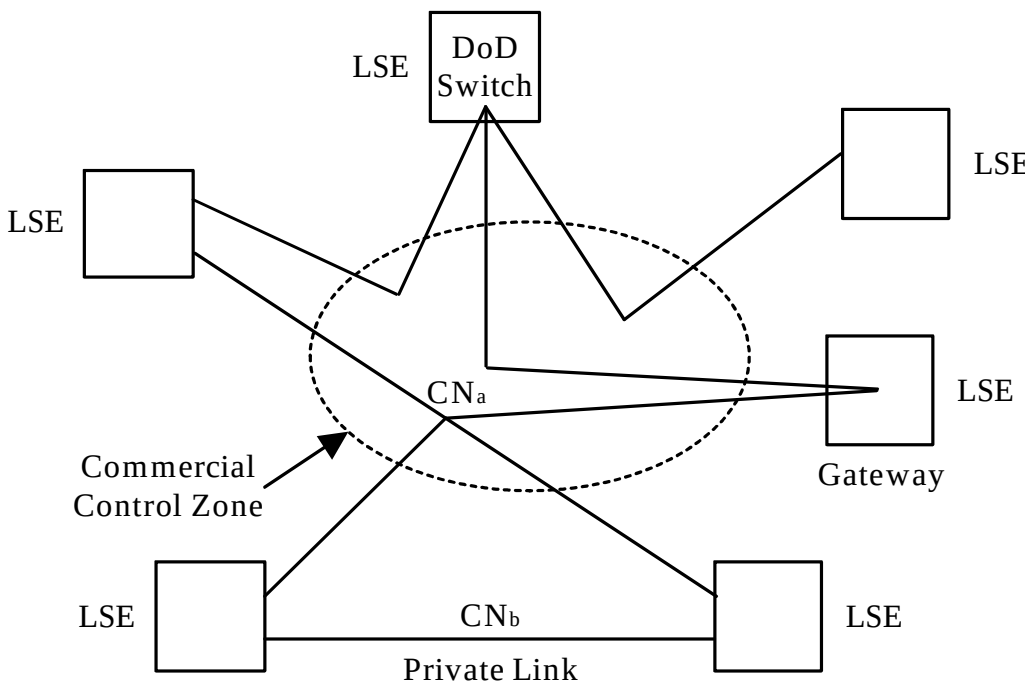

Figure 4-3. Example Variations of Network Elements

Network elements include transmission systems, switches, routers, gateways, and management systems, and, possibly, network-specific servers. Without further definition, there is an overlap in functionality between some LSE components and some CN components. This distinction has always been problematic in the telecommunications industry, but is clarified somewhat through the delineation of ownership, management, and protection responsibilities. In most situations, CNs are owned and operated as utilities that provide services to the LSEs on the basis of fee-

for-service agreements (e.g., leased, payment-on-demand). When CNs are commercially controlled, profound security implications result. Even when the enterprise (i.e., DoD in this case) owns and operates private network elements (e.g., MILNET), their administration is traditionally separated from LSE control. Thus, distinct lines of demarcation are available for management and control of CNs and LSEs. Nonetheless, cooperation must exist among network element management entities to achieve total communications system management. Figure 4-4 illustrates some of the types of CN elements and LSE components within the range of consideration.

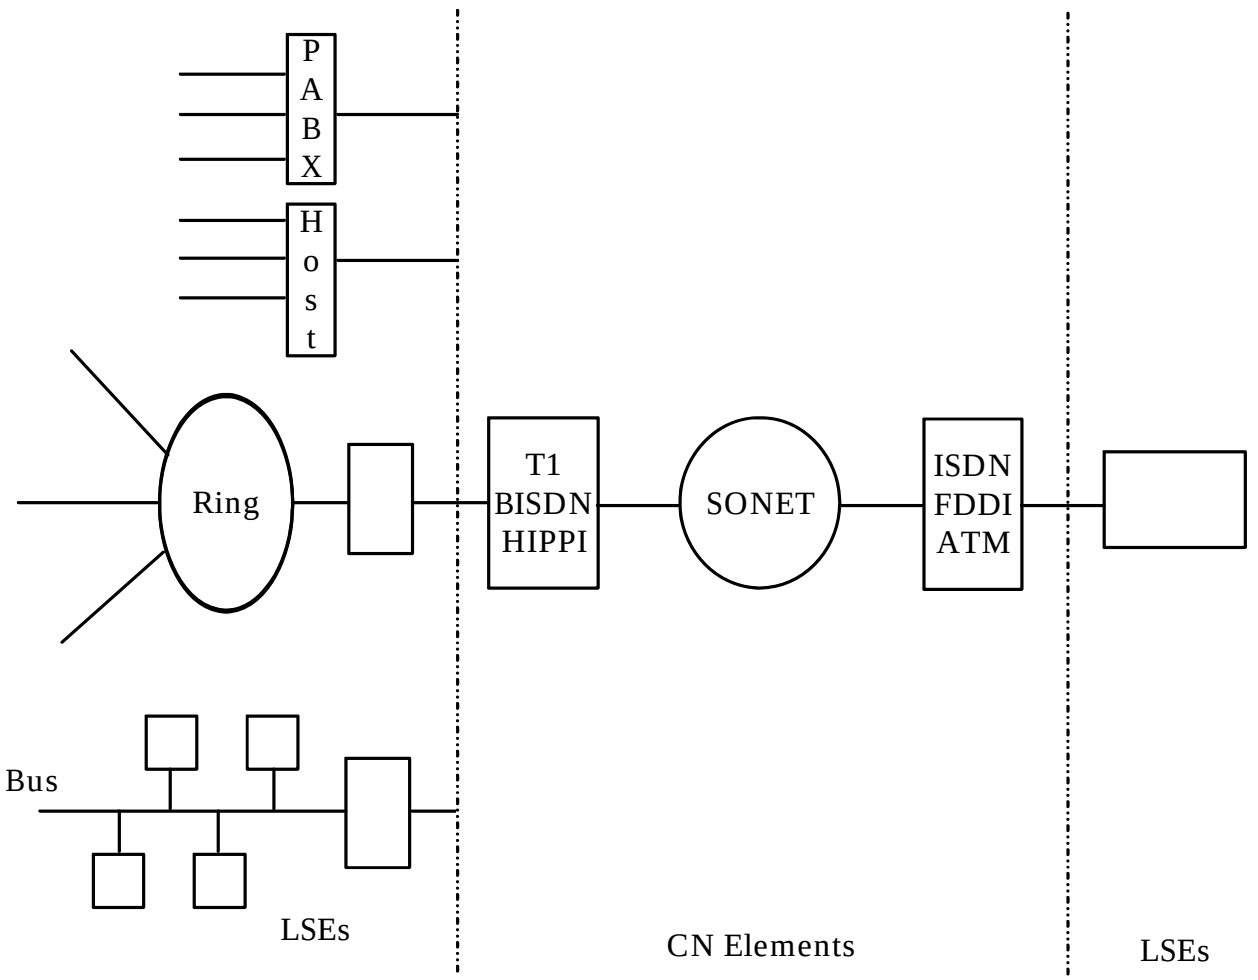

Figure 4-4. Example LSEs and Network Elements

4.1.2.2 LSE Description

Within the definition of the LSE are three generic types of components:

- End systems (e.g., workstations, servers, telephones, radios, mainframes)
- Relay systems (e.g., multiplexers, routers, switches, cellular nodes, message transfer agents)
- LCSs (e.g., rings, buses, wire lines)

The principal distinction between end systems and relay systems is that end systems support users with direct human interfaces and personal or general applications, while relay systems are only indirectly accessible by users and the functionality is limited to information transfer relay functions. Some relay system functions may be performed in many communications protocol layers (see section 7). Transfer system functions are carried out using communications protocols in end systems and relay systems, supported by the communications functions in LCSs and CNs, as illustrated in figure 4-2. The principal example of network service definitions is the ISO OSI protocol stack.

4.1.3 TAFIM IM Integration Model Security View

The end system security view now is mapped to the TAFIM Information Management Integration Model (section 3.1.1). The end system or relay system application layer (i.e., general, management, and transfer system applications) maps to both the mission and support application layers in the TRM. The end system or relay system operating systems and portions of the supporting security functions map to the TRM platform layer. Finally, portions of the end system or relay system transfer system applications and supporting security functions map to the TRM communications layer (e.g., communications, information interchanges, external environment interfaces). Figure 4-5 shows this mapping.

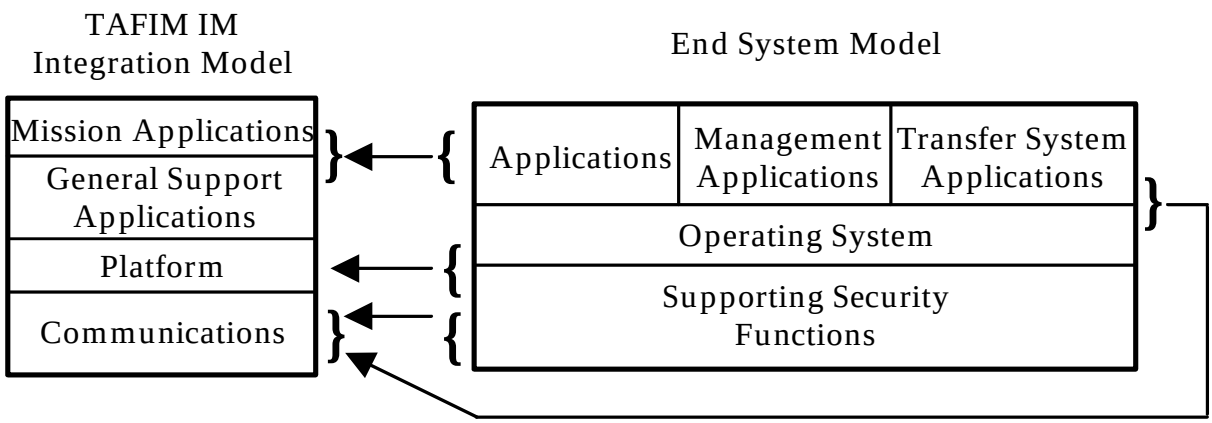

Figure 4-5. TAFIM IM Integration Model and DGSA End System Mapping

4.2 SECURITY SERVICE ALLOCATIONS

The DGSA security services are based on those defined in the ISO 7498-2 (ISO, 1989a) for data communications. These security services include authentication, access control, data integrity, data confidentiality, and non-repudiation. (The *OSI Security Frameworks*, ISO 10181, is a multipart standard that discusses each of these services, plus security audit and key management, in considerable detail.) In the DGSA, availability is considered to be a basic security service, also. Since the DGSA applies to all aspects of information system security, the basic services are considered to apply not only to the transfer system, but are interpreted to apply to the entire LSE. This section discusses security service allocations to LSEs and CNs, which can to traced back to the DIS target architecture and TAFIM communications elements. These allocations establish basic decisions upon which the generic DGSA is based.

4.2.1 CN Security Service Allocation

Figure 4-1 presents the most abstract view of the DIS target architecture—LSEs connected by a CN. In responding to the requirements of section 2, the DGSA makes only a security service allocation of communications availability to CNs. CNs must provide an agreed level of responsiveness, continuity of service, and resistance to accidental and intentional threats to the communications service availability. Any security requirements for complete traffic flow security must be satisfied through measures employed in LSEs.

The reliability, flexibility, contingency actions, management, and preventive maintenance of CNs are some of the factors that will determine the availability of

communications services. Protection of CN resources from accidental or intentional damage is both a security concern and, in the commercial world, a direct financial concern. Well-designed and well-managed CNs should exhibit graceful degradation in service and should provide for establishing priorities of service. CNs are not relied upon for the confidentiality or integrity of the information they transfer. Failures in CNs can only result in the delay, misdelivery, or non-delivery of otherwise adequately protected information. The purpose of CN management, which is to counter these failures, is identical to that of the security service of availability. Some CNs can provide physical protection of service access links (e.g., buried cable, selective protected routing), which may be adequate for some traffic flow security needs.

4.2.2 LSE Security Service Allocation

The remaining security service allocations are to LSEs. The protection of LSEs is accomplished by physical, administrative, and personnel security mechanisms. Physical LSE boundaries can limit facility access to authorized personnel. Protection of LSEs is provided in part by the logistical support system (e.g., configuration management control). In turn, LSEs provide protected environments for their end system, relay system, and LCS components.

The open systems requirement demands that LSEs with highly sensitive information must have the ability to communicate with nonsecure as well as with secure LSEs. The architectural model for such LSEs is shown in figures 4-6 and 4-7.

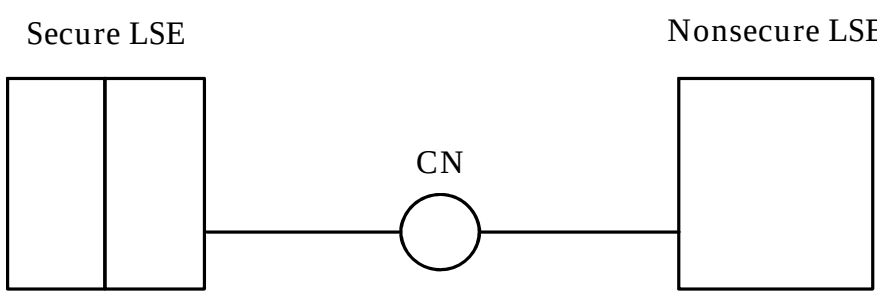

Figure 4-6. Secure-to-Nonsecure LSE Communications

In figure 4-6, a secure LSE is communicating with a nonsecure LSE that must be assumed to include hostile entities if the total information system is truly open. In this situation, no transfer system security services are used to protect information in

transfer because none are needed and the nonsecure LSE offers no such services. The secure LSE must isolate its sensitive information (shown as shaded in the figure) and protect it with its own security mechanisms.

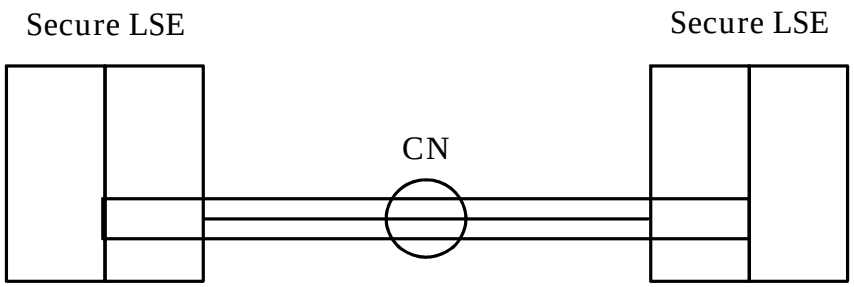

Figure 4-7. Secure LSE Communications

In figure 4-7, both LSEs are secure and cooperate to provide transfer system security services to protect the information in transfer. The secure LSEs must still protect themselves from nonsecure LSEs that are connected to the CN. The requirement for open systems provides serious challenges to the security architecture of LSEs.

As previously discussed, the security service allocations to the LSE (primarily identification and authentication (I & A) of personnel, and access control to facilities) are implemented as doctrinal mechanisms (i.e., physical, administrative, and personnel). These mechanisms are discussed in detail in section 8.

4.2.2.1 LCS Security Service Allocation

The first DGSA security service allocation within LSEs is that LCSs are required only to provide the availability security service for communications among end systems and relay systems within LSEs. (Other security services may be provided in LCSs for local purposes if they do not interfere with other requirements, such as interoperability with other LSEs.) All other security service allocations, including additional support for the availability service, are provided in end systems and relay systems.

4.2.2.2 End System and Relay System Security Service Allocation

The second security service allocations within the LSE are to end system and relay system hardware and software, so the hardware protects the software and the software protects information being processed, transferred, or stored. End system

and relay system hardware and software collectively provide the security services of user identification and authentication, access control, data integrity, data confidentiality, and availability. Details on the end system and relay system security architecture are discussed in section 5.

4.2.2.3 Transfer System Security Service Allocation

Security services implemented within protected end systems and relay systems provide the basis for the protection of information being transferred. The third security service allocations, which complete the provision of security services, comprise the portion of the transfer system that is responsible for peer entity and data origin authentication, access control, non-repudiation, confidentiality, integrity, and availability of information in transfer. The protection of information being transferred enables the protected distribution of security-relevant information as well as user information. The sharing of identification and authentication information, audit records, key management information, and policy and privilege management information among LSEs can be safely accomplished if the transfer system is protected.

The security architecture concepts defined in ISO 7498-2 are expected to form the basis for nearly all future open systems security. This standard is adopted by the DGSA, along with the additions and modifications for local area network (LAN) security (e.g., the series of standards developed in the IEEE 802.10 committees) and the security protocols that are adopted by the ISO community. ISO communication protocols (inclusive of security protocols) and security management applications and protocols are simply a portion of end system and relay system hardware and software compositions. Sections 6 and 7 deal with these subjects extensively.

Security service allocations have been made for the architecture views described in section 4.1. Figure 4-8 summarizes the allocations to the CN (discussed in section 4.2.1) and the LSE and its components (section 4.2.2).

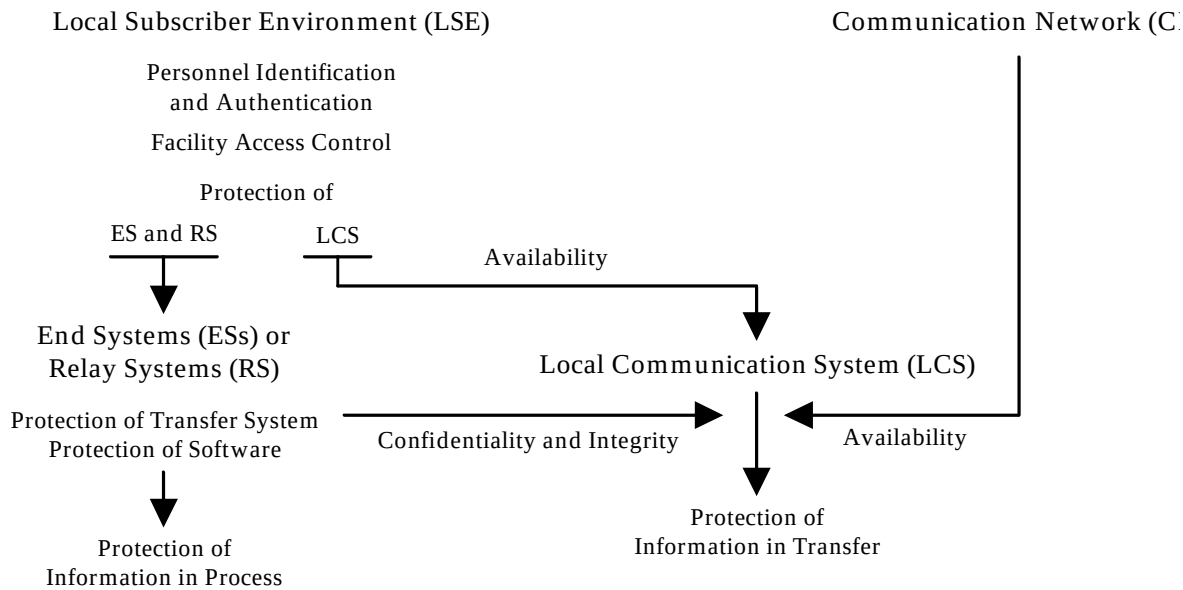

Figure 4-8. Security Service Allocation Summary

4.3 SECURITY CONCEPTS

The most significant capabilities of the DIS target architecture are distributed processing and open communications. The architecture is a rich interconnection of an unbounded number of heterogeneous LSEs. The objectives for security in such an environment are to maintain open and distributed capabilities and yet be able to establish and enforce dynamic mission and information security policies. A simple characterization of such an environment is that resources and information may be shared or isolated as desired. Several security concepts are set forth here that support these objectives: information domains, strict isolation, interdomain information sharing and transfer, multidomain information objects and policies, absolute protection, uniform accreditation, and security management. These security concepts are discussed in sections 4.3.1 through 4.3.7 below. The ways in which these concepts influence and are supported by the DGSA architectural components (end systems and relay systems, security management, transfer system, and doctrine) are detailed in sections 5, 6, 7, and 8.

4.3.1 Information Domains

The management of information is accomplished by individuals and groups of people who create, collect, process, categorize, store, transfer, and communicate particular

information. The value of that information and, therefore, the required protection of that information is determined by the group. The group determines the conditions for authorized access to the information and the conditions for individuals to become members of the group. This concept applies equally to United States national classified information, trade secrets, or proprietary data. Three elements are necessary for this concept to be employed:

- A group must have a defined membership.
- Information objects must be uniquely identified within the domain of the group.
- The security policy regarding the protection of and access to the information objects must be known and agreed to by the membership.

An *information domain* is a set of users, their information objects, and a security policy. An information domain security policy is the statement of the criteria for membership in an information domain and the required protection of the information objects. Information domains are not hierarchically related, nor may they implicitly or explicitly infer a sensitivity related to multiple categories of sensitivity.

In contrast to domains that might be composed of systems or networks, information domains are not bounded by systems or even networks of systems. Information domains are bounded by the presence of their identifiable information objects and may be supported by any system that can meet the protection requirements of the information domain security policy. In this concept, a specific mission security policy may define several information domains, each with its own distinct information domain security policy. The security mechanisms of any number of information systems may be evaluated for their ability to meet these information domain security policies. Through the process of accreditation, these security mechanisms may be usable for part or all of one or more missions.

Each information domain is identified uniquely. The unique identification indicates (directly or indirectly) the sensitivity of all the information objects in an information domain. Any security-relevant attributes of information objects in an information domain must be the same for all information objects in the information domain. That is, there must be no security-relevant distinction made among the information objects in an information domain. Members of an information domain may have different security-related attributes. For example, some members might have only read permission for information objects in an information domain, while other members might have read and write permissions. Since all information objects in an

information domain have the same security-relevant attributes, a user who has read and write permissions in an information domain has those permissions for *every* information object in the information domain.

4.3.2 Strict Isolation

The diversity of missions and the threats to the security of their information will result in information domain security policies with unrelated protection requirements. Thus, information systems that support multiple information domain security policies must adopt a protection strategy that provides a basis for satisfying all of them. One such strategy, termed *strict isolation*, is to isolate one information domain from another, except when there is an explicit relationship established. Under this strategy, an information system must provide mechanisms that maintain separation of information domains in ways that are satisfactory to each of them. The default information system security policy is strict isolation among the information domains supported.

In the absence of any information domain security policy to the contrary, an information object must be isolated. While such a situation is a logical possibility, in practice, all information objects should belong to an information domain that has a defined membership and an information domain security policy. Information domains with no explicit interdomain policies must adopt a policy of strict isolation to be enforced by the systems that support them.

4.3.3 Interdomain Information Sharing and Transfer

Some mission requirements will necessitate the sharing or transfer of information objects among information domains. The establishment of new mission functions, new mission area relationships, or new organizations are examples of events that can create requirements for information sharing and transfer.

The simplest method of sharing information is to accept new members into an existing information domain and to grant access privileges to them. Where a need exists to share some, but not all, of the information objects in one or more information domains with members of other information domains, a new information domain may be created to contain the shared information objects. The new information domain, like any other information domain, requires a security policy. The members of the new information domain may or may not be members of the information domains from which its information objects were obtained.

Information objects can be transferred between two information domains only in accordance with established rules, conditions, and procedures expressed in the security policy of each of them. The transfer can be accomplished only by a user who is a member of both the sending and receiving information domains and, if required by the information domain policies, has been granted the appropriate privileges (e.g., "release authority").

The transfer of information objects between information domains may be implemented as a move operation (in which the information object no longer exists in the originating information domain), or as a copy operation (in which the information object exists in both information domains). Information objects moved or copied from one information domain to another must be relabeled with the label of the information domain to which the information object has been moved or copied.

Interdomain transfers can only occur within an end system or relay system. Interdomain transfers cannot occur among distributed end systems or relay systems; transfers among end systems or relay systems can only occur within the same information domain. These restrictions are consequences of the nature of security contexts and security associations that are used to create an appropriate environment for distributed information domain operations (including the maintenance of strict isolation between information domains). These concepts are explored further in sections 5 and 7.

In the U.S. national classification system, downgrading or declassifying are examples of transfer of information from one information domain to another. This type of transfer is performed to reduce the sensitivity of information objects either to satisfy established policy on information perishability or in careful coordination with the classifying authority. Sanitization is the process of removing certain sensitive portions from an information object, thereby creating new information objects, with the unique identifier of the receiving information domain, that may be transferred to an information domain of lesser sensitivity with a different information domain identifier. Sanitization also requires careful coordination with the original classifying authority.

4.3.4 Multidomain Information Objects and Policies

The missions of most organizations require that their members operate in more than one information domain. The information management activities of a mission may be viewed as taking place in a set of information domains, some of which may be shared with other missions. To carry out their mission information management activities, users may need to process information objects from several information

domains concurrently. Often, a user may have a *perception* that a collection of information objects from different information domains is a single, composite information object. Such a composite information object is referred to as a *multidomain information object*. This perception must be achieved without actually combining real information objects from different information domains to create real multidomain information objects. When creating the perception of multidomain information objects, strict isolation among information domains must be maintained, and the constituent information objects within the multidomain information object must be managed only in accordance with their individual information domain security policies. The purpose of multidomain information objects is to be able to define a collection of information objects to be displayed, printed, or transferred between information systems in a particular order or arrangement.

The creation and use of multidomain information objects must be subject to some security policy. The simplest policy is the one noted above, namely to conform to the policies of the individual information domains. However, in such cases it may not always be possible to print such a multidomain information object or to convey it to another user or information system. A multidomain information object security policy might be based upon some existing policy (e.g., U.S. national security policy) that states a relationship among the constituent information objects. Such a security policy for multidomain information objects is made part of the security policy of the information domains of the constituent information objects. In situations where the security policy for multidomain information objects is complex or involves several information domains, that security policy might be stated in one place in the supporting information system and be referred to by the individual information domain security policies.

Explicit multidomain information object security policies must state the specific privileges a user must have to view, print, create, delete, or transfer a multidomain information object between information systems. To create or otherwise deal with an entire multidomain information object, the user must be a member of each of the information domains in which the constituent parts of the multidomain information object are located. Some multidomain information object security policies might allow access only to the component parts of a multidomain information object for which the user has appropriate privileges, but in many cases this would not result in a sensible multidomain information object.

As noted in section 4.3.1, information domains are not hierarchically related. Nonetheless, security policies for multidomain information objects may recognize marking rules that apply to the entire multidomain object or its parts based on existing policies (such as paragraph and page markings for information subject to

U.S. national classification policy) when printed or displayed. Further, an information domain security policy is not precluded from recognizing that a user security clearance of Top Secret is adequate for access to the information objects in an information domain that contains U.S. national classification Secret information objects, if all other aspects of the information domain security policy are also met. (Note that the apparent hierarchy among U.S. national security policy classifications is actually a property of user privileges, in the form of clearances, rather than a relationship imposed on information of different classifications. Information that is classified Secret is not a subset of information that is classified Top Secret.)

The implementation of multidomain information objects in real information systems has many implications for end system, security management, and transfer system architectures. These implications are discussed further in sections 5, 6, and 7.

4.3.5 Absolute Protection

Since open systems may consist of an unbounded number of unknown heterogeneous LSEs and it may be necessary to communicate with any of them, system security architects must have a rational basis for protection decisions in such an environment. In this environment, it is no longer possible to rely upon the assurances provided by physically separated networks or cryptographically isolated LSEs. Information domains must rely on the protections afforded by a heterogeneous collection of LSEs. The concept of *absolute protection* (which does not imply perfect protection) is set forth to provide a framework for achieving uniformity of protection in all information systems supporting a particular information domain. It directs its attention to the problems created by the interconnection of LSEs that provide disparate strengths of security protection.

In order to support an information domain in multiple LSEs, the overall strength of protection afforded to information objects must be consistent in those LSEs. Strength of protection is a function of the strength and correctness of security mechanisms (including doctrinal mechanisms) implemented in LSEs to satisfy an information domain security policy. The required strength of protection is determined by assessing the value of the information being protected and then assuming a hostile attacker has logical access to the LSE through the transfer system. *The specific mechanisms and their implementations need not be identical in every LSE that supports an information domain, but the implementations must provide at least the required strength of protection.*

If the overall strength of protection provided by each LSE supporting an information domain is successfully evaluated under the assumption that the LSE is logically

accessible to a hostile user, then each of these LSEs can be accredited as being adequate to protect the information domain against the same threats. Protection provided in all the accredited LSEs under these conditions will be absolute, non-relative, and equivalent. Absolute protection assumes minimum, equivalent environmental protection, but is primarily concerned with the vulnerabilities created by connections to communications networks. This concept generally forces stronger mechanisms to be employed for information of a given sensitivity.

For system security architects, implementors, and accreditors to properly apply the concept of absolute protection, different approaches to evaluation of security mechanisms, components, and information systems will be required to determine equivalent protection. A single measure of overall strength of protection is not adequate. Rather, security mechanisms will need to be rated (measured) for their ability to support one or more security services, alone and in combination with other security mechanisms. The required strength of protection for an information domain will be translated to a set of such measures so that an appropriate set of security mechanisms can be chosen. This method of choosing security mechanisms will give security architects, implementors, and accreditors a consistent means for providing equivalent (though not necessarily identical) protection in the LSEs that support an information domain.

4.3.6 Uniform Accreditation

The concepts of information domains, strict isolation, and absolute protection work together to provide the basis for achieving a high degree of uniformity in accreditation. Each LSE is evaluated against the security policy of each information domain that it supports. The objective is to have equivalent protection in all LSEs that support a given information domain. Strict isolation establishes basic independence among all information domains supported by an LSE. This independence allows security services to be provided by different security mechanisms (of different strengths) in each information domain supported by an LSE.

Absolute protection ensures that security services required by an information domain security policy are implemented appropriately and under equivalent threat environments in all supporting LSEs. An information domain accreditor receives an evaluation for each supporting LSE. The evaluation is unaffected by other information domains supported by the same LSEs or by variations in environment. Uniformity in accreditation remains subjective to the degree that evaluators are consistent in their judgments of the effectiveness of security mechanisms and the assumption of risk. Limiting the number of evaluators and accreditors can reduce this subjective variation.

4.3.7 Security Management

Security management is a particular instance of information system management. Security management is central to the proper operation of protected LSEs and their component parts, both separately and jointly. Within the DGSA, security management performs many critical functions in support of its requirements, security service allocations, and security concepts.

Security management is concerned with all aspects of protection of LSEs, protection within LSEs, and protection among LSEs. For example, security management includes the control of physical access to facilities, the maintenance of information domain security policies, the invocation of particular security mechanisms in end systems, and the control and protection of communications between LSEs.

This section introduces terminology and basic security management concepts that are presented in more detail in sections 6 and 7. Generally, the OSI *Management Framework*, ISO 7498-4 (ISO, 1989b), and the security management portion (clause 8) of ISO 7498-2 are adopted for use in the DGSA.

Managed objects are information system resources that may be managed. *Management information* is information associated with a managed object that is operated upon to manage that object. A human administrator employs a *management application processes* (MAP) to use and maintain management information contained in a logical repository called a *management information base* (MIB). (A MAP includes the system management application-entity defined in ISO 7498-4, which is limited to systems management communication.) The contents of a single logical MIB may exist in several LSEs. When it is necessary to refer specifically to the processes and management information for security management, the terms *security MAP* (SMAP) and *security MIB* (SMIB) will be used. Otherwise, statements applying to MAPs and MIBs are understood to apply to SMAPs and SMIBs as well.

To ensure efficient and flexible system management, it is generally required that administrators have local or remote access to MIBs. As a result, MAPs will exist in all LSEs. CNs may also contain MAPs and MIBs associated with their management. LSEs will manage their LCSs and also may need to cooperate with CN management. In most instances, this cooperation will not involve the use of security-related information.

Since management information comprise specially designated sets of information objects, these sets must exist within an information domain. Several possible choices

can be made concerning the information domain in which particular management information objects exist relative to the information domain being managed:

- Each information domain may have a corresponding management information domain (1:1).

- A single management information domain may contain the management information objects for several information domains (1:many).

- The management information objects may be part of the information domain (embedded).

The first two choices are appropriate when the SMIB should not be contained in the information domain to be managed. The last choice, in which the MIB is a part of the information domain being managed, implies that every member of the information domain has the same access privileges to the MIB as any other information objects.

In addition, some management information objects may be associated with an entire information system and its functions. The system MIB might exist in its own management information domain or it might be placed in another management information domain (the latter situation is most likely when a "1:many" management information domains used).

Security management information objects and security management functions are associated with overall end system security management or with information domain security management. Both are concerned with:

- Security policy management
- Security service management
- Security mechanism management (including doctrine)
- Security mechanism support management (e.g., key distribution and management)
- Transfer system security management (security association management)

The ability to manage information systems and information domains effectively will depend on the availability of appropriate management tools. These tools should include not just MAPs that provide the necessary functions, but also convenient user interfaces to those MAPs, tools for expressing information domain security policies as management information objects that can be interpreted by information system functions, and display tools that allow convenient monitoring of system functions,

among others. Standard protocols for the remote invocation of management functions and for cooperative management among LSEs are emerging, but they need considerable further attention.

(This page is intentionally blank.)

5.0 END SYSTEMS AND RELAY SYSTEMS

A generic security architecture for end systems and relay systems must be appropriate for a wide range of applications and environments. Among the many possible implementations, some unifying structure must be created that permits a generic approach to security. This structure must accommodate the requirements of section 2 and the primary security allocations made in section 4. This section refines several concepts presented in earlier sections for end system and relay system architectures, including security allocations, types of functions that are required to support the security allocations, types of devices that make up end systems and relay systems, and technologies that should be considered in specific implementations. Section 5.1 gives an overview of the end system security architecture, and its description is presented in section 5.2. Section 5.3 lists candidate technologies to support implementations. Generally, relay systems provide services that require the same kinds of underlying support as end systems, except that they do not provide support for direct user interactions. Thus, a single security architecture for end systems and relays systems is appropriate. The remainder of this section refers to both end systems and relay systems simply as *end systems*.

Since the DGSA is a generic architecture, not all of its possible architectural choices and alternatives (security services and mechanisms) will be used in every specific implementation. The DGSA allows for a wide variety of specific implementations that will be dictated by missions and threats. Similarly, the end system security architecture must have wide applicability. The end system security architecture described here is a current best estimate of how the DGSA requirements can be met. To the extent that it depends on specific technological directions, it is subject to change as experience and technology dictate. However, the basic architectural decisions described should remain stable.

Much of the end system security architecture is similar to that proposed by Rushby (1984). There are some significant departures from Rushby's proposal, most notably with respect to centralization of security policy-related functions. Rushby argues for such functions to be tailored to and to be implemented with specific resource management functions. This argument is implicitly based on the fact that only a single, access control-based security policy is to be enforced. The DGSA requirement for supporting differing security policies per information domain (which may have other dimensions than simply access control) makes the argument for centralizing the basic security policy-related functions more attractive. More recent proposals for support of multiple security policies suggest architectural approaches which take a middle ground and may offer some performance advantages (see Abrams (1993) for a summary and extension of these approaches).

The end system security architecture focuses on conventional computer systems, which represent a large portion of all end systems. Other end system types may need to implement only portions of the end system security architecture. In extreme cases, such as simple sensor devices, the end system functions may be so limited that only specialized implementations of a small portion of the end system security architecture are appropriate (for example, such a device almost certainly would not need to support multiple information domain security policies). The considerations cited in section 3.3 must be accommodated.

5.1 END SYSTEM SECURITY ARCHITECTURE OVERVIEW

In section 4, fundamental allocations of security services were made to LSEs and to the end systems and LCSs within LSEs. Security service allocations were made to LSEs to protect their resources, including end systems. The end system security architecture makes additional security service allocations to the end system hardware and software. Not every security service allocation needs to be made identically in every system. For example, if electronic emanations are considered to constitute a potential vulnerability, the responsibility for countering it could be assigned to the LSE or to one or more of its components. Similarly, there is flexibility with regard to how protection responsibilities are shared between end system hardware and software.

5.1.1 The LSE Protects the Hardware

As discussed in section 4, the security service allocations to the LSE are implemented as doctrinal mechanisms (i.e., physical, administrative, and personnel). Doctrinal mechanisms are discussed in more detail in section 8. The primary security service allocations to the LSE are access control to facilities and some aspects of authentication of personnel. In addition, some aspects of information confidentiality and integrity, and system integrity and availability may be allocated to the LSE.

5.1.2 The Hardware Protects the Software

Section 4 assigned responsibilities to the end system for all security services. There are a variety of security mechanism choices available between the hardware and software portions of the end system, but certain general allocations and properties can be stated for the hardware.

The hardware is relied upon to function correctly, to enforce isolation of software functions, and to contribute to the protection of the integrity of the system applications and the operating system. It provides protected paths between users

and trusted parts of the software. The hardware indirectly supports the isolation of information processed and stored in the end system by protecting the integrity of the software security. Hardware mechanisms are used to protect the system from radio frequency interference and to prevent undesired emanations. In some environments, specific hardware technologies (e.g., protective coatings, hardened or alarmed containers) may be necessary to protect against tampering with end system components. Availability of an end system may be enhanced through technologies such as fault-tolerant and fault-detecting hardware features. Hardware cryptographic mechanisms are employed as needed to support various security services. Other hardware mechanisms (e.g., memory mapping) support specific aspects of the software architecture and are noted in the detailed end system security architecture discussion (section 5.2). There is an array of equipment available to support the hardware allocations.

5.1.3 The Software Protects Information

The security service allocations made to software are wide ranging. The portion of the transfer system supported by the end system software is responsible for the confidentiality and integrity of information transferred among end systems, for the authentication of end systems to one another, and for user authentication and access control in distributed systems. The details of how the transfer system is supported by end systems are presented in section 7.

Security services and the mechanisms that implement them must be managed. The software applications that support security management in end systems are discussed in section 6 and are extended in section 7 for transfer system support.

The end system software is responsible for user authentication and access control, and for the integrity of information being processed and in storage. Correct operation of certain software is required to ensure end system availability. Additionally, the software is expected to provide functions that support the security policies and requirements stated in section 2 that are not directly expressed as security services, such as support for multiple security policies. The remainder of this section refines the end system security architecture, which primarily is concerned with software structure.

5.2 END SYSTEM SECURITY ARCHITECTURE DESCRIPTION

A generic end system security architecture must respond to the security allocations discussed earlier, and it must be sufficiently flexible to encompass changing technology (of end systems and security mechanisms). The end system security

architecture presented here is not an implementation specification, and it might be realized in several ways. The end system security architecture concentrates on support for multiple information domains with distinct security policies. Attention is paid to strict isolation of information domains, management of end system resources, and controlled sharing and transfer of information among information domains. The end system security architecture also relies upon an engineering approach that seeks to isolate security-critical functions into relatively small modules that are related in well-defined ways. This approach has advantages in implementation, certification, and accreditation by limiting the scope of particular portions of these activities. While there are no existing end systems that specifically implement all of the end system security architecture, several efforts have been documented in the academic and research communities that support various aspects of the end system security architecture. Recently, commercial operating system vendors have adopted design and implementation strategies that share significant aspects of the end system security architecture.

A *security context* is a combination of all the LSE, hardware, system software, user application software, and information supporting the activities of a user (or system function) operating in an information domain. A security context builds on the common operating system notion of a user process space (sometimes called a *context*) as supported by hardware features and operating system functions. The primary distinctions between an ordinary user process space and a security context are that aspects of protection provided by the LSE are explicitly included, and that user applications operate in a controlled process space *subject to an information domain security policy*. Security contexts are described in more detail in section 5.2.2.

A *separation kernel* manipulates the protection features of the end system hardware (e.g., processor state registers, memory mapping registers) to maintain strict separation among security contexts by creating separate address spaces for each of them. A separation kernel also controls communications among security contexts to allow sharing or transfer of information, and to allow services to be performed by one security context for another. All user security contexts and many system function security contexts are constrained to make requests for basic end system services on the separation kernel through a *standard kernel interface*. The separation kernel is described further in section 5.2.1. The functions that make and enforce security policy decisions are intimately related to the separation kernel. These are described in sections 5.2.1 and 5.2.3.1.

Figure 5-1 provides an abstract view of the end system security architecture (care must be taken not to make assumptions about implementation based on this figure).

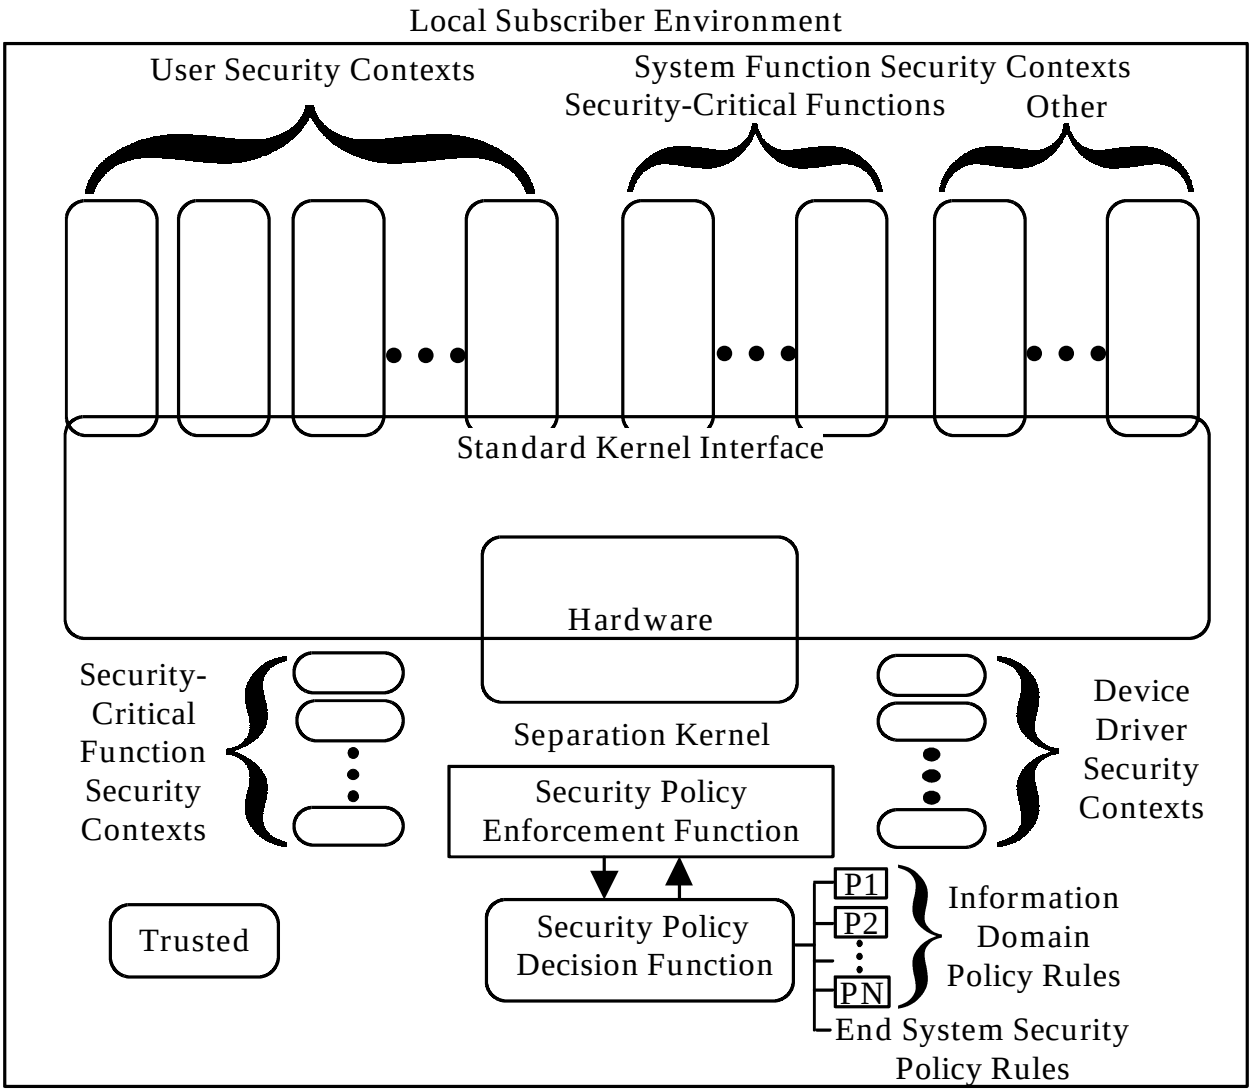

Figure 5-1. End System Security Architecture Abstract View

End system software is divided into trusted (shown in the shaded area) and untrusted parts for practical evaluation. The trusted parts of the software are those that are considered so important to the secure operation of the end system that they must undergo strict evaluation procedures and come under strict configuration management control. The hardware (including any microcode) is considered trusted in the sense that its operation is assumed to be correct. Untrusted software is able to perform operations on basic system resources only through invocations of security-critical functions that are mediated by the separation kernel; inter-security context

operations (e.g., inter-information domain communications) are performed by security-critical functions.

Untrusted *security-related* functions (such as security management applications and portions of transfer system applications) are expected to operate correctly to satisfy user operational needs, but need not be subjected to the rigorous scrutiny applied to the security-critical functions. Security-related software is not assumed to be free of security defects, although it is certainly prudent to obtain such software from reliable sources, test it before use, apply integrity safeguards to ensure it remains unchanged, and apply configuration management to it. (Software obtained from less than reliable sources may need to be inspected more carefully.) Under these conditions, if faulty application software is introduced into a system it will, at worst, prevent certain operations, but information compromise will not result because

of the combination of strict isolation of information domains enforced by the end system, testing, and configuration management. The remaining software is not only untrusted, but is not expected to be examined for any security reasons.

The following subsections provide additional detail on the end system security software components, primarily for the separation kernel, security contexts, security-critical functions, and operating system implementations.

5.2.1 Separation Kernel

Much general operating system research has concentrated on organizing basic operating system functions into a collection called a *kernel*. The kernel presents abstractions of the fundamental resource management mechanisms to other, less primitive, service providers (information system functions and applications). In operating system implementations that attempt to provide a basis for secure information processing, the kernel software is carefully constructed and evaluated. To aid the evaluation process, the kernel functions are implemented as relatively small programs that are independent of one another to the maximum extent possible.

Rushby suggested that significant improvements in secure operating system kernel design and implementation could be achieved by isolating each kernel function in its own process space (i.e., address space). The benefit of this approach is that each operating system function performs a single, well-defined activity and can be understood and evaluated in relative isolation from all other functions. A separation kernel is charged with the critical task of providing separation among process spaces by manipulating the protection features of the end system hardware.

Until recently, most secure operating system designs have been limited with regard to security policy specification and enforcement. Particular limitations include support for only a single security policy (usually an access control policy) and the inability to change security policy conveniently. The DGSA end system security architecture adopts a particular view of operating system kernel design to meet certain of its requirements and concepts, most notably the support of multiple security policies so that a single end system can support users in different information domains simultaneously. The traditional operating system kernel functions are divided among the separation kernel, security policy enforcement and decision functions, and the remainder of the trusted operating system functions, called the *security-critical* functions. The separation kernel serves as the ultimate security policy enforcement function by mediating all use of the basic information system resources. The separation kernel notion is the foundation of the end system security architecture. However, any other information system mechanism that provides equivalent isolation of information domains and control of system resources is appropriate for implementations that conform to the DGSA.

The end system security architecture generalizes a principle that is becoming widely accepted concerning access control, namely the independence between the decision of whether or not an access to a resource is allowed and the enforcement of that decision. The separation of access control decision-making and access control enforcement functions allows the support of multiple access control policies. The ISO Access Control Framework (ISO, 1992b) designates these functions the *access control decision function* (ADF) and the *access control enforcement function* (AEF), respectively. In fact, most existing secure operating system designs have concerned themselves only with access control policy. Since one of the DGSA requirements is to support any security policy, the end system architecture extends the AEF concept to include the enforcement of all aspects of an information domain security policy. The resulting function is called the *security policy enforcement function* (SPEF). Similarly, the ADF concept is extended to a *security policy decision function* (SPDF). (The SPDF is discussed in more detail in section 5.2.3.1.) The separation kernel is the implementation of the SPEF in the end system security architecture.

The separation kernel also is an extension (beyond access control) of the *reference validation mechanism* (RVM) described in the Trusted Computer System Evaluation Criteria (Department of Defense, 1985). The basic properties of the RVM must be applied to any separation kernel implementation: it must be invoked for every security-critical operation, it must be small enough to be verified, and its integrity must be maintained.

In the spirit of several current standardization efforts, a standard kernel interface will be defined to allow open system development of operating systems and applications built on implementations of the DGSA end system security architecture. The standard interface to the separation kernel is the same whether the underlying computer is a large multiprocessor mainframe or a single-processor workstation. This approach allows developers great latitude in implementing the separation kernel and the security-critical functions.

5.2.2 Security Contexts

From the perspective of the separation kernel, a security context is defined by a set of data and programs operating in accordance with an information domain security policy. As noted earlier, a security context also includes the doctrinal mechanisms of the LSE, the hardware-based resources (e.g., registers, memory, disks) that are in use when the end system is serving a particular user (or system function). That is, a security context encompasses all end system resources and security mechanisms (including doctrine) that support the activity of a user operating in an information domain. The separation kernel must maintain all the information needed to isolate one security context from another. When the end system ceases performing operations in one security context and begins performing operations in another security context, no information can be allowed to pass from one security context to the other unless a specific request is made and it is allowable under the security policies of the information domains involved.

Examples of information that end system security-critical functions (including the separation kernel) must maintain to support the operation and isolation of security contexts include:

- A unique identification for each security context
- The identification of the information domain being supported
- Hardware register values related to control of end system resources, including virtual memory and all devices in or attached to the end system
- The authenticated identity of the user being served
- The user's security attributes (permissions)
- Data structures needed to operate security-related functions and other untrusted system applications

Each security context supports a user (or a system function) operating in a particular information domain. Over a period of time, an end system may maintain several security contexts to support one or more users operating in one or more information domains. A particular user might use (simultaneously or serially) security contexts operating in the same or different information domains. Different users may employ security contexts operating in the same or different information domains.

Since security contexts are isolated from one another by the separation kernel, communications among security contexts (requests for service or information transfer) in an end system can only take place in accordance with the security policies of the information domains supported by the security contexts. If the security policies of the supported information domains do not explicitly permit inter-information domain transfer, the SPDF will necessarily deny the request and the separation kernel will enforce that decision. Since an information domain contains the information of a particular user community, it would be unusual for an information domain security policy to prohibit information sharing between two security contexts supporting the same information domain.

Many end system activities are not carried out on behalf of a specific user (either an individual or the entire membership of an information domain as a group), but rather for basic end system operation and management. Examples of such activities include many of the security-critical system functions and end system management activities. These activities are carried out within end system security contexts on behalf of one or more of the information domains supported by the end system. The security policies of these end system information domains are created to exercise appropriate control of end system resources for all of the user information domains supported by the end system. Some example uses of end system information domains include the control and manipulation of multidomain objects, login applications, and management information domains.

Multidomain information objects (see section 4.3.4) never exist in an end system except as displayed (or printed). Nonetheless, in end system implementations, it must be possible for a user to describe the relationships among the components of a multidomain information object so it can be displayed. Some implementations of multidomain information objects will result in the description being represented as an information object. Some security policies may preclude this information object from being held in any of the component information domains. In such cases, the end system must be able to create a system security context in which the description can be used by an appropriate application program that requests the display manager to construct the multidomain information object on a display device. Note that the multidomain information object description could be retained by the end system for

future use by either the creator of the description or by other users who have the necessary information domain memberships. Similarly, the description could be transferred, in accordance with a multidomain object policy, (separately or with the component information objects) to another end system (see section 7).

Before a security context can be created for the activities of a user in a particular information domain, the system must be informed which information domain is to be used. Ordinarily, the user's identity must be obtained and authenticated to determine if the user is a member of the requested information domain. One way of performing this startup function is to create a "login" security context that represents one of the end system security domains. The activities allowed in the login security context are limited to authenticating the user identity and starting a security context for the requested information domain (there might be a default information domain for a user recorded in the end system SMIB).

One useful resource control concept is *type enforcement*. The type enforcement concept generally restricts the input and output of a particular function to be of delineated types. In turn, the functions that are allowed to invoke other functions can be controlled by careful specification of input and output types. It is possible to impose a particular implementation of type enforcement by making specific security-critical functions "members" of particular end system information domains. Thus, only "member functions" of an end system information domain could invoke specific executable end system functions.

A consequence of the strict isolation aspects of the end system architecture is that many aspects of covert channels, both timing and storage, either cease to be concerns or are easily controlled. Possible storage channels are reduced to those between security contexts. If information domain policies are properly stated and the security policy, strict isolation, and interprocess communication functions are performing properly, there will be no covert storage channels available. To exploit timing channels between security contexts requires that a complete security context list is available so that a user can determine which security contexts (including end system security contexts) are in operation. Such information is part of one or more management information domains. It is not likely, and certainly not necessary, that an arbitrary user would be able to access such information. Even for those security contexts in which management information is available to its users, timing information for other security contexts should not be made available to those users.

5.2.3 Security-Critical Functions

The security-critical functions described in this section implement the various security services allocated to the end system and several additional supporting services. This set of security-critical functions is not necessarily complete as presented. Experience through prototyping and experimentation is needed to guide implementations that will meet all of the DGSA requirements, but the functions presented below should provide a sufficient basis for further research.

5.2.3.1 Security Policy Decision Function

The separation of security mechanisms from security policy enforcement and decisions is crucial to the flexibility of the end system security architecture. The SPDF is responsible for making all security policy decisions. The primary role of the SPDF is to isolate the rest of the end system software from knowledge of security policies. The importance of this approach is threefold.

First, the support of multiple information domains with different policies is accomplished easily because the security policies are represented in only one place and are interpreted by only one function. In many current secure system designs, it is difficult to point to the actual software code that implements the single security policy of those systems because it is embedded and scattered throughout code that performs multiple functions.

Second, by keeping security policy representations in one place, it is relatively easy to install, modify, or even replace the security policy for an information domain. It is not necessary to rewrite trusted software that implements the security policy. Rather, the rules that the SPDF interprets for a domain are updated or replaced.

Third, changing the implementation of the SPDF would be transparent to the operation of the remainder of the end system software. Any correct implementation of the SPDF is acceptable, but it may be useful to standardize the representation of security attributes and security policy rules.

The SPDF approach will allow security-critical functions to be implemented independently of particular security policies. There is the potential in this approach that a computer vendor could support its entire customer base within a single end system software design. To illustrate this concept, consider an example of three enterprises with different, or even conflicting, security policies. The first is a DoD organization using a conventional DoD security policy. The second is a corporation with requirements for data integrity and data separation based solely on need-to-

know authorization. The third is a university research laboratory that does not have any special security needs except a basic privacy-based access control policy. Without a policy-independent architecture, these three differing security policies would result in three different operating system implementations that could cause serious compatibility problems for a vendor trying to support all three environments. Using the SPDF approach, any or all of the three policies could be supported by the same end system software. If necessary, the three enterprises could be served by the same end system or (using the transfer system) they could share information as necessary across different end systems.

5.2.3.2 Authentication Function

The authentication function invokes one or more mechanisms used by an end system to identify and authenticate users (and to authenticate an end system to users), and for end systems to authenticate one another in a distributed environment. A common interface to the authentication function is used that is independent of the any information domain security policy or the authentication mechanisms employed. That is, the authentication function is the service interface to the mechanisms used to identify and authenticate users and end systems. The exact mechanisms selected will depend on the information domain policies in effect. An end system supporting multiple information domains may need to implement more than one authentication mechanism.

An authenticated user identity may be passed between information systems rather than the information used to authenticate that identity. That is, an end system supporting a particular information domain would be expected to accept that the authentication function has been performed reliably and correctly by other end systems supporting that information domain (use of the absolute protection concept makes this assumption reasonable). In some cases, it may be necessary to pass information about the authentication mechanisms used to validate the user identity. The transfer system is expected to protect the authenticated user identity as it is passed between information domains. Additional detail about distributed end system interactions is given in section 7.

5.2.3.3 Audit Function

The audit function accepts audit messages from functions in the end system in accord with information domain and management information domain security policies. Audit records may become part of a SMIB that is part of an information management domain (for one or more information domains or end system domains). Audit records may be directed to multiple repositories. In some cases, the audit

information may best be used by an individual user (for example, time and method of most recent end system or information domain use). The audit function guarantees that audit messages cannot be lost and that the ordering of messages is preserved. As part of a distributed audit system, audit functions can forward the audit data they collect to a base-level, regional, or central audit center to alleviate local audit data storage requirements and to coordinate audit information from different end systems or LSEs. Audit data must be protected from unauthorized access or modification.

5.2.3.4 Process Scheduling Function

In operating systems that share the end system processor among multiple processes, the process scheduling function determines which of the processes next uses the processor (or processors in a multiprocessor end system) and for how long. The process scheduling function must be included among the security-critical functions so that no process can deny the processor to other processes either purposefully or inadvertently.

5.2.3.5 Device Management Functions and Device Controllers

The remainder of the security-critical functions are each responsible for a particular class of end system resources described below. These resources include memory, storage devices, display systems, interprocess communications, cryptographic services, and any other input/output devices controlled by the end system.

- The *memory management function* is responsible for controlling the use of memory by all software, including security-critical functions. It maintains memory-mapping information and controls the hardware functions that perform memory mapping.

- The *file management function* is responsible for controlling the use of storage media devices. Like the memory management function, it maintains disk-mapping (or other media-specific) information that provides basic virtualizations of the actual storage media. Other software (e.g., database programs) may build upon these virtualizations to provide even more abstract file structures to applications and users.

- The *display management function* is responsible for controlling the use of display devices (including screens and printers), keyboard devices, and pointing devices (e.g., trackballs, mice). The display management function provides basic display device operations. Because a single display device

may be used to present information from multiple domains at the same time (typically through multiple windows or on paper), the display management function must maintain information that associates particular information to be displayed with the appropriate security context. Other software (e.g., an X Window System implementation) may provide requests to the display management function to achieve a particular display format.

- The *interprocess communication management function* is responsible for controlling the interprocess communication mechanisms (e.g., locks, semaphores, messages) used by all software processes in the end system. In particular, inter-context (e.g., inter-information domain) transfers are carried out through this function.

- The *cryptographic services management function* is responsible for controlling all of the cryptographically based security mechanisms in an end system. The security services it may support include confidentiality, data integrity, data origin authentication, and non-repudiation. The cryptographic management function may control a number of alternative cryptographic mechanisms to support different services and to provide different levels of protection that satisfy different security policies. The choice of mechanism may be based on many factors including the sensitivity of the data being protected, the security service requested, and the mechanisms available on other end systems for data that will be transferred.

- Each of the physical devices in the end system, including memory, disks and other storage devices, displays, cryptographic engines, specific user authentication devices, and communications interface controllers, has a corresponding software program that controls and passes information to and from it. These software programs collectively are called *device drivers*. Every device driver must be considered security critical because this software ultimately determines how a device operates. Although device drivers in older end system platforms were often quite large and complex, many contemporary devices contain much of the former device driver function in the device logic or in their own programs. Thus, many device drivers are now reasonably straightforward and follow well-known paradigms, which make their evaluation easier, although great reliance is placed on the correct implementation of the device.

5.2.4 Security-Relevant Functions

Some software functions within the end system are required to manage information or to provide an interface to the security-critical functions, but are not critical to system security. Of particular interest here are residual operating system functions, security management functions, and transfer system functions.

5.2.4.1 Residual Operating System Structure

Most of the security-critical functions are part of traditional operating system structures. Many other operating system components are not included in the security-critical functions, such as the user interface, utility functions, and high-level abstractions of information. These functions are present in varying forms in all traditional operating systems. The user interface (or shell), the particular utility functions, and the information abstractions provided characterize a particular operating system. That is, they distinguish one operating system from another even though they provide essentially the same services to a user. Because the security-critical functions provide commonly used, low-level services, many different operating systems can be implemented using them. Figure 5-2 is an abstract illustration of the software supporting a single security context.

Since security contexts are separated from one another, each can rely upon a different residual operating system structure. Thus, a single end system can support different operating system environments concurrently. Applications that were written to operate with a particular operating system should not require change unless they were allowed to directly manipulate basic operating system functions now controlled by security-critical functions .

Existing operating system implementations will need to be modified to use the standard kernel interface and the services provided by the security-critical functions. The degree of difficulty in making these modifications will be reduced if the original operating system implementation was well structured and modular. Some existing secure operating system implementations will adapt relatively easily to the use of the standard kernel interface, and many of the security-critical functions will already be present.

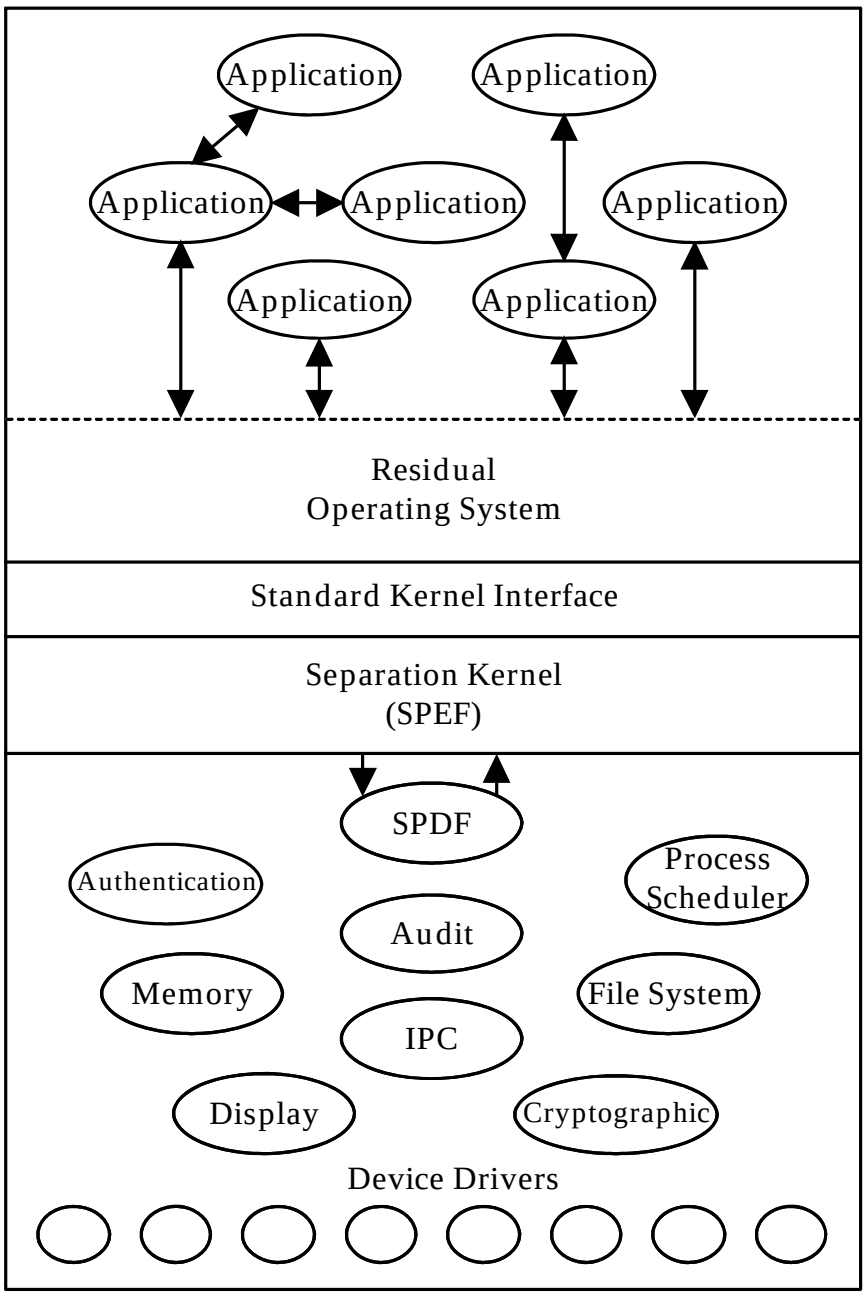

Figure 5-2. Security Context Software Abstract Relationships

Residual operating system implementations structured to use the standard kernel interface to obtain basic services should be able to be moved among different

hardware bases relatively easily since most hardware dependencies will be visible only in the separation kernel and the device drivers. This technique will enable applications to be used even in the face of changing hardware systems.

It should be noted that existing trusted software subsystems (e.g., trusted database applications) also will need to be restructured to fit the end system security architecture. It is possible that such a subsystem might be written to make direct use of the standard kernel interface (rather than calling on the residual operating system) for reasons of efficiency.

5.2.4.2 Security Management Function

The primary role of the security management function is to control information needed by security-critical and security-related functions within the end system security architecture. Security management is a particular instance of general management functions. The concepts and structures defined in ISO 7498-2 and ISO 7498-4, have been adopted for use in the DGSA. Examples of the information manipulated by the security management function include information domain security policy rules used by the SPDF, configuration parameters for security mechanisms (e.g., cryptographic algorithms), configuration parameters for cryptographic mechanisms and end system devices, and audit information. Some information is managed for specific information domains and some is managed for end systems or LSEs. Details on security management are contained in section 6.

5.2.4.3 Transfer System Function

The transfer system is defined in accordance with ISO 7498 and ISO 7498-2. Communication applications (e.g., X.400 electronic mail (CCITT, 1988), X.500 directory services (CCITT, 1992), file transfer) and communications protocols used to communicate with other end systems are implemented as untrusted applications within the end system security architecture. These applications make requests for security services (which process information and generate protocol information) that provide required protection. For information to be transferred between end systems and within an information domain, a distributed security context is established through the use of security management and transfer system applications, and security-critical functions. Details of the transfer system are presented in section 7.

5.3 END SYSTEM SECURITY ARCHITECTURE TECHNOLOGIES

Technologies that must be considered to bring about the end system security architecture affect all of the elements identified in section 5.1 (local subscriber environment, hardware, and software).

5.3.1 LSE

The allocation of security services to the LSE requires that mechanisms must be in place to support those services. Physical, personnel, and administrative doctrinal mechanisms will be used to implement the LSE protections. Some areas and mechanisms identified for additional investigation and research are listed below:

- Biophysical - authentication, physical access control
- Electronic - physical access control

5.3.2 Hardware

The allocation of security services to the hardware requires that mechanisms must be in place to support those services. Some areas and mechanisms for additional investigation and research, and the security services they support, are listed below:

- Fault Tolerance - availability
- Fault Detection - availability, integrity
- Memory Management - strict isolation, integrity
- Protected Mode/Multistate Processors - strict isolation, integrity
- Majority Logic - availability, integrity
- Multiprocessor Architectures - availability, strict isolation, integrity
- TEMPEST - confidentiality
- QUADRANT - availability, integrity

5.3.3 Software

The allocation of security services and other security-critical functions to the software requires that mechanisms must be in place to support those services. Some areas and mechanisms identified for additional investigation and research are listed below:

- Separation kernels - strict isolation and access control

  - Separation kernel interfaces
  - Process subsystems

-- Interprocess communications
-- Buffer caches
-- Security policy enforcement functions

- Security-critical functions - authentication, confidentiality, integrity, access control, non-repudiation

  - Security policy decision functions
  - Audit functions
  - Cryptographic engine functions
  - Device drivers
  - Window managers

- Security-relevant functions

  - Security management functions
  - Transfer system functions

- Trusted applications

  - Databases
  - X Window System
  - Operating systems

(This page is intentionally blank.)

6.0 SECURITY MANAGEMENT

Security management provides supporting services that contribute to the protection of information and resources in open systems in accordance with applicable information domain and information system security policies. Security policies, security services, and security mechanisms are the basic classes of managed objects. This section builds on the definitions and concepts presented in section 4.3.7. In section 6.1, critical aspects of security management are related to architectural elements and concepts of the DGSA. In section 6.2, clause 8 of ISO 7498-2 is used as the basis for presenting details of the DGSA security management architecture. Section 6.3 identifies tools needed by security architects and security administrators, and section 6.4 discusses standards needed to support DGSA security management.

6.1 SECURITY MANAGEMENT RELATIONSHIPS TO DGSA CONCEPTS

The requirement to manage multiple information domains is the most significant impact on traditional approaches to security management. Traditional security management is based on the assumption that all users of an end system are subject to the same security policy, so that a single view of security management is sufficient for the entire end system. End systems that support multiple information domains must provide independent security management for each of the information domains. In addition, the use of security services and security mechanisms shared among multiple information domains requires security management coordination at the end system level. Thus, an end system security policy is necessary to specify how the shared use of security functions and resources among information domains is accomplished. This end system policy also must be managed.

As a result of this focus on security policy management, DGSA security management is mission driven and information oriented because information domains are the reflection of mission decisions on how to organize and control information. Section 2 discussed the relationships among missions, requirements, security policies, and security architectures, but only to the granularity of the entire mission. Information domains typically will reflect a major mission function, so further refinement of the mission-specific security policy into an information domain security policy will be necessary. It is not appropriate to specify how that refinement should be done since no rules exist for creating an information domain. However, a number of elements of information domain and end system security policies will be typical for a wide range of mission functions. Several of these security management elements of security policy are listed below, but the lists are not all-inclusive. Section 9 includes examples of incorporating security management policy elements into information domain and end system security policies.

A typical information domain security policy might include some or all of the types of information listed below. Not all of these information types will be reflected in the information domain security policy rules interpreted by the SPDF in an end system, but they are necessary to the development of those rules. Security management in end systems is concerned with the installation, maintenance, and enforcement of these rules and the information about users, security services, and security mechanisms needed to achieve a security policy. Not all security management activities are performed in end systems and relay systems. There are always supporting security management activities that are related to doctrinal security mechanisms or which are prerequisite to the use of end system security management functions (e.g., issuance of physical credentials to users, hiring and scheduling human guard services, or carrying out routine maintenance of physical barriers). Although these supporting activities are not called out in most parts of section 6, they must be understood to be an integral part of security management. Examples of information domain security policy elements include:

- A brief description of the mission area and a more comprehensive description of the specific mission area function that the information domain supports

- A description of the information objects and their attributes, including rules pertaining to their use within multidomain information objects

- Membership criteria

- Rules for interdomain transfers, if any

- Security service requirements (including strength of service) appropriate to meet the risks determined by a threat analysis. Security services should be allocated to LSEs, end systems and relay systems, and the transfer system

- Criteria for acceptable mechanisms to implement the required security services

- Security management-specific requirements

  - Relationship of the security management information domain to an information domain (1:1, 1:many, or embedded as described in section 4.3.7)

- Criteria for security administrators (e.g., must be a member of the information domain, must not be a member of the information domain)

- Roles, privileges, and duties of security administrators

- Identities of security administrators

- Configuration management requirements for the establishment or modification of information domain security policy rules

• Identification of one or more members of the information domain who are responsible for accrediting information systems that will support the information domain

The security policy for an end system that supports multiple information domains must specify the management rules for conducting the following activities:

• Provide strict isolation among information domains.

• Invoke and manage security mechanisms that implement the security services required by the security policies of the individual information domains.

• Develop rules for the management of multidomain information objects, including criteria for user access, display labeling, and transfers between end systems.

• Control and maintain security management mechanisms and information objects that enable a security manager of a particular information domain to control that information domain independently of others.

The security policy rules for both end system security management and information domain security management are part of their SMIBs. Each SMIB is accessed by a security administrator using a SMAP, as explained in section 4.3.7. A SMIB is a logical collection of information that usually will exist in distributed form among several end systems. For an information domain that is supported in more than one end system, the security administrator may have physical access to only some of those end systems. Thus, the SMAP that operates on the portion of a SMIB in a particular end system must be accessible to the security administrator both locally and remotely. A SMAP is like any other application in that it operates in a security context which represents a security administrator (or process) operating in a

particular security management information domain. Thus, its security policy is interpreted and enforced by the SPDF and SPEF and it is subject to the same strict separation mechanisms as other information domains.

6.2 ISO 7498-2 AND DGSA SECURITY MANAGEMENT CONCEPTS

Clause 8 of ISO 7498-2 addresses many aspects of security management for open systems interconnection. The ISO 7498-2 security management structure is adopted as the basis for the DGSA security architecture and is extended to apply to all aspects of open systems security management.

6.2.1 Information Domains

ISO 7498-2 begins its security management discussion by considering security policy and security domains (clause 8.1.2):

> There can be many security policies imposed by the administration(s) of distributed open systems and OSI security management standards should support such policies. Entities that are subject to a single security policy, administered by a single authority, are sometimes collected into what has been called a "security domain". Security domains and their interactions are an important area for future extensions.

In the DGSA, "information domain" is substituted for "security domain". Some of the future extensions noted above have been included in the OSI Security Frameworks Overview, ISO 10181-1 (ISO, 1993a). The Frameworks Overview allows, but does not require, security domains to have subset and superset relationships. The DGSA does not allow information domains to be hierarchically related, and so has no need for the subset and superset notions. When sensitivity of information objects is a part of an information domain security policy, all the information objects in an information domain have the same sensitivity. The sensitivity of an information object is a consequence of its presence in an information domain. The "single authority" is the membership of an information domain. Usually the authority is delegated to one or more security administrators for day-to-day security management activities. The reference to "security domain...interactions" is accounted for in the DGSA by security policy interdomain transfer rules and their implementation.

6.2.2 Security Management Information Bases

ISO 7498-2 (clause 8.1.4) describes security management information bases as follows:

> The Security Management Information Base (SMIB) is the conceptual repository for all security-relevant information needed by open systems. This concept does not suggest any form for the storage of the information or its implementation. However, each end system must contain the necessary local information to enable it to enforce an appropriate security policy. The SMIB is a distributed information base to the extent that it is necessary to enforce a consistent security policy in a (logical or physical) grouping of end systems. In practice, parts of the SMIB may or may not be integrated with the MIB.

The DGSA uses SMIBs to conduct information domain and end system management, rather than for only end system management as implied above by the "appropriate security policy" for "each end system". As discussed in section 4.3.7, a distinct security management information domain may be responsible for the management of a single information domain (1:1) or several information domains (1:many), or the information domain may contain its security management information domain (embedded). The SMIB in these cases, respectively, contains security information for the single information domain, contains security information for all of the several information domains, or is contained in the information domain with its information objects. In the 1:many case, the information domains may or may not be related to the same mission. This flexibility allows a security administrator (or group of security administrators) to manage more than one information domain from the same SMIB. Also, it implies that each security administrator has the same attributes (privileges) with respect to the management information of all of the information domains that share a management information domain. (However, not every security administrator necessarily has the same attributes as the other security administrators.)

6.2.2.1 Information Domain SMIB Content

The following examples of information objects might be placed in a SMIB to manage an information domain:

- Information domain security policy rules
- Member registration information
- Member authentication criteria (e.g., strength of mechanism required)
- Member authentication information
- Member attributes (privileges) (e.g., access privileges, release authority for interdomain transfers)

- Visible security label information (i.e., what label, if any, is attached to information that is printed or displayed)

- Security service and security mechanism requirements for specific applications, including intradomain communications and interdomain information transfer

6.2.2.2 End System SMIB Content

The end system SMIB contains information for management of security functions and resources shared by several information domains, including hardware resources, security-critical functions (particularly security services and mechanisms), and supporting applications (e.g., key management). More detail is given in later sections on several of the supporting security applications and related functions. The following example classes of information objects might be included in the end system SMIB:

- End system security policy rules

- Security services management information (see section 6.2.7)

- Security mechanisms management information (see section 6.2.8)

- Supporting services and mechanisms management information (e.g., alarm reporting, information system auditing, cryptographic key distribution, security contexts, security-critical functions, security-related applications operating for the end system)

6.2.3 Communication of Security Management Information

ISO 7498-2 (clause 8.1.5) observes the following about the communication of security management information:

> Management protocols, especially security management protocols, and the communication channels carrying the management information, are potentially vulnerable. Particular care must therefore be taken to ensure that the management protocols and information are protected such that the security protection provided for usual instances of communication is not weakened.

Security management information will be protected in accordance with the security policy of each management information domain. Management applications used to

communicate security management will rely upon the same open system protocol infrastructure as other applications. Management applications operate in security contexts. Security associations that ensure secure communication between security contexts in different end systems are described in section 7.

6.2.4 Distributed Security Management Administration

ISO 7498-2 (clause 8.1.6) describes distributed security management administration:

> Security management may require the exchange of security-relevant information between various system administrations, in order that the SMIB can be established or extended. In some cases, the security-relevant information will be passed through non-OSI communication paths, and the local systems administrators will update the SMIB through methods not standardized by OSI. In other cases, it may be desirable to exchange such information over an OSI communication path in which case the information will be passed between two security management applications running in the real open systems. The security management application will use the communicated information to update the SMIB. Such updating of the SMIB may require the prior authorization of the appropriate security administrator.

The DGSA is consistent with this view and uses it as the basis for DGSA distributed security management. Each management information domain uses and maintains the SMIB of the information domain it manages The security administrator may operate locally or remotely, and may use OSI or non-OSI communication to maintain the SMIB. Security administrators may rely on a custodial infrastructure (e.g., communications security custodians). Cooperation with local administrators may be necessary for functions that cannot be managed remotely (e.g., aspects of key management that require physical access and personal accountability dictated by doctrinal considerations).

6.2.5 Security Management Application Protocols

ISO 7498-2 (clause 8.1.7) requires security management application protocols for exchange of security-relevant information:

> Application protocols will be defined for the exchange of security-relevant information over OSI communication channels.

The general management application protocol defined by ISO is the Common Management Information Protocol (CMIP) (ISO, 1987). The General Upper Layer Security (GULS) Security Exchange Service Element Protocol (SESEP) is currently a

Draft International Standard (ISO, 1993b). The Government Network Management Profile (GNMP) (NIST, 1992b) is mandated by the U.S. Government Open Systems Interconnection Profile (GOSIP) (NIST, 1992a). Unfortunately, GNMP does not yet specify particular protocols. It references both the Simple Network Management Protocol (SNMP) (Case, 1989) and CMIP.

Simple but effective management protocols to accommodate all management exchanges are needed. CMIP or the SNMP version 2 (Case, 1991), the successor to SNMP, currently are the best available choices, but the GULS SESEP could become an critical tool in the future. As the GOSIP position on management protocols becomes stable, the DGSA will adopt the chosen protocols.

6.2.6 End Security Management Functions

ISO 7498-2 (clause 8.2.1) observes the following about system security management:

> System security management is concerned with the management of security aspects of the overall OSI environment. The following list is typical of the activities which fall into this category of security management:
>
> a) overall security policy management, including updates and maintenance of consistency;
> b) interaction with other OSI management functions;
> c) interaction with security service management and security mechanism management;
> d) event handling management;
> e) security audit management; and
> f) security recovery management.

As noted previously, the DGSA broadens the view of end system security management to the entire open systems environment, especially with respect to the support of multiple information domains. The topics of event handling, security audit, and security recovery management are interrelated and will be treated together.

ISO 7498-2 (clause 8.3.1) describes event handling management as follows:

> The management aspects of event handling visible in OSI are the remote reporting of apparent attempts to violate system security and the modification of thresholds used to trigger event reporting.

ISO 7498-2 (clause 8.3.2) describes security audit management as follows:

> Security audit management may include:
>
> a) the selection of events to be logged and/or remotely collected;
> b) the enabling and disabling of audit trail logging of selected events;
> c) the remote collection of selected audit records; and,
> d) the preparation of security audit reports.

ISO 7498-2 (clause 8.3.3) describes security recovery management as follows:

> Security recovery management may include:
>
> a) maintenance of the rules used to react to real or suspected security violations;
> b) the remote reporting of apparent violations of system security; and
> c) security administrator interactions.

These security functions are related since the event handling function deals with all the apparent security violations recognized by an end system, the audit function selects those events that will be recorded, and the recovery function acts upon some of the selected events. The selection of audited events and those requiring a recovery action is determined by information domain security policies or by the end system security policy.

Event handling includes local as well as remote reporting of security-related events. Depending on whether a management entity (a security manager or a security recovery application) or a user is expected to examine or act on various alarms or audit records, alarm or audit information objects may be recorded in a particular management information domain SMIB, an end system SMIB, or a user-accessible file in an information domain.

Security recovery actions might include terminating a particular security context, temporarily prohibiting certain activities within an information domain, or disabling a particular communications interface. Some security recovery actions may depend on specialized data structures, such as a compromised cryptographic key material list, which controls continued use of key materials.

6.2.7 Security Service Management

ISO 7498-2 (clause 8.2.2) describes security service management as follows:

> Security service management is concerned with the management of particular security services. The following list is typical of the activities which may be performed in managing a particular security service:

a) determination and assignment of the target security protection for the service;
b) assignment and maintenance of rules for the selection (where alternatives exist) of the specific security mechanism to be employed to provide the requested security service;
c) negotiation (locally and remotely) of available security mechanisms which require prior management agreement;
d) invocation of specific security mechanisms via the appropriate security mechanism function, e.g., for the provision of administratively-imposed security services; and
e) interaction with other security service management functions and security mechanism management functions.

An information domain security policy may be very specific about how security service requirements are to be met (by mandating particular security mechanisms) or it may give only a general requirement for a security service of a particular strength and allow the end system management functions to select an appropriate mechanism from those available. Each of the activities in the list above is concerned with an aspect of determining how security service requirements are satisfied by security mechanisms, as discussed below.

6.2.7.1 Determining and Assigning Strength of Service

Determining security services to be used and their strength primarily is one aspect of developing a security policy for an information domain or an end system. The choices made are dependent on threats, vulnerabilities, and acceptable risk. That is, for large classes of information processing activities, a single determination of required security services can be made in advance because the value of the information being protected does not change often or quickly, nor do the vulnerabilities and risk. There are other classes of information activities for which it may be appropriate for a user to choose whether or not to employ a particular security service. For example, within the same information domain, some electronic mail messages may be of an informal or personal nature and not require a non-repudiation service, but other messages may be official business and may be required (by written policy) to employ a non-repudiation service. In cases like this, the user needs a selective means of invoking the security service, but the strength of the service is likely to be predetermined.

6.2.7.2 Assigning and Maintaining Rules for Mechanism Selection

For a given security service, one or more security mechanisms, alone or in combination with others, may be able to implement it. Some security mechanisms may be able to support more than one security service.

One of the aspects of the principle of absolute protection is that the security services chosen within an information domain security policy each have a minimum strength associated with them. Not all the security mechanisms that support a given security service need to be provided within end systems (or relay systems). In particular, the LSE may employ various doctrinal security mechanisms that contribute to the provision of one or more security services. As a result, the security mechanisms that support a given security service may be different when protecting information within an end system than when protecting information between end systems within the same LSE or between end systems in different LSEs. The resulting security service implementations must provide at least the minimum protection demanded by the security policy in all situations. Thus, to the extent that an end system supports security services with different mechanisms and a SMAP is aware (or can be made aware) of the distinctions among activities within an end system, between end systems in the same LSE, and between end systems in different LSEs, alternate choices of security mechanisms could be made.

The added complexity involved in making such choices might lead information system security architects to use only one set of mechanisms that satisfies an information domain security policy in all cases. However, in some situations this strategy would not be appropriate. For example, if some end systems in the same LSE often exchange large files, but only infrequently with end systems in different LSEs, a confidentiality mechanism necessary in the latter case might introduce an unacceptable performance penalty in the local situation, but doctrinal mechanisms could be relied upon to achieve the required level of protection.

(The question of how the strength of a particular security mechanism is determined is related but separate from making appropriate choices. Information domain security policy makers and implementors will need guidance and tools not currently available. See section 6.3.)

6.2.7.3 Negotiating Available Security Mechanisms

One or more end systems that support the same information domain may be able to support a particular security service with more than one security mechanism, but it may not be known in advance of attempted communications which of these security mechanisms may be implemented in a specific end system. In such cases, the specific security mechanisms to be employed must be negotiated between the SMAPs in the

end systems at the time the security association is established between them. Section 7 discusses security associations, and proposes a security management protocol that supports negotiation.

6.2.7.4 Invoking Security Mechanisms

The invocation of security services and security mechanisms within the end system security architecture involves several functions. Since all security services are security-critical, they are accessible only through the separation kernel, and, for applications, only through the standard kernel interface. Since most applications will rely upon the residual operating system for use of the standard kernel interface, that interaction will be transparent to those applications. If a request for a security service does not specify a security mechanism, the SMAP makes a choice among the available security mechanisms based on the information domain policy and invokes it through an appropriate operating system call. Otherwise, the SMAP invokes the specified security mechanism.

Although each application could make requests for security services and security mechanisms directly to the SMAP, there are significant advantages to adopting an API approach. APIs provide a common set of subroutine calls to a related set of programming functions or services. An API not only relieves application designers of creating a specific set of interfaces, but also allows underlying services to be replaced (by equivalent mechanisms) without affecting the application implementation. Various efforts are defining APIs for the invocation of security mechanisms. One such effort is the General Security Service (GSS) API intended for use with the Internet suite of communications protocols (Linn, 1992). The GSS API can be extended to the invocation of all security functions by making it the standard interface to the SMAP (it could be incorporated into the SMAP). A strawman proposal for an extended GSS API is found in Appendix A. The GULS project in ISO is creating a standard set of protocol elements that will be available to applications for conveying protected information between end systems.

The use of a combination of the extended GSS API, GULS, SMAPs, and the standard kernel interface can contribute to the independence of security services and security mechanisms and to their transparency to users and applications. This independence allows different security mechanisms to be accommodated at various stages in an end system life cycle, and for end systems to accommodate information domains with different security service requirements.

6.2.7.5 Specifying Interactions Among Security Service and Mechanism Management Functions

The use of some security services depends on the results of others. For example, access control usually employs the output of the authentication service. Required security service interactions must be expressed in a security policy. Similarly, some security mechanisms are dependent on others or on supporting security functions, for example, key management for cryptographic security mechanisms. These dependencies must be part of the SMIB so the SMAP can invoke the appropriate security mechanisms and functions.

6.2.8 Security Mechanism Management

ISO 7498-2 (clause 8.2.3) describes security mechanism management as follows:

> Security mechanism management is concerned with the management of particular security mechanisms. The following list of security mechanism management functions is typical but not exhaustive:
>
> a) key management;
> b) encipherment management;
> c) digital signature management;
> d) access control management;
> e) data integrity management;
> f) authentication management;
> g) traffic padding management;
> h) routing control management; and,
> i) notarization management.

The DGSA adopts this list and adds availability management.

6.2.8.1 Key Management

ISO 7498-2 (clause 8.4.1) describes key management as follows:

> Key management may involve:
>
> a) generating suitable keys at intervals commensurate with the level of security required;
> b) determining, in accordance with access control requirements, of which entities should receive a copy of each key; and,
> c) making available or distributing the keys in a secure manner to entity instances in real open systems.

> It is understood that some key management functions will be performed outside the OSI environment. These include the physical distribution of keys by trusted means.
>
> Exchange of working keys for use during an association is a normal layer protocol function. Selection of working keys may also be accomplished by access to a key distribution center or by pre-distribution via management protocols."

The DGSA incorporates standard key management techniques. Specifically, the Security Association Management Protocol (SAMP) is adopted as part of the transfer system. The SAMP is evolving from the IEEE 802.10 Standard for Interoperable LAN/MAN Security (SILS) Part 3 (IEEE, 1993), which has recently become the basis for the key management protocol standard being developed in ISO. SAMP will make use of GULS SESEP elements. It is expected that SAMP will be sufficiently general to support security association establishment as described in section 7.

There is an evolving key distribution system for U.S. Government use, the Electronic Key Management System (EKMS), from which the majority of U.S. Government cryptographic keying materials are generated and distributed. The EKMS Local Management Device (LMD) is the EKMS presence in LSEs. The EKMS is adopted as part of DGSA guidance. Although this is specific guidance, it is necessary because key management and cryptographic systems are being developed independently by vendors. A potential customer might procure several key management devices just to support a large-size, base-level LSE, some of which could be based on proprietary security management systems for vendor-specific end systems or LCS security products. These key management systems would almost certainly be incompatible with one another, thus increasing both initial and life-cycle costs, and impeding interoperability. The clear long-term solution is to develop key management and cryptographic products, including the evolving EKMS, based on the forthcoming international standards.

6.2.8.2 Encipherment Management

ISO 7498-2 (clause 8.4.2) describes encipherment management as follows:

> Encipherment management may involve:
>
> a) interaction with key management;
> b) establishment of cryptographic parameters;
> c) cryptographic synchronization.

> The existence of an encipherment mechanism implies the use of key management and of common ways to reference the cryptographic algorithms.
>
> The degree of discrimination of protection afforded by encipherment is determined by which entities within the OSI environment are independently keyed. This is in turn determined, in general, by the security architecture and specifically by the key management mechanism.
>
> A common reference for cryptographic algorithms can be obtained by using a register for cryptographic algorithms or by prior agreements between entities.

It is expected that new cryptographic products will support multiple algorithms that can be selected by each application. In such an environment, the registration of cryptographic algorithms will be necessary so that algorithm selection can be negotiated between end systems. The ability to select a cryptographic algorithm has implications for the security management of the devices involved, such as determining under what conditions the algorithm can be changed and auditing their use.

6.2.8.3 Digital Signature Management

ISO 7498-2 (clause 8.4.3) describes digital signature management as follows:

> Digital signature management may involve:
>
> a) interaction with key management;
> b) establishment of cryptographic parameters and algorithms; and
> c) use of protocol between communicating entities and possibly a third party.
>
> Note: Generally, there exist strong similarities between digital signature management and encipherment management.

When digital signatures support a non-repudiation service that relies upon a trusted third party, additional security management responsibilities may be added with respect to long-term archiving of keys and algorithm identifiers so that transactions can be verified well after they occur.

6.2.8.4 Access Control Management

ISO 7498-2 (clause 8.4.4) describes access control management as follows:

> Access control management may involve distribution of security attributes (including passwords) or updates to access control lists or capabilities lists. It may also involve the use of a protocol between communication entities and other entities providing access control services.

The "distribution of security attributes" includes their initial installation in a SMIB. Since not all the information in an information domain SMIB is necessarily locally present in every end system that supports an information domain, it may be necessary to convey access control attributes between end systems. Note that user-specific access control attributes may not always be required since an information domain security policy may confer certain access rights on all its members.

6.2.8.5 Data Integrity Management

ISO 7498-2 (clause 8.4.5) describes data integrity management as follows:

> Data integrity management may involve:
>
> a) interaction with key management;
> b) establishment of cryptographic parameters and algorithms; and,
> c) use of protocol between communicating entities.

When using cryptographic techniques to support the data integrity service, similarities exist between data integrity management and encipherment management. In some instances, within a single end system, data integrity can be attained as a by-product of strong access control mechanisms. When a strong communications data integrity service is required, cryptographic mechanisms are likely candidates. The SAMP provides means for selecting algorithms and keys for data integrity.

6.2.8.6 Authentication Management

ISO 7498-2 (clause 8.4.6) describes authentication management as follows:

> Authentication management may involve distribution of descriptive information, passwords or keys (using key management) to entities required to perform authentication. It may also involve use of a protocol between communicating entities and other entities providing authentication services.

Authentication mechanisms rely upon particular authentication information to validate a given identity. The authentication information against which user-supplied authentication information is verified is stored in the SMIB and is subject to similar

considerations as access control attributes. It should be noted that an authenticated individual identity may not be required by some information domain policies since it may be sufficient that an individual has been physically identified and allowed access to an end system to assert membership in an information domain.

6.2.8.7 Traffic Padding Management

ISO 7498-2 (clause 8.4.7) describes traffic padding management as follows:

> Traffic padding management may include maintenance of the rules to be used for traffic padding. For example, this may include:
>
> a) pre-specified data rates;
> b) specifying random data rates;
> c) specifying message characteristics such as length; and
> d) variation of the specification, possibly in accordance with time of day and/or calendar.

Traffic padding in physical layer communications devices is often managed as a configuration parameter. In an open systems environment, traffic padding in the physical layer will occur infrequently. Traffic padding in application layer protocols could be invoked as the result of a user request or as the result of an information domain security policy requirement applied to all or some class of communications. The critical management aspect of satisfying such a request is to assure that the padding is applied at the correct stage of processing with respect to other security services, such as data integrity or data confidentiality.

6.2.8.8 Routing Control Management

ISO 7498-2 (clause 8.4.8) defines routing control management as follows.

> Routing control management may involve the definition of the links or sub-networks which are considered to be either secured or trusted with respect to particular criteria.

Routing control in open systems meeting DGSA requirements will normally be restricted to choosing a particular network interface when an end system is connected to multiple CNs or LCSs.

6.2.8.9 Notarization Management

ISO 7498-2 (clause 8.4.2) defines notarization management as follows.

> Notarization management may include:
>
> a) the distribution of information about notaries;
> b) the use of a protocol between a notary and the communicating entities; and
> c) interaction with notaries.

See section 6.2.8.3.

6.2.8.10 Availability Management

Availability management is not described in ISO 7498-2. Availability management is limited to interactions with the LCS- or CN-provided management facilities for notifications of outages and, if applicable, alternate service information.

6.3 SECURITY MANAGEMENT TOOLS

Security architects will need various tools to enable them to design end systems that will support user requirements as reflected in information domain security policies. Security administrators must have available a set of tools to assist them in performing their functions efficiently and conveniently. Not all of the tools discussed here are available currently, and steps will need to be taken to ensure their timely creation.

6.3.1 Security Policy Rule Specification

To complement the development of the SPDF, a tool must be developed to assist in or perform the reduction of security policies to security policy rules that can be interpreted by the SPDF. The specification of security policy rules is a new endeavor and will require a significant research effort. As a result, the initial tools are likely to be somewhat primitive and will need to be improved for this new technology to be widely accepted.

6.3.2 Security Mechanisms Catalog

The selection of appropriate security mechanisms to implement the security services required by security policies is an activity that will require specific support that does not yet exist. There are several interrelated factors that must be considered.

The first factor is the strength of security mechanisms and other security-critical functions (e.g., separation kernel effectiveness). The second factor is the characteristics of security mechanisms, that is, what they do and do not provide, how security mechanisms interact with one another, and implementation and employment requirements for security mechanisms to work effectively. The third factor is the cost of security mechanisms, both procurement and life-cycle costs (including supporting functions such as key distribution). The fourth factor is user impacts, such as performance penalties.

To an extent, some of these factors are considered in current NSA procedures for evaluating security products. To support security architects in suggesting appropriate security mechanism choices, all of these factors must be considered. Evaluations based on these factors could be performed on implementations of particular security mechanisms or on products that implement multiple security mechanisms. The result of such evaluations would be a security mechanisms and product catalog from which security architects could make appropriate choices.

One significant aspect of the evaluations for such a catalog is that they would not result in a single composite rating for a security mechanism or product. Each security mechanism would be rated for its strength in support of particular security service. A security mechanism that supports more than one security service would have more than one strength rating. The security mechanism might have a different strength rating when used in conjunction with one security mechanism than it would with another. A security product would have strength ratings for each of its mechanisms. Clearly, establishing metrics for these strength ratings will be a formidable and critical aspect of creating the catalog.

6.3.3 Maintenance Applications for Security Administrators

Each of the security management activities discussed in section 6.2 will require automated support for security administrators. The applications that provide this support are concerned with various aspects of SMIB maintenance, key management, and examination, processing, and correlation of information such as audit records. These management applications should work together smoothly, but they must also be separable if it is desired to assign certain activities to specific security administrators. In some instances, it will be necessary to integrate security management applications with other applications. For example, X.500 Directory

Service Agents might be used to store portions of a SMIB so that user certificates are easily available to a user community.

6.4 AREAS FOR SECURITY MANAGEMENT STANDARDIZATION

Standardization of security management functions, data structures, and protocols will enable interoperation of SMAPs across many end system platforms and , thus, allow effective distributed security management. Areas for security management standardization include, but are not limited to the following:

- Security policy rule representations so that security policies can be installed remotely
- Key management functions that support the generation, distribution, and accounting of cryptographic key material
- Audit information formats so security management applications can interpret events occurring on multiple end systems that support multiple security domains
- Protocols for the exchange of security management information and for remote security management operations

7.0 TRANSFER SYSTEM

This section discusses the basic goal of the transfer system security architecture and then the means to achieve that goal. Section 7.1 discusses the basic notion of distributed security contexts and the primary function that supports them, the security association. Section 7.2 describes several supporting functions and tools needed to implement distributed security contexts and security associations. Section 7.3 discusses the relationship of the transfer system security architecture to some specific security-related topics.

In section 4, the transfer system was identified as the LCSs, CNs, and the communications protocols in end systems and relay systems. Security services allocated to the transfer system provide the basis for the protection of information in transfer. Availability is the only security service allocated to CNs and LCSs. Additional security services may be provided by LCSs, but they are only applicable to local communications.

The portion of the transfer system in end systems and relay systems consists of open system networking applications and communication protocols (including some security protocols). These applications and protocols are executed in the same security context as other user applications for a user operating in a particular information domain. Except for transfer system functions that are among the security-critical functions (e.g., network interface device drivers, cryptographic functions), transfer system software does not need to be trusted. The transfer system must be managed, so the SMAP and SMIB of section 6 are extended to account for transfer system functions.

The primary goal of the transfer system security architecture is to provide protection of information in transfer to support information sharing and distributed processing within the security architectures of the other DGSA elements and the fundamental concepts. The basic approach to achieving this goal is to enable security contexts in different end systems or relay systems (that support the same information domain) to communicate as if they were in the same end system or relay system. The transfer system security architecture must fit within the end system and relay system architecture of section 5 and the security management architecture of section 6, and it must extend the support of fundamental DGSA concepts to communications, especially information domains, strict isolation, multidomain information objects, and absolute protection. The remainder of section 7 addresses various concepts and functions needed to achieving the transfer system goal.

7.1 DISTRIBUTED SECURITY CONTEXTS

The generic transfer system security architecture seeks to create structures in which applications in security contexts in different end systems or relay systems (that support the same information domain) communicate with the same assurance as if they were in the same end system or relay system. Such structures are referred to as *distributed security contexts*. There are two basic classes of communication that must be considered, *interactive* and *staged delivery*. Staged delivery refers to communications in which the information being transferred is sent from the originating end system application to a relay system application, in its entirety, and then is sent from the relay system application to the destination end system application. (There may be several relay system applications involved before the information is finally delivered to the destination end system application.) The most common example of staged delivery is electronic mail. Interactive communications include all non-staged delivery applications. The means used to create distributed security contexts are different for interactive and staged delivery communications and will be discussed separately.

7.1.1 Distributed Security Contexts for Interactive Communications

An *interactive distributed security context* is formed when two security contexts in different end systems are joined securely using a set of mechanisms that is referred to as a security association. A *security association* is the totality of communication and security mechanisms and functions (e.g., communications protocols, security protocols, doctrinal mechanisms, security-critical mechanisms and functions) that securely binds together two security contexts in different end systems or relay systems supporting the same information domain.[7-1] A security association extends the protections required by an information domain security policy within an end system to information in transfer between two end systems and it maintains strict isolation from other information domains. A security association can be considered an extension or expansion of an OSI application association. OSI application layer entities in different end systems employ application associations to communicate. An application association is composed of appropriate application layer functions and protocols plus all of the underlying communications functions and protocols at other layers. A security association is an application association that includes additional support from security functions and mechanisms.

[7-1] Note that the DGSA meanings of security association, agreed set of security rules, security association identifier, and security association management protocol are different and more general than their meanings in existing protocol specifications.

The security management information for a security association is contained in a SMIB data structure called the *agreed set of security rule*s (ASSR). The ASSR includes all the security-relevant attributes required to establish and maintain a security association, such as the information domain label and secure communication attributes (e.g., cryptographic algorithm identifiers and keys). Each end system supporting a security association chooses a local *security association identifier* (SAID). The pair of SAIDs uniquely identifies the particular security association and links it to the ASSR for that security association. Thus, any security function supporting the security association obtains necessary information from the ASSR.

7.1.1.1 Security Association Establishment

A security association is established using a SAMP. The originating end system SAMP implementation, invoked by a SMAP, creates an OSI application association with the destination end system peer SAMP implementation.[7-2] The SAMP implementations (hereafter referred to as SAMP machines) cooperate to establish a security association through a set of SAMP exchanges. The SAMP exchanges include three basic functions. Initially, the originating SAMP machine makes known its secure communications capabilities to its peer SAMP machine in the form of one or more object identifiers (i.e., an Abstract Syntax Notation.1 *Sequence of Object Identifiers* (ISO, 1988d)). An entire set of capabilities may be referenced by a single object identifier that corresponds to a specific registered ASSR. The destination SAMP machine subsequently generates the appropriate response (i.e., a positive acknowledgment to continue the SAMP exchanges or an error response). This first paired SAMP exchange is always conducted in the "clear" (i.e., it is not cryptographically protected). A positive response does not indicate that the security association has been established nor that the destination intends to accept the security association, but only that the subsequent exchanges can proceed. (Some of the considerations involved in deciding whether to allow a security association to be established are considered in the next subsection.)

The primary purpose of the second SAMP exchange is to establish the keys needed for cryptographic security mechanisms. Generally, security associations will rely upon cryptographic mechanisms so that sufficient strength and assurance for the security services is provided. One or more mechanisms may be chosen to support a particular information domain security policy. Either asymmetric or symmetric key

---

[7-2] Like the end system security context that is dedicated to user login activities, there will be an end system security context dedicated to processing incoming SAMP exchanges since the information domain to be supported will not be known reliably until the security association has been established.

generation, coordination, and exchange techniques may be employed. The SAMP must be general enough to support any standard key management technique. Depending upon the key management technique selected, various processes or key management protocols may be executed to form a traffic key. Examples include the Secure Data Network System (SDNS) Key Management Protocol for exchange of certificates and associated user keying material, and the X9.17 (ANSI, 1985) and Rivest-Shamir-Adelman (RSA, 1978) key management protocols and associated techniques developed to support commercial cryptography. This second SAMP exchange may or may not be conducted in the clear, depending on the key management technique employed. Other security functions may be performed in conjunction with the second SAMP exchange, such as access control checks based on information conveyed in a certificate delivered by the SAMP. If multiple cryptographic algorithms or keys are required to support a security association, it may be possible to convey the required information in one exchange or it might be necessary to repeat the second SAMP exchange several times.

The third SAMP exchange employs the encryption algorithms and keys established in the second exchange to test the state (liveness) of the security association. This exchange also provides peer entity authentication between the cooperating SAMP machines, reliably sends any remaining security attributes needed to operate the security association which are not already in the ASSR (e.g., the specific security services to be supported by the security protocols selected in the second exchange), and may validate the information used in the earlier SAMP exchanges. When high assurance is required that a security association has been established between two end systems that are accredited for support of a particular information domain, the peer entity authentication must be based on cryptographic techniques. The peer entity authentication is not assured until the algorithm and keys have been tested between the two cooperating SAMP machines.

When the third SAMP exchange has been successfully completed, the security association is established. The destination end system then creates a security context for the appropriate information domain and initiates the execution of the applications necessary to communicate with the originating end system applications, including the communications protocols and the SMAP for the information domain. The security contexts in the originating and destination end systems are now joined by the security association to form the interactive distributed security context.

7.1.1.2 Additional Aspects of Interactive Distributed Security Contexts

The decision to allow establishment of a security association may require several related functions to be performed such as the exchange and processing of security

attributes of the user (e.g., authenticated identity, access privileges). These attributes might be contained in a security certificate such as that defined in the X.509 Directory Services Authentication Framework (CCITT, 1992). The information contained in an X.509 certificate may be signed by any number of hierarchically related certificate-issuing authorities, down to an information domain-specific certificate-issuing authority if that level of granularity is required. This signature verification adds greater assurance to the credibility of the information contained in the certificate.

Multiple security protocols may be included in a single security association to provide a combination of security services. For example, a network layer protocol might provide continuous end system origin authentication and data integrity, while a presentation layer protocol might provide selective field data confidentiality. Some lower layer security protocols can multiplex several security associations between the same end systems. The security associations share the same cryptographic algorithm and keys. This arrangement may be appropriate for interactive distributed security contexts that support the same information domain, but it is unlikely to be acceptable for different information domains because of strict isolation requirements.

In some instances, an interactive distributed security context will be formed between end systems that employ no security protocols and may not even require an authenticated user identity. Such instances include access to public information utilities (e.g., a news wire service feed) or completely unprotected end systems. In these instances, an end system that supports other information domains, then the end system strict isolation mechanisms will be entirely responsible for maintaining the isolation of unprotected information domains from other information domains.

Some communications between end systems involve information that is not ordinarily stored in an end system, for example, real-time voice and video applications. In these cases, users must monitor and enforce the accuracy of the security context and association established for the distributed security context. That is, humans must ensure that information exchanged belongs to the information domain represented by the distributed security context as is currently done when using Secure Telephone Unit-IIIs for secure voice or data communications.

Figure 7-1 illustrates the relationships among the primary components that create a security association.

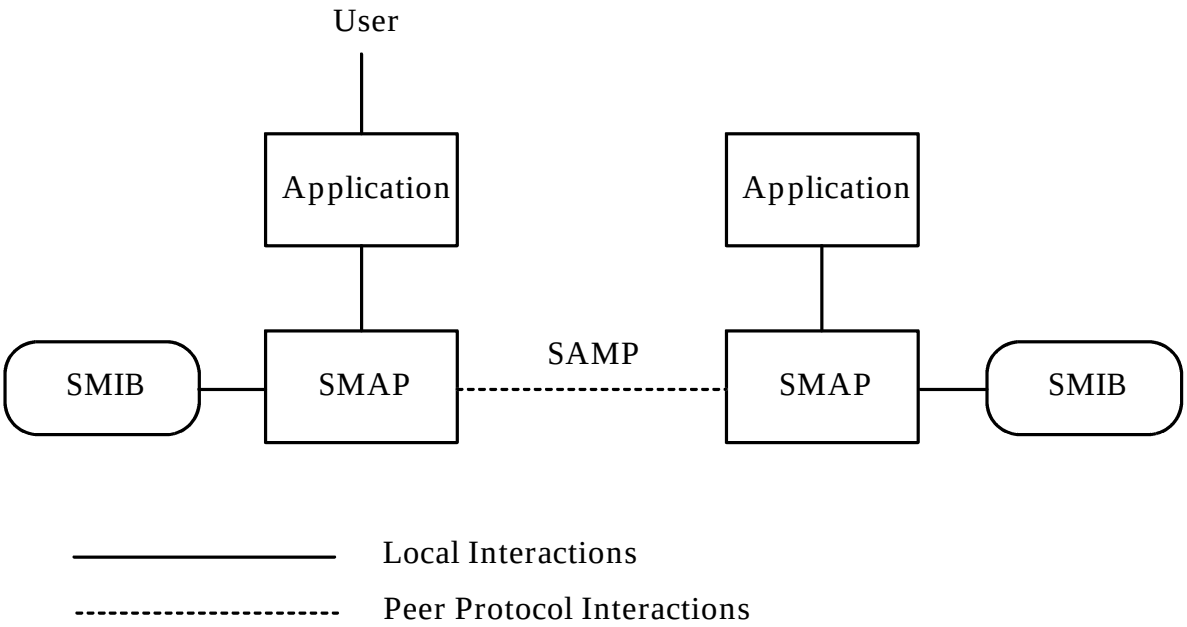

Figure 7-1. Component Relationships for Creating a Security Association

Appendix B provides a detailed example of the establishment of an interactive distributed security context.

7.1.2 Staged Delivery Distributed Security Contexts

A staged delivery distributed security context is transferred from the originating end system to the destination end system. This is accomplished by an application in the originating end system cryptographically wrapping the information to be transferred in a form that allows the destination end system to reconstitute the security context in which the information was wrapped. The wrapped information is transferred (in stages) from the originating end system to the destination end system. Ideally, the wrapping process should provide all security protection of the information while in transfer. No security services (other than availability) should be expected of the application relay systems involved in the staged delivery because they might be provided by common carrier providers, as is the case for CNs. If the wrapping process cannot provide all the necessary security protection, the application relay systems will have to be implemented to support the DGSA and interactive distributed security contexts between end systems and relay systems will have to be used to ensure the secure staged transfer of information.

There is an existing specification for a secure electronic mail service that satisfies the requirements for staged relay distributed security contexts. This document is the SDNS Message Security Protocol (MSP) specification (NSA, 1992). MSP can provide authentication, access control, message confidentiality and integrity, and non-repudiation security services. MSP allows delivery of the same message to multiple recipients supported by several end systems without creating multiple copies of the message in the originating end system. Multiple messages created in different security contexts can be combined in a single MSP transfer. The wrapping of the messages takes place in the originating end system in an MSP user agent. The wrapped message is submitted to the message transfer system, which consists of a group of untrusted cooperating message transfer agents. The message is delivered to one or more destination MSP user agents, which unwrap the message. For details of how secure staged delivery can be achieved, the MSP specification should be examined. MSP will be the basis for secure messaging in DoD as Phase II of the Defense Message System is implemented and deployed.

7.1.3 Other Aspects of Distributed Security Contexts

7.1.3.1 Multidomain Object Transfer

Section 4.3.4 defined and discussed multidomain objects and noted that their purpose is to display or print related information objects from several information domains in an ordered format. Section 5.2.2 discussed some high-level implementation aspects of multidomain objects. The transfer of a multidomain object between end systems requires that both the component information objects and the description of their relationships be transferred. Since a distributed security context supports transfer of information within a single information domain, one distributed security context is used for each of the component information domains. If the description of the component relationships is contained in an information object in a separate information domain, another distributed security context is required for its transfer. An application similar to those used to display or print multidomain objects is needed to coordinate the transfer of the component information objects.

7.1.3.2 Distributed Security Context Single Information Domain Restriction

The definition of a distributed security context restricts it to joining end system or relay system security contexts that support the same information domain. In principle, this restriction could be removed, however, there are practical reasons for retaining it. One of the principle functions of a distributed security context is to maintain strict isolation of information in transfer. Within an end system, the separation kernel (or other strict isolation mechanism) controls all interactions

between security contexts. As noted earlier, it is expected that cryptographic mechanisms will be the usual means to maintain strict isolation for information in transfer. The use of such cryptographic mechanisms requires shared use of keys and other supporting information between security contexts in the communicating end systems. If those security contexts support different information domains, sharing of the keying information is difficult. There will also be additional complexity introduced into many communications and security protocols that will result in trusted implementation of additional functions. The restriction that distributed security contexts support transfers within a single information domain is intended to simplify implementations that support the DGSA concepts.

7.2 TRANSFER SYSTEM SUPPORT

This section describes several elements needed to support the basic transfer system activities.

7.2.1 Security Management Application Process

In addition to the SMAP functions described in section 6, it also controls the establishment and termination of all security associations and distributed security contexts, and all transfer system security services and mechanisms. Additional transfer system-related SMAP functions and interfaces support the following activities:

- End system communications applications requests (through the extended GSS-API)

- Additional SMIB information object use and maintenance (e.g., to access information for remote security administration maintenance, security protocol and algorithm operation, certificate processing)

- Maintenance and retrieval of security information from the X.500 Directory using the directory access protocol

- MSP processing for staged delivery secure messaging for both transmission and receipt

- SAMP operations for establishment of interactive distributed security contexts, including security protocol operation, termination, and recovery, plus maintenance of the SAID and ASSR structure for each security association established

- Cryptographic and key management functions for security service and security protocol operation

- General-purpose management protocol operation (e.g., CMIP) to accomplish secure exchange of security information between distributed SMAPs or network management information requested by network management systems

7.2.2 Security Management Information Base

Additional information is required in the end system SMIB and the information domain SMIBs to support transfer system operations.

Additional information domain SMIB information items include:

- X.509 certificates to carry appropriate security information, such as SDNS key management certificates

- User access control information for distributed operations

- Traffic and message keys

- Accumulated audit data, including records of distributed security context utilization

Additional end system SMIB information items include:

- Key management, encipherment, integrity, and signature algorithm identifiers, and security protocol objects

- End system access control information for distributed operations

- Encryption algorithm initialization information

- Security association configuration information (e.g., ASSRs, SAID tables)

- Compromise action information (e.g., revoked certificates lists)

- Contingency plan parameters (e.g., auto-purge and security policy replacement actions under emergency conditions)

Some SMIB items may be held in Directory Service Agents for ease of access by many users. Such items might include key management information (i.e., SDNS certificates and user keying material) used by MSP implementations. SMIB information stored in X.500 Directories must be integrity protected.

7.2.3 Security Protocols

Several security protocols, either existing or in development, are candidates for use in end systems implementing the DGSA. Others may be added over time.

The Transport Layer Security Protocol (TLSP) is an ISO standard (ISO, 1992c) and the Network Layer Security Protocol (NLSP) is a Draft ISO standard (ISO, 1992a). The IEEE 802.10 SILS Secure Data Exchange (SDE) protocol standard (IEEE, 1992) is appropriate for LCS security services (beyond availability) when needed. MSP is the DoD standard for electronic messaging.

No current SAMP meets DGSA needs, but an ISO project under way that is developing such a protocol. The initial version is based on IEEE 802.10 SILS Part 3. The GULS SESEP will carry SAMP exchanges. SESEP is expected to be the carrier for all new application and presentation layer security protocols, so it will be included, at least implicitly, among security protocols for DGSA implementation. If the SAMP standard that emerges from ISO contains the functionality as currently planned, it will be a suitable protocol for DGSA implementation.

There are many existing and planned physical layer encryption and transmission security devices (which are necessarily communications technology-specific). When traffic flow security services are required, these devices may be used (see section 7.3.1).

Table 7-1 identifies the security services supported (actual or planned) by each of the security protocols discussed.

Table 7-1. DGSA Security Protocols and Security Services

Security Services \ Security Protocols	SAMP (Protocol Layer 7)	MSP (Protocol Layer 7)	TLSP (Protocol Layer 7)	NLSP (Protocol Layer 4)	SILS SDE (Protocol Layer 3)	Link Encryptors (Protocol Layer 2, 1)
Peer Entity Authentication	Y	N	N	N	N	N
Data Origin Authentication	Y	Y	I	I	I	I
Access Control	Y	Y	Y	Y	N	N
Connection Confidentiality	P	N	Y	O_1	O	N
Connectionless Confidentiality	N	Y	N	O_2	N	N
Selective Field Confidentiality	N	N	N	N	N	N
Traffic Flow Confidentiality	N	N	N	N	N	Y
Connection Integrity with Recovery	N	N	O_1	O_1	N	N
Connection Integrity without Recovery	N	N	O_1	O_1	O	N
Selective Field Connection Integrity	N	N	N	N	N	N
Select Field Connectionless Integrity	N	N	N	N	N	N
Connectionless Integrity	N	Y	N	Y	N	N
Non-Repudiation (Proof of Origin)	N	Y	N	N	N	N
Non-Repudiation (Proof of Delivery)	N	Y	N	N	N	N

Cell Legend:

Y	Yes, the service is supported by the security protocol
N	No, the service is not supported by the security protocol
I	The service is implicitly provided, based on unique key
O	The service is optionally provided
O_1	Connection-oriented service option
O_2	Connectionless service option
P	The service is provided only for parts of the last protocol exchange

7.2.4 Cryptographic Support

The creation of distributed security contexts, which provide communications security services and strict isolation adequate for sensitive information, is usually dependent

on cryptographic mechanisms. Thus, the availability of low-cost cryptographic devices is a critical element of the DGSA. These cryptographic devices must be sufficiently flexible to support requirements of different information domains in the same end system.

This flexibility will be achieved if the devices accommodate multiple cryptographic algorithms and multiple key management schemes, including public key encryption schemes and various key distribution center schemes Otherwise, a multiplicity of cryptographic devices will be needed, resulting in increased costs. To manage these devices, there must be a registry of cryptographic algorithms and key management schemes so that the specific choices can be negotiated for a particular security association.

Currently available cryptographic and key management devices do not meet these flexibility criteria. Very large scale integration (VLSI) chip technology may now have reached a sufficient density to achieve a cost-effective single-chip design which can support multiple algorithms and a variety of key management schemes, along with a cache memory capable of handling reasonable quantities of key material. The cryptographic devices must be capable of a minimum throughput rate of 10 megabits per second to be useful with high-performance workstations. Isolation techniques must accommodate concurrent algorithm execution and Red/Black separation (in software, hardware, or both). The DGSA is achievable only if this kind of cost-effective technology is available. In addition to creating low-cost devices, COMSEC custodial functions must be minimized through the use of electronic key management technology.

7.2.5 Distributed Management Systems

Distributed management of information systems both supports the transfer system and relies upon the transfer system for its operation. Management systems will rely upon the same transfer system security structures (distributed security contexts, security associations, and security protocols) as any other application.

When distributed information systems become very large, their management becomes very complex. To make the complexity manageable, hierarchical management approaches are often adopted. It then becomes necessary to coordinate the levels of delegated management authority. The coordination is achieved by the way management information is organized and through the control of that information as required by security policies. Hierarchical management relationships are not reflected in the way management applications communicate with one another. That is, management protocols are peer oriented, not hierarchically related. When the

term "hierarchical management system" is used, it must be understood that a set of information relationships is being described, not a communications structure. This means that the hierarchical aspect of management is a human, organizational function. The organizations and administrators that manage information systems may be organized hierarchically. Management information may reflect that organization, but the end systems in which management applications are implemented only communicate as peers.

Management systems are composed of management applications implemented in end systems. Some management applications must coexist with other applications in end systems, but for logistical reasons it may be desirable to dedicate some end systems to management system activities. Management systems can be grouped into three categories based on the particular type of management function being performed. While these categories are logically separate, they often support one another. The three categories are network management, security management, and information management.

Traditional network management systems are network control centers that monitor and configure network components, perform fault isolation functions, and collect accounting and performance information. Security management systems typically provide information to support security services and mechanisms in end systems and relay systems. Most often the support is for cryptographic mechanisms. Example systems are the DoD EKMS, BLACKER Access Control Centers and Key Distribution Centers, and the CANEWARE security management components. Information management systems include are X.500 Directory systems, the Internet Domain Name Service and the Network Information Center.

Although these three logical categories of management systems could be implemented in end systems dedicated to the functions of only one of them, as a practical matter, some of the functions can be expected to be supported on common end systems. However, each logical category may require unique technical administrative expertise. In some cases, it will not be prudent to assign multiple administrative functions to individuals because too much control might be entrusted to them.

7.3 DGSA TRANSFER SYSTEM ISSUES

Two aspects of the DGSA transfer system deserve further discussion. One is the use of traffic flow security mechanisms, and the other is potential limitations on distributed processing functions.

7.3.1 Traffic Flow Security in Open System Communication Environments

The DGSA open system and common carrier communications requirements result in the allocation of security of information in transfer to LSEs, particularly end system security support for the transfer system. The use of common carrier CNs precludes the use of full traffic flow security (TFS) mechanisms. Full TFS mechanisms operate at the physical protocol layer. Only if communications facilities are owned or controlled by user organizations can full TFS be applied.

The clear cost disadvantages of owning and operating private CNs means that there must be a careful examination of threats and vulnerabilities to determine whether full TFS is required. Unless it is necessary to subject all communications to full TFS, the DGSA requirements for open system and common carrier communications can be met with multiple communications connectivity. The strict isolation mechanisms required in end systems make it possible to support multiple communications connections among the information domains supported. Partial TFS mechanisms should be considered as alternatives to full TFS when judged to be appropriate to the known threats and vulnerabilities.

7.3.2 Limitations on Distributed Processing

Some communications technologies are inherently of a broadcast nature (e.g., radio, broadband LANs). Broadcast technologies make it possible to communicate with any end system that has access to the medium without the need to explicitly address information to specific end systems. Broadcast-like effects, called multicasts, can be achieved over non-broadcast communications systems through various methods that address and send information to (possibly large) groups of recipient end systems or users (e.g., groups of electronic mail recipients).

Certain limitations are encountered if cryptographic mechanisms are used to support security services for broadcast (and some multicast) communications. There are two basic choices. First, for true broadcasts, a single encryption key must be shared among all recipients. The use of a shared key among large numbers of recipients not only increases the likelihood that the key will be compromised, but the distribution and use of one or more shared keys is difficult to coordinate. (The same considerations apply to multicast services that depend on broadcast media.)

Second, for multicasts that are addressed to the group of recipients, a single key can be used for the security mechanism applied to the information to be sent and that key can be replicated and protected with a cryptographic mechanism using a

different key known to each recipient (e.g., MSP confidentiality service for multiple recipients).

Thus, if it is desired to broadcast information to all the members of an information domain, group multicasts are likely to be sufficient for most purposes since the member addresses are known. The only real limitation on broadcast communications is that the inherent broadcast capabilities of some media cannot be used.

(This page is intentionally blank.)

8.0 SECURITY DOCTRINE

The achievement of total security for an information system must rely on people and the environment. When products are designed and deployed in information systems, conditions of their use must be met to complement the protection afforded by any hardware and software security mechanisms employed in those products. The specification of the conditions for the use of a component, facility, or system is found in the *security doctrine*, which is part of the security policy. The security policy specifies how security requirements are to be met and, thus, is an element of the security architecture. As with any security architecture design element, there will be different types of security doctrine allocations, each with different degrees of specificity, which eventually lead to the satisfaction of the required security services through the choice of appropriate security mechanisms. In the case of security doctrine, security services are provided by physical, administrative, personnel, and operational security mechanisms. The DGSA assigns certain fixed security service allocations to doctrinal security mechanisms. For the designer of more specific security architectures, choices must be made regarding security service allocations and types of security mechanisms. This section presents the fixed security service allocations made to security doctrine and examples of doctrinal security mechanisms that are permissible and consistent with the DGSA.

An LSE and its components must satisfy each of the information domain security policies for which it is accredited. The doctrinal security mechanisms employed may vary among information domains. For example, one information domain may require authentication of the identity of an individual through cryptographic-based mechanisms, while another may rely on the simple possession of a badge. The LSE is the principal location for direct implementation of doctrinal security mechanisms, but local security mechanisms may rely upon remote systems to provide initial capabilities and life-cycle support (e.g., key management systems, personnel investigations, shrink-wrapped software, security inspection and testing, security training and awareness).

8.1 SECURITY SERVICES

The security services of ISO 7498-2 (authentication, access control, data confidentiality, data integrity, and non-repudiation) are extended in their definitions for use in the DGSA beyond only communications. The DGSA also includes availability among the security services. In section 4, only availability is allocated to the LCS in the LSE, while all the security services are allocated to the environment and to the end systems and relay systems. The environment protects the end systems, relay systems, and the LCS. Security services are implemented in LSEs in

the form of physical, personnel, and administrative security mechanisms. In addition, some types of physical security mechanisms may be incorporated into the hardware of components within the LSE.

8.2 DOCTRINAL SECURITY MECHANISMS TO PROVIDE SECURITY SERVICES

All, some, or none of the responsibility for provision of each of the security services may be allocated to doctrinal security mechanisms.

8.2.1 Mechanisms for Identification and Authentication

Authentication of the claimed identities of individuals, as individuals or as members of a group, is a typical security policy requirement. Authentication mechanisms provide varying degrees of credibility that such claims are correct. Authentication responsibilities are often shared between doctrinal, hardware, and software mechanisms. Probably the most common mechanism is the picture badge and the human recognizer. The picture on the badge matching the appearance of the holder affirms the association of the individual with what the badge represents. The identity of the individual is thereby authenticated and, in some cases, the possession of the badge establishes further claims. The reading of the magnetic code on a badge matched with the entry of a personal identification number is similar in capability to picture confirmation. Similarly, the matching of fingerprints or retina images authenticates the identity of an individual.

The use of keys with locks, passwords, or cipher lock codes authenticates identity only to the extent of the probability that the presenter is a valid holder of the object or information. That probability is based on the administrative handling and physical protection of such mechanisms or information. The same considerations apply to the use of smart cards, cryptographic ignition keys, and other credentials that make no positive connection with the holder. In general, non-forgeable information bound to the holder is the strongest type of authentication mechanism. Security mechanisms for authentication depend upon system security administrators who perform the initial assignment of the badge or other credential to an individual.

8.2.2 Mechanisms for Access Control

Access control mechanisms enforce security policy requirements for the isolation of assets and information from people and their agents. Access control mechanisms also permit authorized access to assets and information. The first line of protection for the LSE is through mechanisms that control access to the facilities (e.g., buildings, rooms) containing the end systems, relay systems, and LCSs. The human security

guard is one of the most familiar types of access control mechanisms. Key, combination, and cypher locks are common mechanisms for controlling access to facilities. Placing an entire LSE within a vault is an extreme form of facility control. With the assumption that only authorized people are in the LSE, surveillance of their activities by security administrators or by co-workers can form the next line of protection. Areas may be declared to require at least two people to be present when activities are in progress ("no-lone" zones).

The next line of protection involves the use of approved containers (e.g., combination safes and locking cabinets) for the protection of system assets. Such containers can be used to protect entire system components (end systems, relay systems, and LCSs) or information storage media (e.g., disks, tapes). Finally, the components themselves may contain access control mechanisms such as power locks, two-person-control devices, and sealed housings.

Within and beyond these lines of protection, access control becomes the responsibility of hardware and software features of the end systems and relay systems. Access control mechanisms can also contribute to the provision of confidentiality, integrity, and availability services; independent aspects of these services are presented in the following sections.

8.2.3 Mechanisms for Confidentiality

Confidentiality mechanisms satisfy security policy requirements to protect information from unauthorized disclosure. The major applications of doctrinal confidentiality mechanisms in LSEs involve video displays, printing devices, sounds, and non-video electromagnetic emanations.

Users and security administrators can control when, where, and in whose presence video information is displayed. Video display emanations can be controlled through screen filters and shielded enclosures. Printer ribbon handling, copy counting, and labeling requirements can be controlled by users, operators, and system administrators. The control of trash and the destruction of paper and other media are important procedures. Paper shredders may be useful. Procedures for handling and mechanisms for erasure of persistent storage media can be critical to confidentiality. Sound insulation and sound masking can be used to control disclosure through conversations and machine noises. Electromagnetic emanations, either radiated or conducted, can be confined by shielding rooms and by filtering signal and power wiring using standard TEMPEST features. The presence of copiers and photographic equipment in LSEs requires careful control. Proper wrapping prior to shipping or mailing of paper and other media devices is an important function.

8.2.4 Mechanisms for Integrity

Integrity mechanisms respond to security policy requirements to protect information and other system assets from unauthorized modification. The major applications of doctrinal integrity mechanisms in LSEs involve the correctness of end system and relay system hardware and software, and the correct functioning and use of other doctrinal security mechanisms. System components may have features that permit security diagnostic checking of hardware (for example, through comparison of diagnostic known-answer tests compared with off-line security check mechanisms). Non-forgeable seals and protective coatings may be used on hardware components and subcomponents to detect or prevent alteration. Cryptographic and non-cryptographic checkvalue mechanisms can be used to ensure the integrity of software packages as delivered and as used.

Regular inspections of facilities and system components is an important part of using integrity mechanisms. Protected storage of devices used for integrity checking is of value. The protected storage of software master copies and small system components while not in use is also valuable. Protection from electromagnetic interference can be accomplished by filtering and shielding.

8.2.5 Mechanisms for Non-Repudiation

Non-repudiation mechanisms support security policy requirements for proof of delivery and proof of origin of information transactions. Non-repudiation mechanisms may include the contents of a transaction. For paper transactions, notary services and personal signatures are useful mechanisms in providing non-repudiation services. Non-repudiation mechanisms, such as hash coding of data and digital signatures, can be used to validate the source of software packages. Non-repudiation mechanisms could be used for verifying that hardware is unchanged from its manufactured state.

8.2.6 Mechanisms for Availability

Availability mechanisms in communications networks and LSEs satisfy security policy requirements for availability of communications and processing resources. The ability of communications networks to provide timely and regular service depends upon the total security architecture, implementation, and management of those systems. The techniques of redundancy, diversity, contingency reserves, and contingency planning play a large part in communications network availability. Within LSEs, the LCS must be similarly designed and protected to avoid failure

outages. Generally, the physical protection and integrity checking of the end systems, relay systems, and LCSs will provide for their availability.

8.3 COTS PRODUCT CONSIDERATIONS

Section 8.2 has presented a number of doctrinal security mechanisms appropriate for use in LSEs in support of security policies. Generally, COTS products have few, if any, built-in security mechanisms. Thus, COTS products are vulnerable to component modification and substitution. Any user not under surveillance may be able to modify or substitute COTS product components to his benefit. The doctrinal reaction must be that COTS products can be physically accessed only by persons authorized for access to all information in the component unless escorted by someone who is so authorized. At the other extreme, when sufficient built-in isolation mechanisms exist (in GOTS products or custom-designed products), then all communities of interest can be satisfied that physical access is permissible by persons authorized in only one information domain of all those supported.

8.4 SECURITY MANAGEMENT

Security management, as presented in section 6, includes security service management, security mechanism management, and security of management. All of these functions are performed within the LSE. The information domain security manager is an administrator who is authorized to perform installation and maintenance of the information domain security policy representation, access control lists, and other items of the SMIB, such as cryptographic keys. The security manager is provided tools, such as the SMAP, to perform these tasks. The security manager is ultimately responsible for checking personnel clearances, monitoring guard activities, performing audits of security-relevant records, and, in general, supporting all other security mechanisms.

The security of system management is no different from any other application which requires protection. The system must have a security policy, and doctrinal mechanisms will be given specific responsibilities in support of system management. A critical aspect of security management is the training of security administrators and users so that they understand their responsibilities as part of the entire security posture.

(This page is intentionally blank.)

9.0 EXAMPLE LOGICAL ARCHITECTURE DEVELOPMENT PROCESS

The purpose of this section is to illustrate, by example, how the principles and concepts of the DGSA may be applied to specific organizations and their missions. The knowledge to be gained from this section is the process by which the DGSA may be applied, an understanding of the principles and concepts described in this report, and the specific security policies and architectures.

The process for developing a logical or specific architecture, in accordance with the DGSA, begins with the mission. An understanding of the mission's goals, users, information, and inter-mission interactions must be achieved before the process of designating information domains can begin. Once the information domains have been identified, the development of information domain security policies can commence. Finally, information domains must be implemented in real systems. A specific system's security policy and architecture should reflect the requirements of the information domain security policies supported by the particular system.

In the following sections, the mission functions and the relationships among the example mission areas are used to illustrate:

- The correlation among specific sets of mission functions
- The derivation of information domains, information domain security policies, and system security policies to support the mission functions
- The identification of methods for sharing information
- The accomplishment of security management through alternate techniques for forming security management domains
- The use of information domains and their policies to develop a specific system security architecture
- The application of an interdependency analysis to measure the effectiveness of the system security architecture (i.e., how residual risk can be formulated and addressed before actually implementing the security architecture)

The missions chosen for this example, while designed to reflect real-world situations and requirements, are strictly fictitious. Any references to specific organizations, mission threats, system capabilities, or interoperability requirements are for example

purposes only and are not accurate reflections of the actual information processing capabilities or environments of any of the organizations.

9.1 MISSION DESCRIPTIONS

Three specific mission areas are the focus of this section: intelligence, tactical, and drug interdiction. Each of these mission areas are assumed to have existed independently before additional requirements for cooperative activities were imposed. The intelligence mission function must provide photographic reconnaissance information to specific tactical and drug interdiction mission area functions. The tactical mission area function is to conduct flight planning for naval air sorties that will be used to enforce a "no-fly zone" over a particular geographic area. In turn, the tactical mission function provides information on ground-based activities to the intelligence mission function. The drug interdiction mission area function is targeting and tracking foreign and domestic illicit drug activities. (There is no relationship between the specific tactical and drug interdiction mission functions.) Figure 9-1 illustrates the mission functions and their relationships.

These three (hypothetical) specific mission functions were derived from the mission charters of each broad mission area (i.e., tactical, intelligence and drug interdiction) which were in turn derived from U.S. National and Foreign Policies. Table 9-1 provides a summary of the hierarchical relationships among national level policies, responsible action offices, and missions used within this example.

9.1.1 Intelligence Mission Area Function Description

The intelligence mission area function produces intelligence information based on data received from a photographic reconnaissance satellite. The data contains imagery and text information, including a time stamp and other data (e.g., source, accuracy). This information is considered to be extremely sensitive and classified Top Secret (TS). Full disclosure of source, accuracy, image resolution, or the photo image data to unauthorized persons could result in extreme damage to national security. The imagery information will be downgraded before being transferred to other mission areas by removing the information's source and method, and possibly lowering the resolution of the image to remove knowledge of the actual resolution capability of the intelligence source.

The specific responsibilities within the intelligence mission area function include the following categories (positions): mission control, satellite uplink/downlink command and collection, collection post-processing and production control, intelligence analysis

and reporting, and security management. Table 9-2 provides a description of the responsibilities associated with each of the intelligence mission area positions.

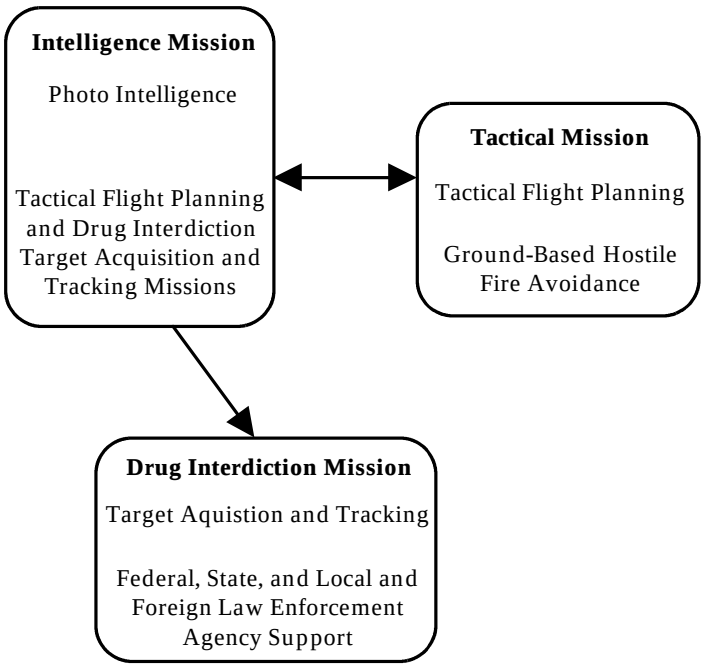

Figure 9-1. Example Mission Functions and Their Relationships

The individuals who fill these positions may have different privileges assigned to them to carry out their responsibilities. Some individuals may fill different positions at different times so their privileges must be sufficient to carry out all their responsibilities. All of the individuals associated with the intelligence mission function have at least a Top Secret clearance. There is no compartmentalization required within this mission area operation. All participants have a need-to-know authorizations for all mission area information.

Table 9-1. Hierarchical Policy/ Mission Relationships

Responsible Action Office	Intelligence Mission (Foreign Intel Policy Driven)	Responsible Action Office	Tactical Mission (Foreign Policy Driven)	Responsible Action Office	Drug Interdiction Mission (National/Foreign Policy Driven)
USG Executive Branch	Support all U.S. Government agencies with intelligence information required to fulfill their missions as mandated by the Executive Branch	U.S./DoD	Enforce UN resolutions by providing a tactical air presence	U.S. DOJ and DEA (National Action Agents) DOS (Action Agent for Foreign Drug Interdiction Elements)	Decrease availability of illicit drugs within and the U.S. ("war on drugs" policy). Applies to all law enforcement agencies, DEA, and cooperating foreign government drug interdiction elements
DCI and SECDEF	Support U.S. foreign policy with tactical intelligence information and drug interdiction information	U.S. and Coalition Forces	Enforce "no-fly zone" in country for which UN sanctions are mandated - Tactical Mission	DCI and SECDEF (Action Agents)	Utilize U.S. defense and intelligence agencies to support target acquisition and tracking of foreign drug activities
USFIB (Task Agent) USAF (Action Agent)	Collect, analyze, and disseminate photographic reconnaissance information; based on USFIB and foreign reconnaissance targets	U.S./DoD (Task Agent) USN (Action Agent)	Conduct tactical reconnaissance and enemy engagement over the "no-fly zone" with naval fighter bomber squadrons based on U.S. carriers	DOJ (Task Agent) DEA (Action Agent)	Carry out operations to decrease the quantity of drugs entering the U.S.
Mission Areas	Imagery for tactical flight planning Analyze data (only) for drug interdiction targeting and tracking	Mission Area	Perform flight planning function	Mission Area	Acquire, disseminate and maintain database of drug movement activities(some information acquired from U.S. Intelligence sources)

Table 9-2. Intelligence Mission Area Positions and Responsibilities

Position	Responsibilities
Mission Control	• Manages overall intelligence mission area • Controls receipt of collection, analysis and reporting requests from higher authority or other mission areas (e.g., tactical) • Transmits requests to other intelligence mission area positions
Satellite Uplink	• Controls receipt of collection request parameters from Mission Control • Transmits requests to satellite
Satellite Downlink	• Controls receipt of raw intelligence data from the satellite • Prepares, packages, and transfers satellite information to Collection Post-Processing and Production Control position
Collection Post-Processing and Production Control	• Separates raw intelligence data into appropriate collection transaction packages (collection request/response pair) • Performs quality-control checks on each transaction package • Transfers transaction package to appropriate (mission-specific) Intelligence Analysis and Reporting work center
Intelligence Analysis Work Center (includes supervisory and analyst positions)	• Performs analysis operation for a specific mission (e.g., tactical, drug interdiction) • Supervisory position disseminates transaction packages to appropriate analysts; controls release of intelligence reports • Analyst position performs data analysis and produces intelligence reports
Security Management (includes security officer and security administrator positions)	• Security officer position (designated by Mission Control) maintains the security policy, including modifications in the event of changing mission requirements • Security administrator position (determined by the security officer) maintains rules of enforcement and security attributes for intelligence mission function

9.1.2 Tactical Mission Area Function Description

The tactical mission area function generates flight plans that enable U.S. and coalition force aircraft to avoid ground-based fire. All information contained in this tactical mission area function is classified Secret (S). The flight plans are generated using geographical information, information developed and maintained within the tactical community, and downgraded imagery information from the intelligence mission function area. This mission area function also has a requirement to be able to make requests directly to the intelligence reconnaissance mission area function.

The major responsibilities within the tactical flight planning mission include flight planning, flight planning administration, and security management. Table 9-3

provides a description of the responsibilities associated with each of the tactical mission area positions.

Table 9-3. Tactical Mission Area Positions and Responsibilities

Position	Responsibilities
Flight Planning	• Develops air sortie flight plans that identify specific flight lanes over the "no-fly zone" • Identifies enemy weapon coordinates and types of weapons
Flight Planning Administration	• Generates and transmits requests to intelligence mission area for collection and analysis operations • Controls receipt of intelligence information
Security Management (includes security officer and security administrator positions)	• Security officer position (designated by tactical mission commanding officer) maintains the security policy and any modifications in the event of changing mission requirements • Security administrator position (determined by the security officer) maintains rules of enforcement and security attributes for tactical flight planning function

The individuals who fill these positions may have different privileges assigned to them to carry out their responsibilities. All of the individuals associated with the tactical mission function have at least a Secret clearance. There is no compartmentalization required within this mission area operation. All participants have a need-to-know authorization for all mission area information.

9.1.3 Drug Interdiction Mission Area Function Description

The drug interdiction mission area function has been formed to promote sharing of information previously maintained separately by various agencies. (Non-automated sharing of some of this information may have taken place in the past.) Providers and users of this information include the Drug Enforcement Administration (DEA), Federal Bureau of Investigation (FBI), Secret Service, Customs Service, Department of Justice (DOJ), other U.S. law enforcement elements (federal, state, local), and counternarcotics elements of various foreign countries. The DOJ oversees all drug interdiction activities with the assistance of the Department of State (DOS) for foreign law enforcement coordination activities.
The shared information contains the locations of drug production and processing facilities, and the identification and tracking of drug shipment activities. The shared information for the drug interdiction mission area function will be referred to as the *drug interdiction database*. The responsibility for management of the drug interdiction

database has been assigned to the DEA. (This database is conceptual and might be implemented in a single LSE or it might be distributed.) In this example, the focus will be on the functions performed by the DEA and DOJ.

The DEA maintains large quantities of data relevant to the drug interdiction mission. Some of this data is closely held by DEA, while some is shared with other U.S. and foreign government law enforcement agencies. Specifically, the DEA maintains a database (herein referred to as the *DEA-Only database*) containing locations where foreign drug production facilities reside, farmlands where illicit crops are grown, and specific target and tracking information of bulk drug movement activities. Access to this information is restricted to a select group of DEA employees. Portions of the DEA database are based on downgraded intelligence reports obtained from the intelligence mission. The DEA-Only database does not receive imagery data from the intelligence mission, but rather receives reports on new and known targets developed by intelligence analysts from the imagery data.

The DOJ maintains a set of files that contains investigative and case history information related to drug interdiction. Only select members of the DOJ have access to this information. Portions of the drug interdiction database are generated by the DOJ based on this set of files.

The other members of the drug interdiction community also maintain their own databases, accessible only by members of their particular organization, from which contributions are made to the drug interdiction database. However, these databases are not considered in this example.

The specific responsibilities within the drug interdiction mission are assigned to database administrators and database users for each of the databases considered above (DEA-Only, DOJ-Only, and drug interdiction), and security managers. Table 9-4 provides a description of the responsibilities associated with each of the drug interdiction mission area positions.

The individuals who fill these positions may have different privileges assigned to them to carry out their responsibilities. All participants have a need-to-know authorization for all mission area information, but need not hold a (U.S.) security clearance. The sensitivities and labels of information in the drug interdiction mission area are discussed in the next section.

Table 9-4 Drug Interdiction Mission Area Positions and Responsibilities

Position	Responsibilities
DEA-Only Database Administrator	• Maintains DEA-Only database (e.g., updates, changes) sdfg • Controls receipt of downgraded intelligence reports from the Intelligence mission area • Filters downgraded intelligence reports and transmits to Drug Interdiction Administrator position • Transmits select DEA information (e.g., farm lands) to Drug Interdiction Administrator position
DEA-Only Database User	• Queries and updates the database
DOJ-Only Database Administrator	• Maintains DOJ-Only database (e.g., updates, changes) • Transmits select DOJ case files to Drug Interdiction Administrator position
DOJ-Only Database User	• Queries and updates the database
Drug Interdiction Database Administrator	• Maintains Drug Interdiction database (e.g., updates, changes) • Controls receipt of downgraded intelligence reports from DEA-Only Database Administrator • Controls receipt of DEA information (e.g., farm lands) from DEA-Only Database Administrator • Control receipt of DOJ case files from DOJ-Only Database Administrator
Drug Interdiction Database User	• Queries and updates the database
Security Administrators (one per agency)	• Maintains the security policy and any modifications in the event of changing mission requirements • Maintains rules of enforcement and security attributes for drug interdiction mission function

9.2 MISSION AREA INFORMATION DOMAINS

Based on the description of the mission areas, their relationships, their information object sensitivities, and their functional requirement descriptions, decisions have been made about the information domains needed to support the mission functions in this example. The sensitivity of the intelligence information precludes bringing members of the other missions into an existing intelligence information domain. Since the tactical and drug interdiction missions are the only users of intelligence information filtered for their specific purposes, creating new information domains for only the filtered information is unnecessary. For the information to be shared by the several agencies involved in the drug interdiction mission, it is appropriate to create an information domain for that purpose. In this case, there is a need to merge

information from several source agencies to be used by some members of those agencies. Each of the other information domains described below existed before the requirements to cooperate were put in place:

1. The *Intelligence information domain* contains Top Secret photographic reconnaissance information objects. All users hold at least a Top Secret clearance and have a common need-to-know authorization regarding all aspects of this mission area function.

2. The *Intelligence Security Management information domain* is the security management information domain associated with the Intelligence information domain. It provides security management application and administrator access to security management objects.

3. The *Tactical information domain* contains Secret flight-planning information objects (including downgraded imagery from the intelligence domain). All users hold at least a Secret clearance and have a common need-to-know regarding all aspects of this mission area function.

4. The *Tactical Security Management information domain* is the security management information domain associated with the Tactical information domain. it provides security management application and administrator access to security management objects.

5. The *DEA-Only information domain* contains Unclassified but Sensitive information objects specifically associated with the drug interdiction mission (including downgraded intelligence reports from the intelligence domain). All users have a common need-to-know authorization regarding all aspects of this mission area function. The label for this information domain is *N/DI* (i.e., Unclassified but Sensitive/Drug Interdiction).

6. The *DOJ-Only information domain* contains information objects consisting of investigative files and active case histories, as well as DOJ security management information (embedded). All users have a common need-to-know authorization regarding all aspects of this mission area function. The sensitivity of the information in this domain is outside the national security classification system. It is labeled *DOJ-Only*.

7. The *Drug Interdiction information domain* contains information objects (including some information from the DEA-Only and DOJ-Only information domains). All users (both U.S. and foreign) have a common

need-to-know authorization regarding all aspects of this mission area function. The label for this information domain is *DI*.

8. The *DEA/DI Security Management information domain* serves both the DEA-Only and Drug Interdiction information domains. It existed previously for the DEA-Only information domain, and will now also provide security management application and administrator access to security management objects for the Drug Interdiction information domain.

Also included is the public information domain containing information used in support of the missions. The public information domain is a general-purpose domain in which there is no characterization of information sensitivity. Use of the public information domain infers all of the vulnerabilities associated with public information systems (e.g., MCI Mail, Prodigy, Compuserve). All of the missions cited in this example have access to the public domain.

Table 9-5 summarizes the mission area information domain membership and requirements. Figure 9-2 illustrates the example mission area information domain interrelationships. This example has been designed to include a shared domain (the drug interdiction information domain) and to convey why and how such an information domain can be created. Many missions will not need shared information domains (i.e., it is sometimes more reasonable simply to assign new members to existing domains). In this example, the need-to-know isolation requirements provide the impetus for creating the shared information domain.

The interdomain transfers (shown by directional lines in figure 9-2) identified for this example include:

- Bi-directional transfers between:
  - Intelligence and Tactical information domains
- One way transfers:
  - From the Intelligence information domain to the DEA-Only information domain
  - From the DEA-Only information domain to the Drug Interdiction information domain
  - From the DOJ-Only information domain to the Drug Interdiction information domain

Table 9-5. Summary of Mission Area Information Domains

Domain Number	Mission Area Function	Mission Requirements	
		Users/ Positions	Data
1	Intelligence	Mission Control Users, Satellite Uplink/Downlink Users, Collection Post-Processing and Production Control Users, Intelligence Analysis Users	Intelligence Parameters, ID Number, Time Stamp, Photographic Reconnaissance Images
2	Intelligence Security Management	Security Officer, Security Administrator	Security Management information (relevant to the Intelligence mission)
3	Tactical	Flight Planning Users, Flight Planning Administrators	Tactical Flight information (includes downgraded imagery)
4	Tactical Security Management	Security Officer, Security Administrator	Security Management information (relevant to the Tactical mission)
5	DEA-Only	Database Administrator, Database Users	Drug Interdiction Targets and Movement Tracking (includes reports from Intelligence domain)
6	DOJ-Only	Database Administrator, Database Users, Security Administrator	Ongoing DOJ Investigative and Active Case History Information
7	Drug Interdiction (shared by law enforcement agencies)	Database Administrator, Database Users	Drug Interdiction Targets and Movement Tracking information (includes information from DEA-Only domain); Drug Interdiction- Specific Investigative and Active Case History Information (includes files from DOJ-Only domain)
8	DEA/DI Security Management	Security Administrator	Security Management information (relevant to the DEA and Drug Interdiction missions)
9	Public Domain	All users and/or administrators	Data available through public domain (e.g., wire services, commercial servers, and bulletin boards)

The example also illustrates the three ways that security management information domains (shown with dashed lines in figure 9-2) can be structured: as separate security management domains for each information domain, as "composite" security

management domains that serve more than one information domain, and embedded security management information within an information domain itself.

The first is the simplest and most straightforward method because it keeps all security management information (i.e., the SMIB) isolated from other information domain objects. The second approach reduces the number of security management domains, but implies that the security managers of each of the information domains share functions and information or that one security manager is responsible for multiple information domains. In the final approach, the security management information is treated as any other information in the domain (e.g., all users have access to the SMIB).

9.3 INFORMATION DOMAIN SECURITY POLICIES

All information domains require a security policy. Members of the information domain are responsible for the development of the security policy. An information domain security policy should identify the following types of information:

- Relationship of the information domain to the mission function areas
- Users (membership requirements) and user privileges
- Information objects
- Threats, operational requirements, and DISSP requirements
- Security services necessary to meet the threats and strengths needed in each security service to meet an acceptable risk posture for the information domain
- Interdomain transfer requirements (if applicable): import and export conditions (rules of transfer) for each information domain for which information object transfers are allowed. Rules of transfer must be explicit with regard to flow (one way or bi-directional), and whether sanitization is required (including explicit conditions of sanitization)
- Security management requirements (i.e., what approach, or combination of approaches, is required for security management). Alternatives include separate, shared, or embedded management domains. Rules for invocation of management functions by information domain members must be stated.

- Contingency planning requirements

Each of the existing information domains in this example had a security policy before the requirements for interdomain transfers were added. Threat analyses were performed taking the interdomain transfers into account and, where necessary, the information domain security policies were modified. For the newly formed Drug Interdiction information domain, a security policy was created. The security policies presented for the example information domains are the final policies.

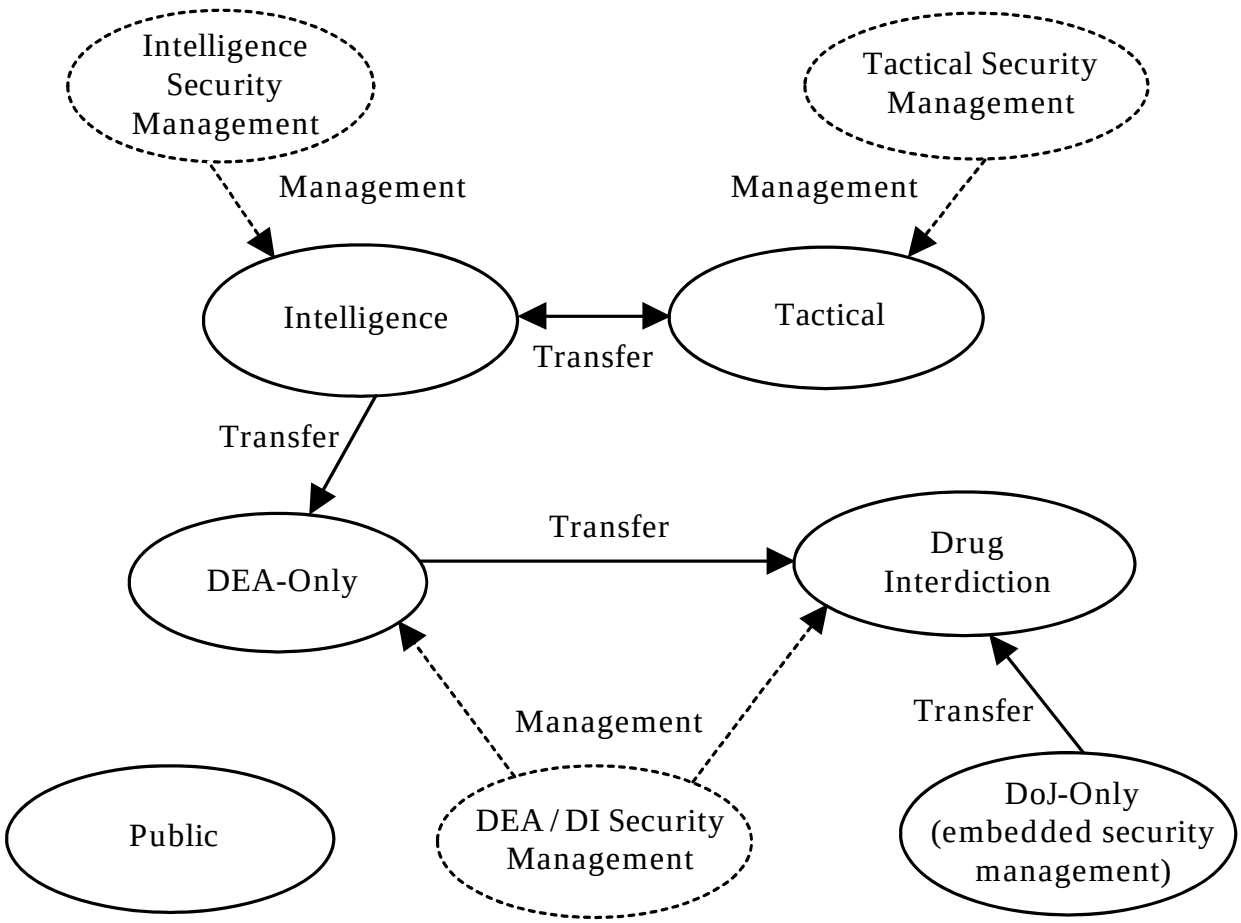

Figure 9-2. Example Mission Area Information Domain Relationships

9.3.1 Intelligence Mission Function Information Domains

The threat analysis performed when the requirements to support the tactical and drug interdiction missions were added resulted in modifications to the Intelligence information domain security policy. Each interdomain transfer requirement was addressed from a threat and risk perspective, and then the aggregate of the interdomain transfer additions was examined for threat and risk, before adding the necessary security requirements to the Intelligence information domain security policy.

9.3.1.1 Threats Related to the Tactical Mission Interdomain Transfer Requirement

The threats to the Intelligence information domain as a result of adding an interdomain transfer to the Tactical mission, specifically to the Tactical information domain, relate to disclosure. The threats to the Tactical information domain, in turn, relate to deception and availability (e.g., denial of service). These threats must be considered when making modifications to the information domain security policies.

The threat of disclosure to the Intelligence information domain is from intentional or accidental release of unauthorized information during downgrading from the Intelligence information domain to the Tactical information domain. The accidental release of information, primarily through human carelessness, is considered to be of moderate risk. The intentional release of information is considered to be of moderate risk primarily since personnel are assumed trustworthy because they hold high-level security clearances. The low-to-moderate risk of intentional disclosure is concerned with source and method capabilities. The security requirements developed to lower these risks (to an acceptable level) are applied to the imagery and associated parameter data to be transferred to the tactical flight planning information domain as follows:

- An automated sanitization (prohibited word scan) prior to transfer (i.e., to check for source or method indicators, or improper classification labels)

- Image resolution editing ("blurring") invoked by the intelligence analyst prior to transfer to reduce resolution to an acceptable level (to hide real resolution capability)

- Recording information on all interdomain transfer operations in an audit log as evidence of possible unauthorized disclosure (accessible only by a security officer or delegate); recorded information includes the intelligence

analyst's identification, and the date-time group the transfer operation occurred

The Tactical information domain threats of deception and denial of service must be addressed in both the Intelligence information domain and the Tactical information domain (i.e., both domains must share the responsibility for abating risk in these threat areas).

The deception threat includes unauthorized imagery or parameter data modifications by intelligence analysts or flight planners intent on endangering air sorties by providing incorrect enemy location information. This threat has a low-to-moderate risk in the Intelligence information domain because of the special background investigation associated with the level of security clearance held by the intelligence analysts. In the Tactical information domain, the risk is considered to be moderate because the flight planners hold only Secret security clearances. The risk is considered significant enough that the threat must be countered in both information domains. Reliance on the accuracy and integrity of the imagery data is critical to the safe operation of the "no-fly zone" tactical air sorties. The security requirements for the Intelligence information domain developed to lower the risks associated with the deception threat are as follows:

- The transfer of imagery information and associated parameters to the Tactical information domain shall be subjected to a two-party control procedure. That is, an intelligence analyst will filter the information intended for transfer to the Tactical information domain and then another intelligence analyst will review and confirm that the filtering has been performed properly before invoking the transfer function.

- A data integrity seal shall be placed on the released information so that modification of the information can not be made without detection.

- The information released to the Tactical information domain shall be signed (using a high-grade signature algorithm) by both analysts. All release activities will be audited.

The non-availability of critical flight planning imagery information from the Intelligence information domain is considered a high-risk threat to the tactical mission that must be abated in both domains. The security requirement applied to both information domain policies is that the collection, processing, analysis, and communications components of both domains shall include availability assurance measures to provide operational accessibility ninety-nine percent of the time.

In addition, only certain named individuals in the Tactical information domain shall have the authority to prepare collection requests (signed using a high-grade signature algorithm) intended for the Intelligence information domain. The transfer of such requests can only be carried out by a member of the Tactical information domain who is also a member of the Intelligence information domain.

9.3.1.2 Threats Related to the Drug Interdiction Mission Interdomain Transfer Requirement

The threat to the Intelligence information domain associated with adding an interdomain transfer to the drug interdiction mission, specifically to the DEA-Only information domain, is disclosure and is considered a high-risk threat. No threats to the DEA-Only information domain resulting from the interdomain transfer carry an unacceptable risk. The downgrade process will be a two-party control procedure which is the same as that used between the Intelligence and Tactical information domains. The security requirements developed to lower the risk associated with the disclosure threat are stated as follows:

- Only selected text portions of the targeting and tracking reports shall be transferred to the DEA-Only information domain. None of the imagery information in these reports may be transferred.

- The transfer of reports to the DEA-Only information domain shall be subjected to:

  - Automated sanitization checks prior to transfer (e.g., prohibited word scan to check for presence of source or method indicators, or ensuring that the correct label has been applied to the sanitized information)

  - Audit of all interdomain transfer operations as evidence of possible unauthorized disclosure (accessible only by a security officer or delegate); recorded information includes the intelligence analyst's identification and the date-time group the transfer operation occurred

- The transfer of intelligence reports to the DEA-Only information domain shall be subjected to a two-party control procedure.

Information objects will not be imported from the DEA-Only information domain to the Intelligence information domain. All collection requirements in support of the drug interdiction mission will be coordinated and processed by the intelligence

mission function from an external source authorized to task collection assignments (e.g., the U.S. Foreign Intelligence Board).

9.3.1.3 Intelligence Information Domain Security Policy

Table 9-6 summarizes the security policy for the Intelligence information domain, including the interdomain transfer requirements to the tactical and drug interdiction mission areas. Specific DISSP requirements were stated in section 2 of this document.

9.3.1.4 Intelligence Security Management Information Domain Security Policy

The Intelligence information domain shall have a separate Intelligence Security Management information domain Table 9-7 reflects the security policy of the Intelligence Security Management information domain.

9.3.2 Tactical Mission Function Information Domains

In the Tactical information domain, flight plans will be generated using the imagery and associated parameter data from the Intelligence domain, and geographical and other information developed and maintained within the tactical community. The users in this domain will forward collection tasking requests to the Intelligence information domain. Only designated users in the Tactical information domain may generate and communicate collection requirements. Collection requirements will be prepared as fixed-format messages, which must be signed by the authorized releaser of the message.

The threats to the Tactical information domain resulting from interdomain transfer requirements were discussed in 9.3.1.

Like the intelligence mission area, the Tactical information domain shall have a separate security management information domain.

9.3.2.1 Tactical Information Domain Security Policy

Using the Intelligence information domain security policy as a template, table 9-8 summarizes the security policy for the Tactical information domain, including the interdomain transfer requirements to the intelligence mission area. Specific DISSP requirements were stated in section 2 of this document.

Table 9-6. Intelligence Information Domain Security Policy

Membership Requirements

All members must fill one or more of the mission operational positions and must hold a TS clearance.

Information Objects

Information objects shall include mission control parameters for collection management, satellite control parameters, raw collection imagery with time stamp and coordinates, intelligence analyst reports, and end product imagery.

Operational Requirements

All members shall have read, write, and delete privileges. All write or delete operations will be audited. All interdomain transfer operations will be audited.

Security Services and Strengths Required

Authentication: Authentication shall be required for all users and end (or relay) systems that access or support this information domain. The strength shall be basic (e.g., name and password) for direct user access to end systems, strong (e.g., signed X.509 certificate) for remote user access, and strong (e.g., cryptographically bound) for communicating end systems. Physical authentication will be required for personnel at facilities that process, store, and communicate information in this information domain.

Access Control: Access control shall be based on strict isolation of information domains and on the ability to restrict which users will be allowed to perform interdomain transfers. Physical access controls shall be employed to adequately protect the facility, commensurate with national and DoD doctrine for TS information. All personnel authorized to access the facility who do not have the requisite security clearance shall be continuously escorted while in the facility.

Data Integrity: Data integrity shall be provided for all information domain objects and process resources within local processing environments. Interdomain transfers shall provide data integrity that is strong enough to guarantee that the object sent is precisely the object received. All application software that operates on the system shall be under strict configuration management control. Users will not be permitted to enter unauthorized software into the system, or modify existing software.

Confidentiality: Confidentiality shall be provided for all communications. When employing communication networks that are not physically secured, cryptographic mechanisms shall be employed. The strength of cryptographic algorithm and key management strategy shall be consistent with national policy and doctrine. Personnel entering and leaving the facility shall be subject to search.

Non-Repudiation: Non-repudiation shall be provided for all intradomain staged delivery (messaging) service communications. Proof-of-origin non-repudiation service shall be mandatory. Proof-of-delivery non-repudiation service (i.e., a signed receipt) will be optional.

Audit: All security-relevant events shall be audited, including user login attempts, all information object write and delete operations, all interdomain transfers (also archived in conjunction with the audit log for a period of not less than 30 days), failed security association attempts, cryptographic alarms and faults, and signature verification failures. The audit log shall be reviewed on a daily basis. Significant audit events shall be archived and retained for a period of not less than one year.

Availability: Critical system resources and all communications shall have 99 percent availability for the intelligence and tactical mission domains, and 90 percent for the drug interdiction mission domain.

Table 9-6. (Concluded)

Security Services and Strengths Required (continued)
Interdomain Transfer Requirements
Tactical Information Domain Transfers: Exports shall be subject to automated sanitization, reduction of image resolution, audit, two-party release, release authority signature, and non-alterable integrity seal applied to a released information object. Imports shall be subject to tactical domain release authority signature check, integrity check of received information object, information object upgrade to TS. **DEA-Only Information Domain Transfers:** Exports shall be subject to automated sanitization scanning (text only, no imagery), audit, two-party release, and release authority signature. No automated imports shall be allowed. All collection requirements shall enter this domain through mission control external input sources only.
Security Management Requirements The Intelligence information domain shall have a dedicated information domain for security management (see table 9-7).
Contingency Plan If a threat of intention related to disclosure is discovered or strongly suspected between the Intelligence information domain and the Drug Interdiction information domain, this capability shall be disabled until a full investigation has been completed and resolved.

Table 9-7. Intelligence Security Management Information Domain Security Policy

Membership Requirements All members of this information domain shall have at least a TS clearance. Security officers shall maintain security policy and security administrators shall maintain rules of enforcement and security attributes for intelligence mission function.
Information Objects Information objects are the SMIB objects associated with the Intelligence information domain.
Operational Requirements All members shall have read, write, and delete privileges. All write or delete operations for all SMIB objects shall be audited. All end systems which support this information domain shall provide a SMAP interface to the SMIB. For end systems that support this and other information domains, an end system SMAP shall be provided to manage system security-critical functions for this security policy in isolation from other information domains.
Security Services and Strengths Required Same as Intelligence information domain.
Interdomain Transfer Requirements None.
Contingency Plan In wartime conditions, the audit resources shall be disabled to increase performance.

Table 9-8. Tactical Information Domain Security Policy

Membership Requirements All members of this information domain must hold at least a Secret clearance and be authorized by the Commanding Officer to have access to this domain.[9-1]
Information Objects Geographical, imagery, and other information maintained within the tactical community to aid in the successful planning of air sorties.
Operational Requirements All members have read, write, and delete privileges. All write or delete operations are audited. All interdomain transfer operations are audited. Mobile operation must be support . The environment is assumed hostile.
Security Services and Strengths Required Security services and strengths required are the same as in the Intelligence information domain security policy (see Table 9-6) except that security labels are Secret rather than TS.
Inter-Domain Transfer Requirements **Intelligence Domain Transfers:** Exports are subject to message signature for collection requirements messages. Only authorized tactical domain members are permitted to generate and forward collection requirement messages.[9-2]
Security Management Requirements The tactical information domain shall have a dedicated information domain for security management. (Uses same approach Intelligence community.
Contingency Plan Alternative methods for obtaining the intelligence information will be used. Specifically, redundant equipment will be maintained shipboard for satellite data acquisition.

9.3.2.2 Tactical Security Management Information Domain Security Policy

The domain security policy for the Tactical Security Management information management domain is similar to the Intelligence Security Management information domain (see section 9.3.1.4) with the exception that the classification of objects and the clearance level required are Secret rather than Top Secret.

[9-1] Members include those Intelligence information domain members with release authority to the Tactical information domain.

[9-2] A member of the Intelligence information domain who is a member of the Tactical information domain must perform the transfer between the two information domains.

9.3.3 Drug Interdiction Mission Function Information Domains

The information domains in this example that are considered part of the drug interdiction mission function are the DEA-Only information domain, the DOJ-Only information domain, the Drug Interdiction information domain, and the DEA/DI Security Management information domain.

No new threats were judged to exist as a result of the requirements for interdomain transfers with the Drug Interdiction information domain.

9.3.3.1 DEA-Only Information Domain Security Policy

Table 9-9 summarizes the security policy for the DEA-Only information domain, including the interdomain transfer requirements to the Intelligence information domain and the Drug Interdiction information domain. Specific DISSP requirements were stated in section 2 of this document.

9.3.3.2 DOJ-Only Information Domain Security Policy

Table 9-10 summarizes the security policy for the DOJ-Only information domain including the interdomain transfer requirements to the Drug Interdiction information domain. Specific DISSP requirements were stated in section 2 of this document. Security management requirements will be satisfied internally for the DOJ-Only information domain. That is, no separate dedicated management information domain has been defined. No distinction exists between security management information objects and DOJ case history information objects.

9.3.3.3 Drug Interdiction Information Domain Security Policy

Table 9-11 summarizes the security policy for the Drug Interdiction information domain, including the interdomain transfer requirements from the DEA-Only and DOJ-Only information domains. Specific DISSP requirements were stated in section 2 of this document. Security management requirements will be satisfied for the Drug Interdiction information domain by the DEA/DI security management information domain.

9.3.3.4 DEA/ DI Security Management Information Domain Security Policy

Table 9-12 summarizes the security policy for the DEA/DI Security Management information domain. Specific DISSP requirements were stated in section 2 of this document.

Table 9-9. DEA-Only Information Domain Security Policy

Membership Requirements All members of this information domain must be employees of the DEA and authorized to receive downgraded information from the Intelligence information domain.[9-3] No security clearance shall be required, but a need-to-know authorization shall be required. There are two positions defined for the DEA-Only information domain: database administrators and database users (Intelligence information domain members shall be considered database administrators). The database administrators shall maintain the DEA-Only database and shall be responsible for the receipt and filtering of downgrade intelligence information. The database users shall be permitted to query and update the database.
Information Objects Information objects shall include locations of foreign drug production facilities, farmlands where illicit crops are grown, specific target and tracking information, and downgraded intelligence reports.
Operational Requirements All members shall have read, write, and delete privileges. Only database administrators shall have the ability to participate in an interdomain transfer. All write or delete operations shall be audited. All interdomain transfer operations shall be audited.
Security Services and Strengths Required Security services and strengths required shall be the same as in the Intelligence information domain security policy (see table 9-6) except that security labels shall be N/DI and availability of critical system resources shall be 90 percent.
Interdomain Transfer Requirements **Intelligence Information Domain Transfers:** Automated exports shall not be allowed to be received by the Intelligence information domain. For imports, information objects must be signed by an authorized releasing authority. The source of this information shall not be disclosed outside of the DEA-Only information domain. **Drug Interdiction Information Domain Transfers:** Exports shall be subject to review by authorized DEA personnel, and all information objects shall be signed by a releasing authority. No automated imports shall be allowed. All collection requests shall enter this domain through external sources only.
Security Management Requirements The DEA-Only information domain shall share an information domain for security management with the Drug Interdiction information domain (see table 9-12).
Contingency Plan If the information transferred from the Intelligence domain is not available, the DEA-Only information domain will continue to transfer drug interdiction-related information to the Drug Interdiction information domain.

[9-3] Members will include those Intelligence information domain members with release authority to the DEA-only information domain. They shall be considered database administrators.

Table 9-10. DOJ-Only Information Domain Security Policy

Membership Requirements All members of this information domain shall be employees of the DOJ. No security clearance shall be required, but a need-to-know authorization shall be required. Two positions shall be defined for the DOJ-Only information domain: database administrators and database users. The database administrators shall maintain the DOJ-Only database and determine which information may be transferred to the Drug Interdiction information domain. The database users shall be permitted to query and update the database.
Information Objects DOJ files that contain investigative and case history information.
Operational Requirements All members shall have read, write, and delete privileges. All write or delete operations and interdomain transfer operations shall be audited.
Security Services and Strengths Required **Authentication:** Authentication shall be required of all users and end or relay systems that have access to or support this information domain. The strength shall be basic (e.g., name and password) for user access and for communicating end and relay systems. Physical authentication shall be required for all personnel entering a facility that processes, stores, and communicates information in this information domain. **Access Control:** Access control shall be based on strict isolation of information domains and on the ability to restrict which users are allowed to perform interdomain transfers. **Data Integrity:** Data integrity shall be provided for all information domain objects and process resources within local processing environments. All application software that operates on the system shall be under strict configuration management control. **Audit:** All transfers to the Drug Interdiction information domain shall be audited. The audit log shall be reviewed weekly. Significant audit events shall be archived for a period of one year. **Availability:** Critical system resources and communications shall have 90 percent availability.
Interdomain Transfer Requirements **Drug Interdiction Information Domain Transfers:** Exports shall be subject to review by authorized DOJ personnel, and all information objects shall be signed by a releasing authority. There shall be no automated imports. All collection requests shall enter this domain through external sources only.
Security Management Requirement All security management information shall be embedded within the DOJ-Only information domain. That is, all members of the DOJ-Only information domain shall have the same access to the security management information (i.e., the SMIB) as they do to other objects in that information domain.
Contingency Plan If interdomain transfers to the Drug Interdiction information domain are not operational, external means will be used to transfer information.

Table 9-11. Drug Interdiction Information Domain Security Policy

Membership Requirements
Members shall include employees of DEA, DOJ, FBI, Secret Service, Customs Service, other U.S. law enforcement elements (Federal, state, local), and counternarcotics elements of various foreign governments. No security clearance shall be required, but a need-to-know authorization shall be required. Two positions shall be defined for the Drug Interdiction information domain: database administrators and database users. The database administrators shall maintain the drug interdiction database. The database users shall be permitted to query and update database.

Information Objects
The Drug Interdiction database shall contain selected downgraded intelligence reports (without source indication), locations of foreign drug production facilities, farmlands where illicit crops are grown, select target and tracking information, and drug offender case history files and investigative reports .

Operational Requirements
All members shall have read, write, and delete privileges. All write or delete operations and interdomain transfer operations shall be audited.

Security Services and Strengths Required
Security services and strengths required shall be the same as in the DOJ-Only information domain security policy (see table 9-10).

Interdomain Transfer Requirements
DEA-Only Information Domain Transfers: Exports shall not be permitted. Imports are subject to DEA-Only release authority signature check.
DOJ-Only Information Domain Transfers: Exports shall not be permitted. Imports shall be subject to a DOJ-Only release authority signature check.

Security Management Requirements
The Drug Interdiction information domain shall share an information domain for security management with the DEA-Only information domain (see table 9-12).

Contingency Plan
If interdomain transfers to the Drug Interdiction information domain are not operational, external means will be used to obtain the information.

Table 9-12. DEA/ DI Security Management Information Domain Security Policy

Membership Requirements All members of this information domain must have a valid need-to-know authorization. Security administrators shall maintain the security policies and rules of enforcement for the drug interdiction mission function.
Information Objects The information objects are the SMIB objects associated with the DEA-Only and Drug Interdiction information domains.
Operational Requirements All members shall have read, write, and delete privileges. All write and delete operations shall be audited. All end systems that support this information domain shall provide a SMAP interface to the SMIB. For end systems that support this and other information domains, an end system SMAP shall be provided to manage system security-critical functions for this security policy in isolation from other information domains.
Security Services and Strengths Required Same as the DEA-Only information domain.
Interdomain Transfer Requirements None.
Contingency Plan When systems are overloaded, the audit resources will be disabled to increase performance.

9.4 LOCAL SUBSCRIBER ENVIRONMENTS

Once the information domain security policies have been defined, LSEs, CNs, and their components must be determined. An information domain can be supported by one or more end systems, which in turn may be considered parts of one or more LSEs or CNs. Identifying the LSEs and CNs that will support each information domain is the next step in the architecture development process. Subsequent steps include defining and updating LSE and CN security policies, and developing specific LSE and CN security architectures. Because the DGSA assumes that the only security service provided by CNs is availability, this example focuses on LSEs.

Three distinct organizations, functions, and funding sources lead to the distribution of systems along mission lines (i.e., three categories of LSEs, one for the intelligence, tactical, and drug interdiction mission areas, respectively). The photographic intelligence system exists within an LSE at an Air Force Base in the northeastern United States, the tactical flight planning system exists within an LSE aboard a U.S. aircraft carrier located off shore of the "no-fly" area, and the drug interdiction targeting and tracking system exists within an LSE at the DEA headquarters building in Washington, D.C. The LSE for the DOJ is also an important source for the drug

interdiction mission area and therefore requires examination. Many other drug interdiction mission LSEs support users of the drug interdiction targeting and tracking information, some of which are owned by foreign government LSEs. The example LSEs are designated as follows:

- LSE-A - Intelligence site
  - Air Force base in northeastern United States
  - Fixed-site LSE
- LSE-B - Tactical site
  - U.S. aircraft carrier off shore of "no-fly" area
  - Mobile LSE
- LSE-C - DEA site
  - DEA headquarters, Washington, D.C.
  - Fixed-site LSE
- LSE-D - DOJ site
  - DOJ headquarters, Washington, D.C.
  - Fixed-site LSE
- LSE-E - Representative of Drug Interdiction database user sites
  - Various locales worldwide (U.S. and foreign law enforcement agencies)
  - Fixed-site LSE

End systems within LSE-A communicate with end systems in LSE-B using voice and data applications, and with end systems in LSE-C using data applications only. End systems within LSE-C communicate with LSE-D and LSE-E using voice and data applications. Figure 9-3 illustrates the various LSEs, their locations, mission applications, and communications. The dashed lines indicate the logical flow of information between information domains, and the solid lines indicate communication paths.

All personnel within LSE-A have at least a Top Secret clearance, personnel at LSE-B have a range of clearances (some have no clearances and some have Secret clearances), and personnel in LSEs-C, D, and E do not hold U.S. government security clearances.

LSE-A supports the Intelligence, Tactical, and Drug Interdiction information domains (including associated security management information domains). LSE-B supports the Tactical and Intelligence (optional) information domains (including associated security management information domains). LSE-B also supports other information domains that have no relationship to any of the example missions; some of these information domains require users to have Secret clearances, while others do not require users to have security clearances. LSE-C supports the DEA-Only, Drug Interdiction, and composite DEA/DI Security Management information domains. LSE-D supports the DOJ-Only, Drug Interdiction, and DEA/DI Security Management information domains. LSE-E supports the Drug Interdiction information domain (including the DEA/DI security management information domain). In addition, all LSEs support the Public information domain. The information domains supported by the LSEs are summarized in table 9-13.

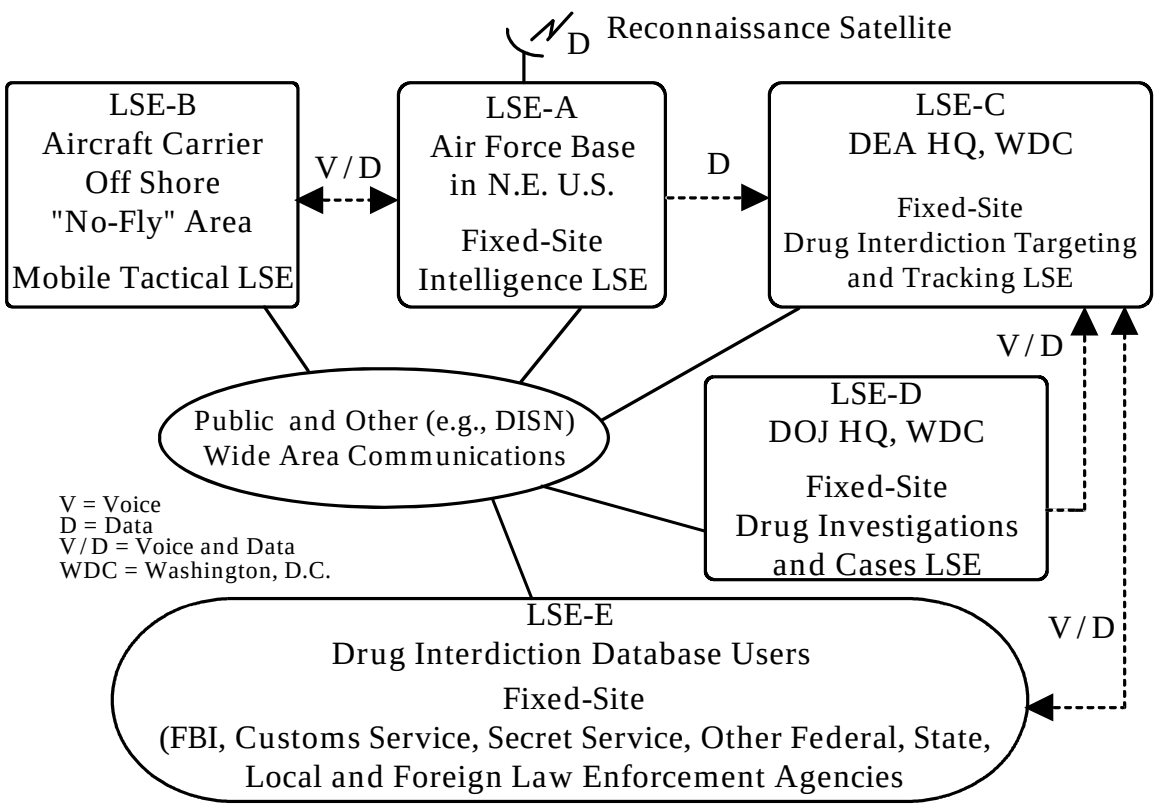

Figure 9-3. LSE Locations, Mission Applications, and CommunicationsA

Table 9-13. Domain Allocations by LSE

Domain ID Number	Domain Name	LSE-A Intelligence Site	LSE-B Tactical Site	LSE-C DEA Site	LSE-D DOJ Site	LSE-E DI User Site
1	Intelligence	x	(optional)			
2	Intelligence Security Management	x	(optional)			
3	Tactical	x	x			
4	Tactical Security Management	x	x			
5	DEA-Only	x		x		
6	DOJ-Only				x	
7	Drug Interdiction			x	x	x
8	DEA/DI Security Management	x		x	x	x
9	Public Domain	x	x	x	x	x

9.4.1 Intelligence Site LSE Architecture (LSE-A)

The intelligence site LSE (LSE-A) was described from a "position" perspective in table 9-2, and is illustrated by functional component in figure 9-4. This LSE has been configured so that new mission functions may be added and deleted in a short period of time. Intelligence analyst workstation (IAWS) end systems can be moved from one mission-specific support center to another (after sanitization), or they may be configured to support multiple missions. The reconnaissance satellite is also considered an end system of the intelligence site LSE. Optionally, if availability requirements demand, the satellite may also be an end system of the mobile tactical site LSE.

The intelligence site local communications system configuration is a network of Fiber Distributed Data Interface (FDDI) rings and Ethernet branching buses. The relay systems include routers using the stream transfer protocol and the Connectionless Network Layer Protocol (CLNP) (ISO, 1988a) in layer 3, X.400 Message Transfer Agent and X.500 Directory Service Agent components, an asynchronous transfer mode (ATM) switch (serving as a relay to a commercial Synchronous Optical Network (SONET)), and an optional commercial satellite communications relay system to support tactical mission areas. The stream transfer protocol is used for real-time voice or video applications. CLNP is used for interactive data communications applications requiring a connectionless layer 3 service. The ISDN Terminal Access Control (TAC) system is an end system with relay functions, as are

the Workcenter Control IAWSs. The transfer system for the intelligence site is depicted in figure 9-5.

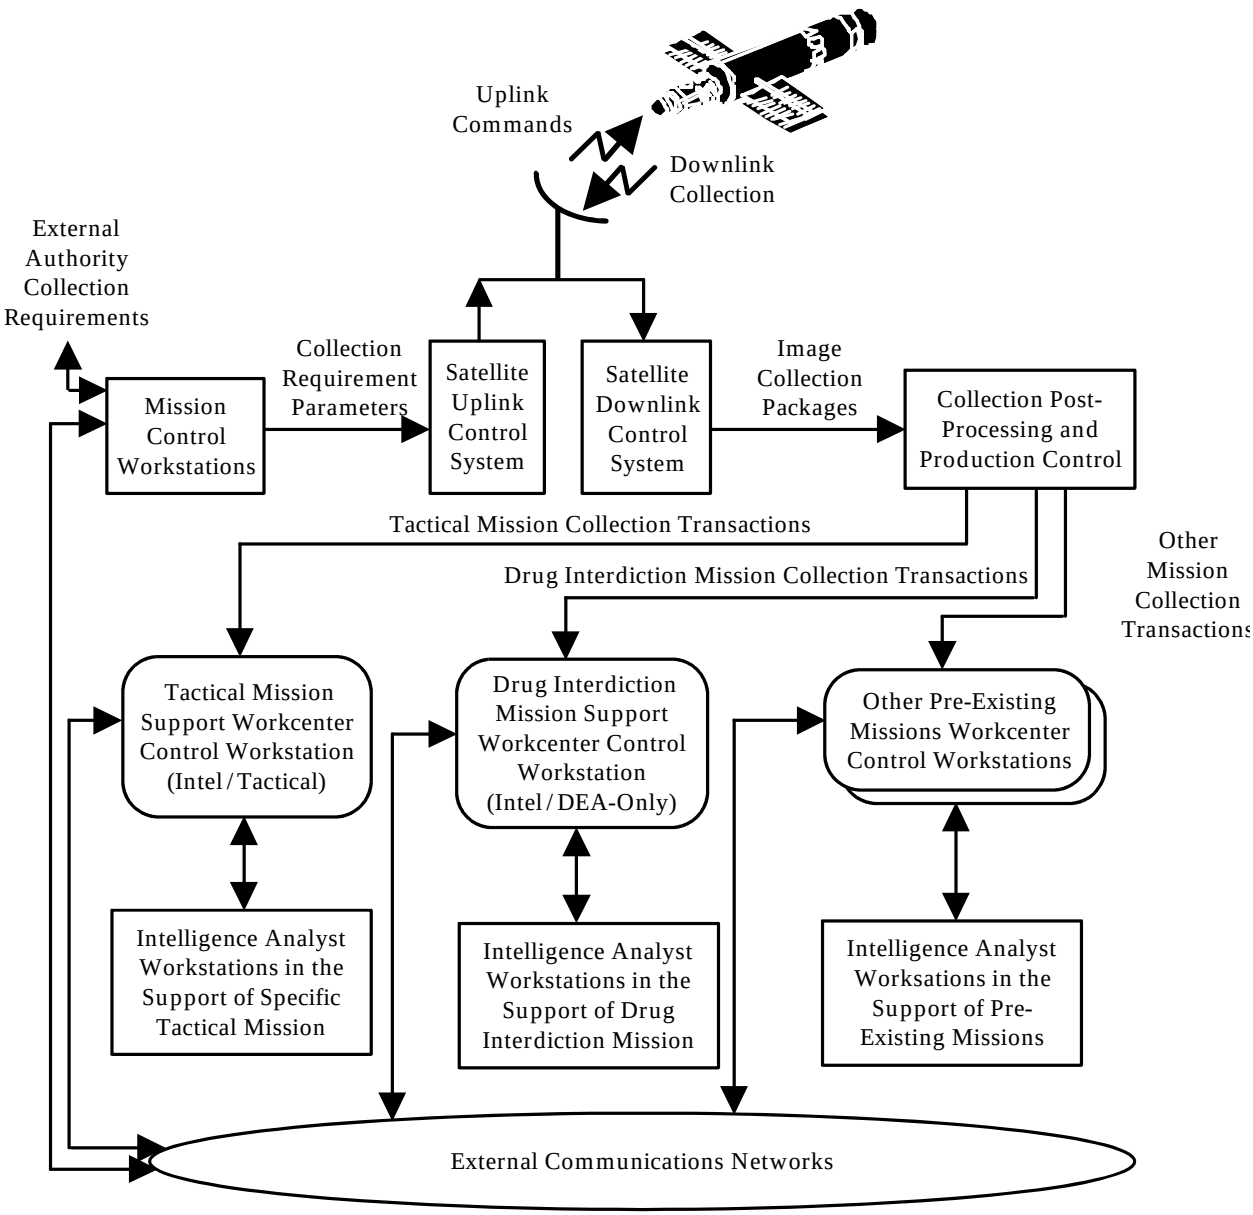

Figure 9-4. Intelligence Site LSE Functional Architecture (LSE-A)

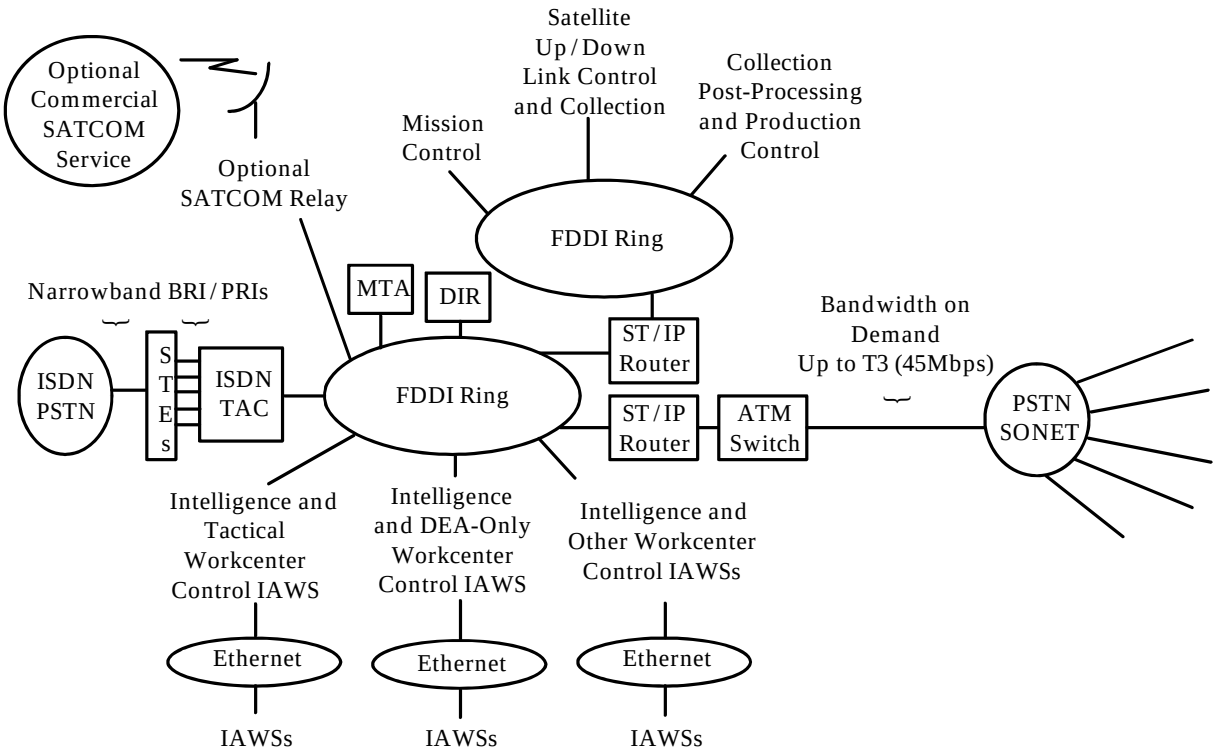

Figure 9-5. Intelligence Site Transfer System Architecture (LSE-A)

File transfers among intelligence site end systems use the ISO File Transfer and Access Management (FTAM) protocol (ISO, 1988b). Inter-mission and inter-site messaging services use the X.400 P3 (submission and delivery for the message transfer system) or P7 (submission and fetch for message stores) protocols. The Directory Access Protocol is used to support messaging. The Common Management Information Protocol (CMIP) is used by management entities to exchange information. The Association Control Service Element (ACSE) is used by all application layer protocols to form application associations (ISO, 1988a). Relay systems accommodate LAN interoperability with a range of wide area communications technologies (ISDN, ATM, Freespace (radio)). The security protocols that are supported by the intelligence site are illustrated in figure 9-6.

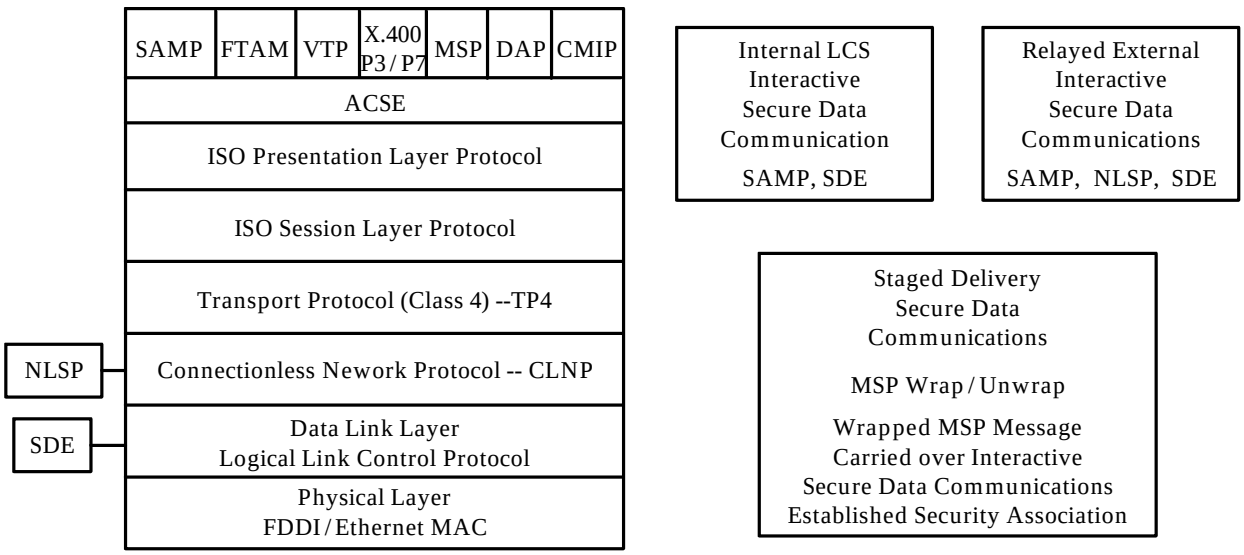

End/Relay System Security Protocol Allocations

Mission Control Workstations	-- SAMP, NLSP, SDE, MSP
Satellite Control & Collection	-- SAMP, SDE
Post-Collection Processing	-- SAMP, SDE
Intel Analysis Product Control	-- SAMP, SDE
Intel Analysis Work Center Control IAWS	-- SAMP, NLSP, SDE, MSP
Intel Analysis Work Center IAWSs	-- SAMP, SDE
Recon Satellite	-- SAMP, TLE (SAMP provides OTAR)
ISDN TAC	-- SAMP, SDE<-->SAMP
MTA Components	-- SAMP, NLSP
DIR Components	-- SAMP, NLSP
Routers: FDDI - ATM	-- SAMP, SDE<-->SAMP, ATM LE -- NLSP passes through end-to-end for interactive data communications -- real -ime voice carried over ST as an end-to-end pass-through
Commercial SATCOM Relay	-- SAMP, SDE<-->SAMP, link encryptors

Figure 9-6. Intelligence Site Security Protocols (LSE-A)

9.4.2 Tactical Site LSE Architecture (LSE-B)

The tactical site LSE (LSE-B) is composed of flight-planning workstations, a tactical flight-planning database server, a flight-planning data management workstation, and an all-source incoming tactical support information repository component.

Optionally, the tactical site LSE may implement intelligence domain collection and analysis components to augment the intelligence site (LSE-A) support to the tactical mission (i.e., to enhance availability and timeliness of intelligence support information). The functional architecture for LSE-B is illustrated in figure 9-7.

The tactical site transfer system functional architecture is composed of end system communications protocol stacks, local communications systems, and relay systems for intra-LCS and wide area communications connectivity. The tactical site transfer system functional architecture is illustrated in figure 9-8.

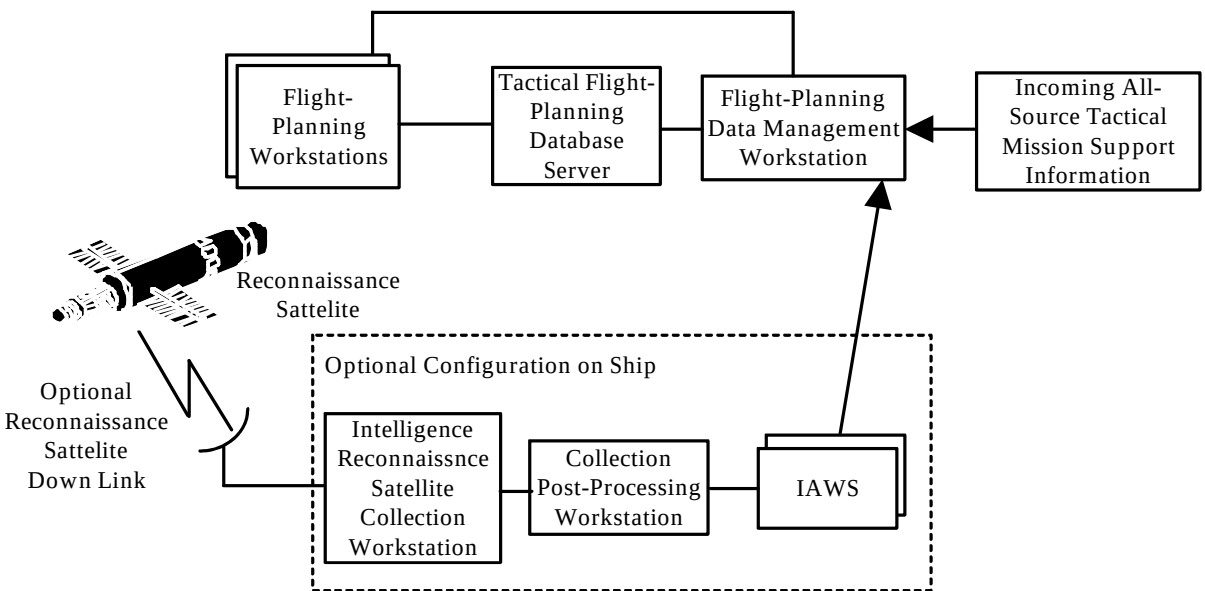

Figure 9-7. Tactical Site Functional Architecture (LSE-B)

The tactical site local communication system is composed of multiple FDDI rings, and Ethernets supporting specific mission support enclaves. Its relay systems are routers, protocol tunneling bridges, and X.400 MTA components. Wide area communications, at sea, are entirely freespace (extremely high frequency and super high frequency satellite communication systems and very high frequency/ultra high frequency ship-to-air communications systems). Wide area communications, while the ship is in port, are ISDN narrowband using the ISDN TAC end system.

Redundant FDDI rings function as the aircraft carrier backbone, and separate FDDI rings function as the radar, weapon systems, and bridge control local communications system. There are also various point-to-point local communication

system interfaces associated with critical radar, weapon system, and bridge control components, serve as fail-safe backup interfaces and communications mediums.

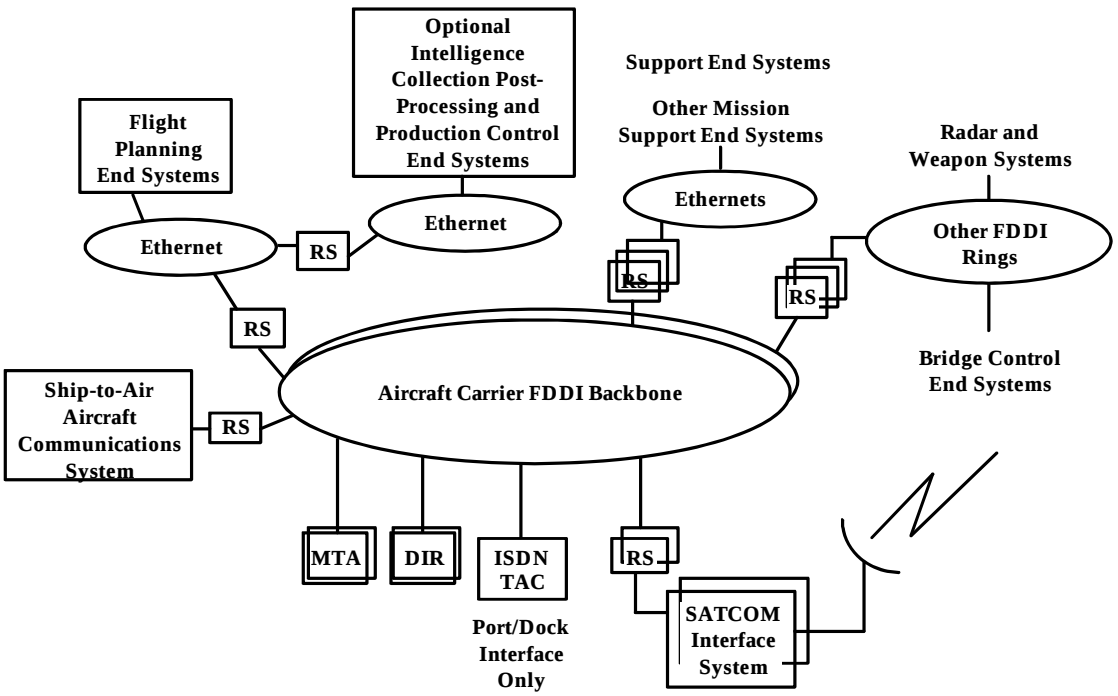

Figue 9-8. Tactical Site Transfer System Architecture (LSE-B)

End system and relay system protocol stacks for interactive data, staged delivery data, and interactive voice communications are the same as the intelligence site (see figure 9-6). LSE-B communicates with LSE-A using X.400 messaging and (optionally) voice applications.

The optional intelligence collection end system communicates directly with the reconnaissance satellite (downlink only). If this option is exercised onboard the aircraft carrier, the reconnaissance satellite is considered an end system of the tactical LSE (i.e., the satellite is a mobile end system of both the intelligence site and the tactical site).

9.4.3 DEA Site LSE Architecture (LSE-C)

The DEA site end systems include workstations and database servers. There are two database servers, one which is accessible to the general drug interdiction community, and the other which is accessible to only the DEA. DEA workstations receive intelligence reports which then are formatted as DEA-Only database updates. In this example, only DEA workstations and DOJ workstations are permitted to make updates to the general drug interdiction database. There is no direct flow of information between the database servers. The functional configuration in support of the DEA drug interdiction mission function is illustrated in figure 9-9.

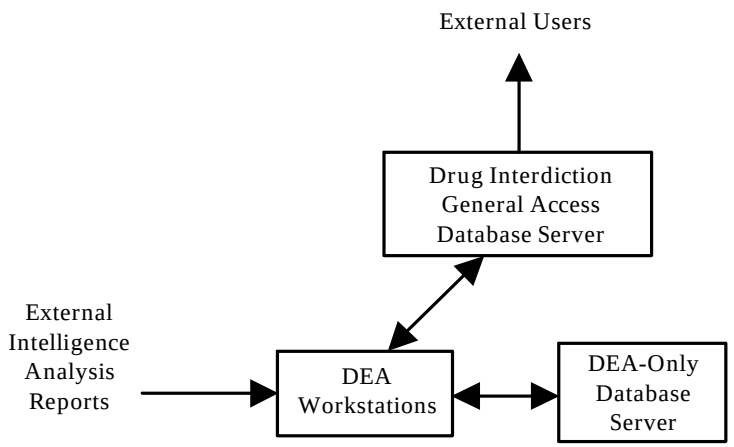

Figure 9-9. DEA Site Functional Architecture (LSE-C)

The DEA site transfer system is composed of Ethernet and ISDN backbones at DEA headquarters. Wide area communications for the DEA facility are through the plain old telephone system (POTS) direct dial service, using a TAC end system interface, or through narrowband ISDN basic rate interface (BRI) public switched telephone network (PSTN) service. All end systems are connected either to the headquarters ISDN local cable plant or to the Ethernet backbone. Voice communications are accommodated only through ISDN-connected end systems. Information flow between LSE-A and LSE-C is by X.400 messaging, only. There is no interactive data or voice communications between LSE-A and LSE-C. Information flow between LSE-D and LSE-C is via voice, interactive data, and X.400 messaging. Information flow among the Drug Interdiction community LSEs is through voice and interactive database query/response applications. Figure 9-10 illustrates the DEA site transfer system architecture.

The protocol stacks in DEA end systems are the same as the Tactical and Intelligence site end systems for data communications (see figure 9-6). For voice communications, a condensed stack for ISDN is used (i.e., layers 7, 2, and 1, with layer 1 only used once a call has been established).

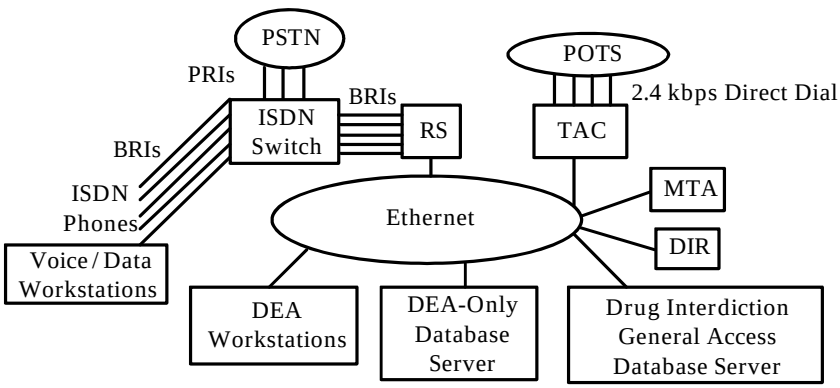

Figure 9-10. DEA Site Transfer System Architecture (LSE-C)

9.4.4 DOJ Site LSE Architecture (LSE-D)

The DOJ site (LSE-D) is located at DOJ headquarters in Washington, D.C. The end systems, local communications and wide area communications systems are identical to those at DEA Headquarters (see section 9.4.3) with one exception, unlike LSE-C, the DOJ site has only one database server for DOJ-Only use.

9.4.5 Drug Interdiction User Site LSE Architecture (LSE-E)

The drug interdiction user sites (LSE-E) are composed of various configurations of end systems and transfer system. End systems typically are workstations or personal communication and computer systems. The workstations interact with the drug interdiction database at DEA headquarters using direct-dial (POTS), or narrowband ISDN (BRI). These workstations may be configured to work with other drug interdiction community workstations using an ATM interface adapter or various LAN interfaces. The personal communications and computer systems are hand-held devices or notebook computer systems with a radio communications capability (e.g., cellular, digital cellular, VHF), and optionally may have LAN interfaces. The drug interdiction functional architecture is depicted in figure 9-11.

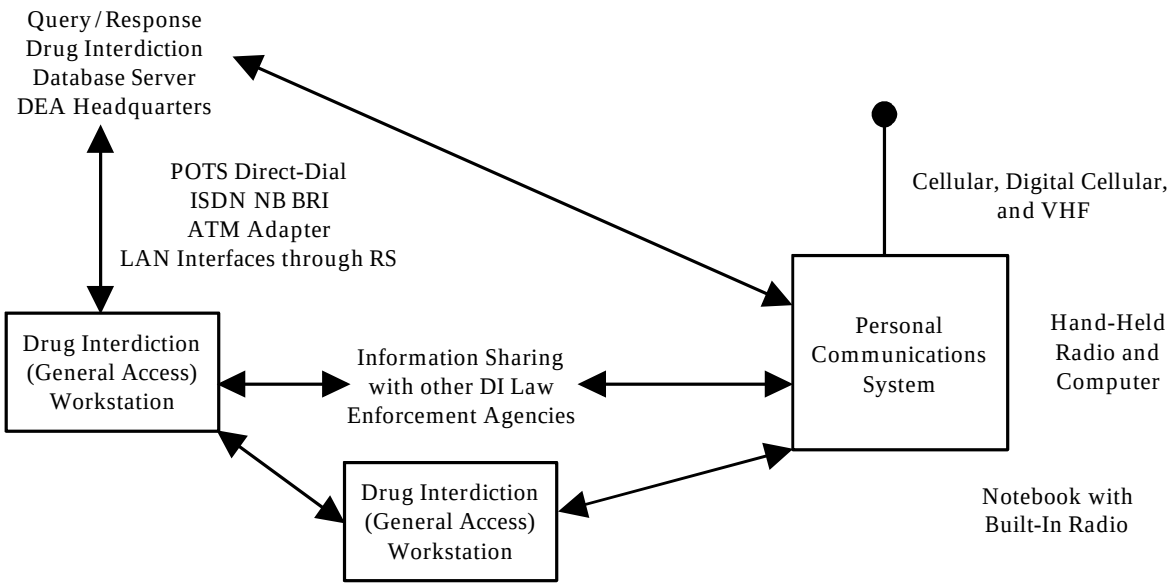

Figure 9-11. Drug Interdiction Site Functional Architecture (LSE-E)

9.5 SYSTEM SECURITY POLICIES

This section uses the information presented in sections 9.1 through 9.4 to develop system security policies for each of the LSEs in this example.

9.5.1 Intelligence Site System Security Policy (LSE-A)

The intelligence site (LSE-A) system security policy is a composite of the information domain security policies supported by the site information systems, namely, the Intelligence, Intelligence Security Management, Tactical, Tactical Security Management, DEA-Only, DEA/DI Security Management, and public information domains. The site system security policy includes the security services (and their strengths) required to satisfy the composite of the information domain security policies supported in LSE-A . The intelligence site processes, stores, and communicates several national classification levels ranging from Unclassified to Top Secret. Table 9-14 presents the intelligence site (LSE-A) system security policy.

9.5.2 Tactical Site System Security Policy (LSE-B)

The tactical site (LSE-B) system security policy is a composite of the information domain security policies supported by the site information systems, namely, the Tactical, Tactical Security Management (optionally Intelligence, Intelligence Security

Management), and public information domains. The tactical site processes, stores, and communicates national classification levels ranging from Unclassified to Top Secret (Top Secret only if the optional intelligence information domain is employed). LSE-B must support the same security services as the Intelligence site LSE (see table 9-14) except that if the Intelligence information domain is not employed, the strengths of service for authentication and access control are reduced from high to medium.

9.5.3 DEA Site System Security Policy (LSE-C)

The DEA site (LSE-C) system security policy is a composite of the information domain security policies supported by the site information systems, namely, the DEA-Only, Drug Interdiction, DEA/DI Security Management, and Public information domains. The DEA site processes, stores, and communicates information of national classification level N/DI as well as information which is labeled DI. Table 9-15 presents the DEA site system security policy.

9.5.4 DOJ Site System Security Policy (LSE-D)

The DOJ site (LSE-D) system security policy is a composite of the information domain security policies supported by the site information systems, namely, the DOJ-Only, Drug Interdiction, DI/DEA Security Management, and public information domains. The DOJ site processes, stores, and communicates information that is labeled Drug Interdiction. Table 9-16 presents the DOJ site system security policy.

9.5.5 Drug Interdiction User Site System Security Policy (LSE-E)

The drug interdiction sites (LSE-E) include U.S.-only and foreign systems. Although these systems may support information domains related to their specific mission functions, the Drug Interdiction information domain, its associated DEA/DI security management information domain, and the Public information domain are the only information domains considered in this example. LSE-E processes, stores, and communicates information that is labeled DI. Table 9-17 presents the Drug Interdiction user site system security policy.

Table 9-14. Intelligence Site System Security Policy (LSE-A)

Security Service	Strength of Service
Authentication	Range: high to medium Environment: high Local Security Context: medium Remote Security Context: high Security Association: high Staged Delivery Security Context: high
Access Control	Range: high to medium Environment: high Strict Isolation of Domains: high Local End System Interdomain Transfers: high Security Context Privilege/Role Management: medium Security Association: high Staged Delivery Security Context: high Interdomain Transfers Intelligence to Tactical: high Tactical to Intelligence: medium Intelligence to DEA-Only: high Access to Associated Security Management Domains: high
Data Integrity	Range: high to medium Environment (i.e., configuration management): high Implicit from Strict Isolation: medium Security Association: high Staged Delivery Security Context: high
Non-Repudiation	Range: high to medium Environment (i.e., physical signature): medium Staged Delivery Security Context: high
Availability	Range: high to medium Environment (e.g., backup power, redundant components, multiple wide area communication access paths): high - 99 percent Network Management: medium (e.g., fault isolation, alternate path redirection)
System Integrity	Security Critical Function Trusted Software and Hardware: high Security Management: high (requires separate information and associated security management domains; relationship is 1:1)

Table 9-15. DEA Site System Security Policy (LSE-C)

Security Service	Strength of Service
Authentication	Range: high to low Environment: low to medium Local Security Context: medium Remote Security Context: medium Security Association: high Staged Delivery Security Context: high
Access Control	Range: high to medium Environment: medium (access to headquarters building) Strict Isolation of Domains: high Local End System Interdomain Transfers: medium Security Context Privilege/Role Management: medium Security Association: high Staged Delivery Security Context: high Interdomain transfers - DEA-Only to Drug Interdiction: medium Access to Associated Security Management Domain: medium
Data Integrity	Range: high to medium Environment: high Implicit from Strict Isolation: medium Security Association Enforcement: high Staged Delivery Security Context: high
Non-Repudiation	Range: high to medium Environment: medium Staged Delivery Security Context: high
Availability	Range: medium to low Environment: low Network Management: medium
System Integrity	Security Critical Function Trusted Software and Hardware: high Security Management: medium (requires separate information and associated security management domains; relationship is 1:many)

9.6 SYSTEM SECURITY ARCHITECTURES

9.6.1 Intelligence Site System Security Architecture (LSE-A)

The intelligence site system security architecture (LSE-A) reflects the security mechanisms used to provide the security services necessary to support the system security policy and each of the information domain security policies. The four parts of the architecture are addressed separately (i.e., end system, transfer system, security

management, environment), followed by an aggregate mapping back to the system and information domain security policies.

Table 9-16. DOJ Site System Security Policy (LSE-D)

Security Service	Strength of Service
Authentication	Range: high to low Environment: low Local Security Context: medium Remote Security Context: medium Security Association: high Staged Delivery Security Context: high
Access Control	Range: high to medium Environment: medium (access to headquarters building) Strict Isolation of Domains: high Local End System Interdomain Transfers: medium Security Context Privilege/Role Management: medium Security Association: high Staged Delivery Security Context: high Interdomain transfers - DOJ-Only to Drug Interdiction: medium
Data Integrity	Range: high to medium Environment: medium Implicit from Strict Isolation: medium Security Association Enforcement: high Staged Delivery Security Context: high
Availability	Environment: low Network Management: low
System Integrity	Security Critical Function Trusted Software and Hardware: low Security Management: low

9.6.1.1 Intelligence Site End System Security Architecture (LSE-A)

All IAWSs will be evaluated products which will provide the DGSA strict isolation target assurance measurement (i.e., can be used for domains ranging from Unclassified through Top Secret). These IAWSs may be configured to accommodate any number of different information domains, regardless of classification and category sensitivity. For LSE-A, IAWSs form workcenters that support different mission functions. Therefore, the addition of the tactical and DEA-Only information domains has created the need for tactical flight planning and Intelligence/DEA

workcenters. There are no changes to the other end systems in the intelligence mission LSE.

Table 9-17. Drug Interdiction User Site System Security Policy (LSE-E)

Security Service	Strength of Service
Authentication	Range: high to low Environment: low Local Security Context: medium Remote Security Context: medium Security Association: high
Access Control	Range: high to low Environment: medium to low Strict Isolation of Domains: high Security Context Privilege/Role Management: medium Security Association: high
Data Integrity	Range: high to medium Environment: low Implicit from Strict Isolation: medium Security Association Enforcement: high
Availability	Environment: low Network Management: low
System Integrity	Security Critical Function Trusted Software and Hardware: low Security Management: medium to low

The intelligence site LSE (LSE-A) has end systems which will be used for varying information domain configurations. Specifically, the end systems in this LSE will support the following information domains:

- End system configuration 1:Intelligence and Intelligence Security Management
- End system configuration 2:Intelligence, Intelligence Security Management, and Public
- End system configuration 3:Intelligence, Intelligence Security Management, Tactical, Tactical Security Management, and Public
- End system configuration 4:Intelligence, Intelligence Security Management, DEA-Only, DEA/DI Security Management, and Public

End system configurations 1 and 2 existed in LSE-A prior to the need for joint mission operations with the tactical and drug interdiction communities. These end systems are based on compartmented mode workstation (CMW) security technology. The CMW end systems support the satellite uplink/downlink control operations, collection post-processing and production control, and select intelligence analyst workcenter positions within LSE-A. End system configurations 3 and 4 will be added to LSE-A. These are state-of-the-art multidomain end systems which can accommodate (securely) any range of classified information and any number of information domains.

The ISDN TAC, defined in the LSE-A transfer system architecture (above), is an end system based on the state-of-the-art multidomain workstation. This end system acts as the secure relay between FDDI-connected end systems and ISDN-connected end systems (systems in other LSEs that use narrowband ISDN PSTN services).

All of the system components within LSE-A are located in a high-assurance, physically and procedurally protected, secure facility. This view of the end system information domain allocation for LSE-A is illustrated in figure 9-12.

LSE-A end system security architecture elements are summarized in table 9-18. Although the transfer system, security management, and environmental architectural elements are inherently part of the end system security architecture, for clarity these components are addressed separately (see sections 9.6.1.2 through 9.6.1.4).

9.6.1.2 Intelligence Site Transfer System Security Architecture (LSE-A)

The intelligence site transfer system (LSE-A) is composed of the communications protocol stacks in end systems, relay systems, the local communications system, and the connectivity to the wide area communications systems. All systems (except the CMW end systems) are capable of multimedia communications; most employ both data and voice. Some only employ data but are capable of accommodating video communications.

Intradomain communications occur within the intelligence information domain using file transfers. Communications with the intelligence security management information domain use virtual terminal applications and CMIP exchanges.

Interdomain transfers from the Intelligence information domain to the Tactical information domain occur within IAWSs at the Intel/tactical workcenter. Intradomain communications use file transfers between analyst IAWSs and the workcenter control IAWSs.

Intradomain communications between the Intelligence information domain Intelligence/Tactical workcenter control IAWS and the Tactical information domain flight-planning workstations is provided by staged X.400 messaging (DMS compliant). Interdomain transfers from the Intelligence information domain to the DEA-Only information domain follow a similar. Security mechanisms used for the transfer system in LSE-A is primarily protocol based. The security protocols supported by the intelligence site were identified in table 9-6.

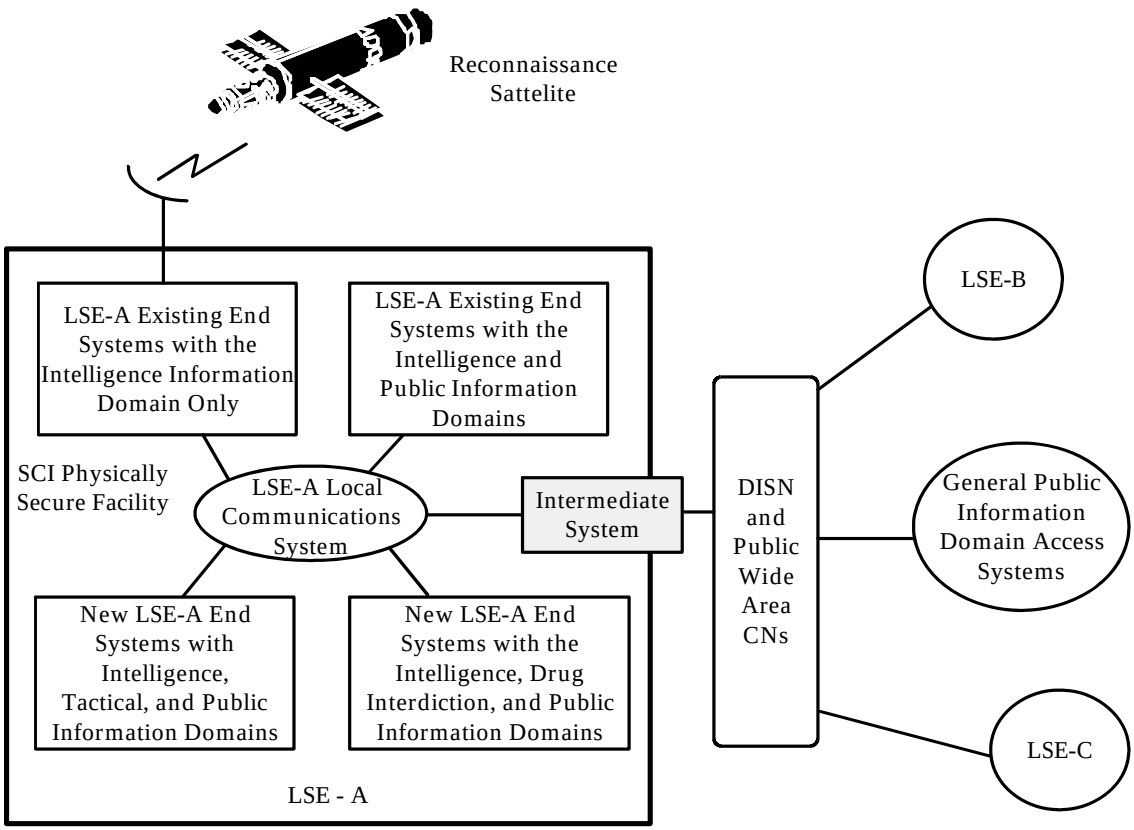

Figure 9-12. LSE-A End System Information Domain Allocations

Table 9-18. LSE-A End System Security Architecture Elements

Service and Strength	Selected Mechanism(s)
Local User Authentication (Medium)	Password per user, changed periodically. Passwords are maintained in the security management domains associated with the intelligence, tactical , and DEA-Only information domains. For end systems which support the public information domain, a user must first be authenticated to the system.
Access Control (High)	Strict isolation (DGSA requirement). All information objects in their respective information domains are labeled uniquely for that particular information domain. All access control decisions (including interdomain transfers) are mediated by the Security Policy Decision Function (SPDF). Only specific workstations and user identities are permitted to release downgraded information objects to other information domains. Strict isolation applied to end system configurations 1 and 2 is enforced by the mandatory access control policy in the CMWs.
Data Integrity (Medium)	Data integrity is attained through access control mechanisms and through well-defined procedural configuration control mechanisms.
Availability (High)	End system availability must be 99 percent. It is achieved by redundant critical system components, a continuous electronic backup of information objects, and a ready supply of spare systems and parts.
Pervasive Services (High) - Hardware Protection, Audit, and Event Handling	Hardware Protection: By tamper resistant packaging and facility security. Audit: Audit data is created and controlled by the SMAP for each security management information domain. Audit tools are available to the administrator for each of the security management information domains. Accumulated audit data for a particular information domain (i.e., Intelligence, Tactical, DEA-Only) is forwarded from every end system's associated security management domain SMAP to the specific mission function (i.e., mission control, Intel/tactical, Intel/DEA-Only) SMAP and SMIB, where that information domain's security administrators reside. Event Handling: Security-critical events are monitored by an end system security management information domain and sent directly to the information domain where the security administrators reside.

9.6.1.3 Intelligence Site Security Management Architecture (LSE-A)

The security management architecture for the intelligence site (LSE-A) is distributed across various positions within the intelligence mission. The Mission Control position is responsible for the overall management (including security) of the intelligence mission function. It is responsible for designating a security officer and the security officer, in turn, designates security administrator(s), as required. LSE-A supports the Intelligence Security Management, Tactical Security Management, and DEA/DI Security Management information domains.

Because of the addition of the tactical and drug interdiction missions to LSE-A, the intelligence site has moved from a centralized security management architecture to one which is more distributed. Mission Control, although still responsible for the overall management of LSE-A, has delegated security management responsibilities for the tactical interface to the supervisor of the intelligence/tactical workcenter, and for the drug interdiction interface to the supervisor of the intelligence/DEA-Only workcenter. These supervisors are considered security administrators and are responsible for maintaining the rules of enforcement and for downloading the SMIBs to all applicable end and relay systems within the intelligence site.

Those end systems which support the Intelligence, Tactical, and Public information domains (end system configuration 3) will have two SMAPs: one dealing with the Intelligence information domain and the other dealing with the Tactical information domain. The public information domain has no SMAP associated with it. There is strict isolation between the public information domain and all other information domains supported by this LSE. Likewise, those end systems which support the Intelligence, DEA-Only, and Public information domains (end system configuration 4) will also have two SMAPs. Finally, those end systems which do not participate in interdomain transfers (end system configurations 1 and 2) will have only one SMAP dealing with the intelligence mission which is administered by Mission Control.

Note that the SME for the tactical information domain could be located within LSE-B; however, because of the unique downgrade nature of exporting downgraded Intelligence information domain objects into the Tactical information domain, it is more appropriate to maintain a Tactical information domain security administrator in both LSE-A and LSE-B. Audit review, for example, should be conducted by both LSE security administrators (i.e., two party review). The Intelligence information domain security officer maintains those portions of the Intelligence information domain security policy that allow information object downgrade and transfer to the Tactical information domain, and for information object (text only) downgrade and transfer to the DEA-Only information domain. The supervisor of the intelligence/tactical workcenter in LSE-A maintains that portion of the Tactical information domain security policy which allows information object imports from the Intelligence information domain.

9.6.1.4 Intelligence Site Environmental Security Architecture (LSE-A)

The stringent environmental protection of LSE-A provides physical, personnel, administrative, and procedural security mechanisms in accordance with intelligence community doctrine.

Access to the air base is protected by armed guards at all entry and exit points. Only authorized personnel may enter the base. Authentication is provided by vehicle

registration and the use of a specific sticker displayed on the windshield of the vehicle. Visitors must register the visitor center outside the main entrance gate to the base. Drivers licenses are used to verify the identity of the registered visitor. All vehicles entering and exiting the air base are subject to search and seizure.

Access to the intelligence mission's operational facilities requires that personnel have an identification badge with picture and other identifying information. On the back of the badge is a magnetic strip that is interpreted by a badge reader and challenge/response device. The badge reader devices require the holder of the badge to enter a Personal Identification Number (PIN). An armed guard monitors the badge reader entry system and stops anyone unable to pass the badge reader from entering the facility. All badged personnel must have a Top Secret security clearance, which also includes a special background investigation. Visitors to the facility must be registered for access by those cleared personnel sponsoring the visit.

Entry into specific areas within the intelligence facility (e.g., Mission Control, various workcenters) requires knowledge of the cipher lock combination on the entrance. Only assigned personnel are provided the combinations to these facilities. Emergency exit doors have electronic alarm devices, triggered when the door is open. The doors cannot be open from the outside.

The entire intelligence facility, with exception of the entrance and exit gate, is surrounded by double 12-foot metal fencing. The perimeter between the two fences is 20 feet and the perimeter area has an electronic alarm system.

9.6.2 Tactical Site System Security Architecture (LSE-B)

LSE-B has end systems which will be used for varying information domain configurations. Some end systems in this LSE will implement only the tactical flight-planning mission. As an option, other end systems in this LSE will implement the intelligence mission function and the tactical flight planning mission, to enhance availability and timeliness. Other end systems, which have no direct relationship to the mission functions identified in the example, but which already exist in this LSE, are those requiring a SECRET clearance, and those that require no security clearance (e.g., Unclassified but Sensitive information domains, and the Public information domain).

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss the tactical site system security architecture.)

9.6.3 DEA Site System Security Architecture (LSE-C)

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss the DEA site system security architecture.)

9.6.4 DOJ Site System Security Architecture (LSE-D)

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss the DOJ site system security architecture.)

9.6.5 Drug Interdiction System Security Architecture (LSE-E)

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss a representative drug interdiction site system security architecture.)

9.7 INTERDEPENDENCY ANALYSIS

Interdependency analysis is the final examination and scrutiny of the developed LSE security architectures. The purpose of the interdependency analysis is to validate the credibility of the combined system security architectures. This last step of the security architecture development process focuses upon the architectures' ability to support interoperability and to ensure that sufficient protection is provided that satisfies the minimum assurance requirements established by the corresponding information domain and system security policies.

This aspect of security analysis examines the strengths and weaknesses of the chosen security mechanisms, security service-to-security mechanism interactions, and security mechanism-to-security mechanism interactions. Discovered vulnerabilities can create unacceptable residual risk which must be dealt with before advocating the security architecture for implementation. As a result, vulnerability analysis, residual risk assessment, and techniques used for risk abatement are all major activities of interdependency analysis. Modifications to the system security architecture will be required if the residual risk is deemed unacceptable.

The results of interdependency analysis can be useful in developing (or updating) system (LSE) accreditation plans and can also help security certifiers understand the security interrelationships and the purported strengths and weaknesses of the security architecture. The results will also help system designers and integrators understand how security mechanisms must fit into the end system and component specifications to realize the overall operational system. In its final form, the security architecture becomes an integral part of the overall system architecture.

The scope of interdependency analysis and the organization of the remainder of this section address the following topics:

- Interoperability imperatives
  - Analysis of security-relevant and security-critical software function interactions
  - Adherence to security-relevant and security-critical data standards and specifications, including communications protocols, related security protocols, and the algorithms and key management techniques to be used by the security protocols
- Analysis of security service-to-security mechanism interactions, as well as security mechanism strengths and weaknesses for intradomain and interdomain protection relationships, as these relationships correspond to LSEs for which security architectural choices have been made
- Analysis of vulnerabilities and determination of residual risk associated with any security mechanism and any security mechanism interaction weaknesses.

9.7.1 Interoperability Analysis

Interoperability, in general, is concerned with standardization in three areas: function and service *definitions*, mechanism, data, interface, and communications protocol *specifications*, and conforming *implementations* of the specifications.

Interoperability is critical because the system simply will not work if the security elements that need to interact can not.[9-4] Attempting to retrofit interoperability features into otherwise non-interoperable security mechanisms is a potentially dangerous practice. If retrofits for interoperability are the only viable solution, then they should be prototyped and tested, with results rigorously evaluated before considering any alterations to the security architecture.

[9-4] It is envisioned that the development of security profiles will assist in understanding interoperability constraints among products and mechanisms. One method this information could be conveyed is in a security mechanisms catalog (as discussed in section 6.3.2) from which a security architect would chose mechanisms. The use of such profiles could significantly reduce the time and cost associated with security architecture development and interdependency analysis

For the purpose of interdependency analysis, the interoperability examination relates to intradomain and interdomain interactions. This equates to the following elements of the security architecture to which standardization applies:

- security critical/relevant software functions which must interoperate properly and cooperate appropriately; and,

- security critical/relevant data structures, such that they can be interpreted and acted upon properly, in a uniform way.

Further, intradomain interoperability relates to communication and associated security protocol standard choices, including the choice of cryptographic algorithms and key management technique(s). This aspect of interoperability is not required for interdomain interactions, because interdomain transfers and interprocess communications can only occur within the bounds of a single end system; interdomain security associations are forbidden by the DGSA.

9.7.1.1 Security-Critical and Security-Relevant Software Functions and Data Structures

Interoperability aspects of security critical and relevant software functions relate to three specific areas: the specific services the software functions must perform, the interfaces utilized to intercommunicate between these functions, and the information that is to be conveyed (data structures) for interprocess (security function) communication. These interoperability aspects relate to both intradomain and interdomain information object transfers and interprocess communications as they are handled through cooperation of both security-relevant (e.g., an information domain SMAP and SMIB) and security-critical (e.g., reference monitor mediated) functions (and their security-critical data structures).

For the security architecture defined for each LSE in this example, the security management application process is a GOTS portable set of software modules that have been chosen for each end system implementation. The SMIB data structures are all standards which the SMAP software interprets and handles in a common, uniform way. Although the reference monitor utilized in the end systems is different from system to system (from both an implementation and an assurance perspective), its interface and its IPC mechanism interfaces are standard (i.e., the SMAP software interface to the reference monitor in all cases is common and uniform.

9.7.1.2 Intradomain Communications Interoperability

Table 9-19 summarizes the intradomain interoperability areas related to communications using the transfer system. Secure voice communication is only applicable to information domains 1 and 3. All mission areas have nonsecure voice and data communications capabilities using the Public information domain (domain 9).

Information domain 1 in the Intelligence site LSE utilizes file transfers for intra-LSE communications. This LSE obtains collection requirements for its mission function from a remote LSE via X.400 secure messaging, following DMS standards. This LSE optionally communicates with the Tactical site LSE (which also supports information domain 1) via X.400 secure messaging, and also has the option for secure voice communications. This option available because the Tactical site LSE only supports information domain 1 as a part of a contingency plan. The contingency plan may never become operational.

Information domain 2 in the Intelligence site LSE (and optionally the Tactical site LSE if the contingency plan is activated), utilizes the CMIP and security associations using SAMP and SDE. The security administrators for information domain 1 are the only members of Domain 2, and reside in one or more of the Intelligence LSE mission control positions. The information domain 2 SMIB is administered from these position only.

Information domain 3 in both the Intelligence site LSE and the Tactical site LSE utilizes file transfers for intra-LSE communication, and security associations using SAMP and SDE. For inter-LSE communications (but still within information domain 3), there is two-way X.400 secure messaging, following DMS standards. There is also a secure voice communications capability between the two LSEs in this information domain.

Table 9-19. Intradomain Communications Interoperability Summary

Information Domain	Communications Applications	Security Protocols	Cryptographic Algorithms	Key Management Techniques
1	• Intra-LSE Interactive Data (FTAM, P3/P7) • Inter-LSE Staged Delivery Messages (DMS Compliant) • Inter-LSE Voice with Encryption Conversions at an Appropriate Intra-LSE Relay System	SAMP-SDE MSP Organizational Message UA-UA Protected SAMP-SDE (over ST2) Converted at ATM Relay to ATM Link Encryption or at ISDN TAC to ISDN Narrowband Link Encryption Using STE	Military Grade (MG) Type 1 MG Type 1 with Other MG Algorithms Possible MG Type 1 and Type 2	SDNS MG Type 1 SDNS MG - Type 1 SDNS MG - Type 1
2	Interactive Data SMAP-SMAP using CMIP	Intra-LSE Uses SAMP-SDE Inter LSE Uses SAMP-NLSP	MG Type 1	SDNS MG - Type 1
3	Same as 1	Same as 1	Same as 1	Same as 1
4	Same as 2	Same as 2	Same as 2	Same as 2
5	• Intra-LSE Interactive Data (FTAM, P3/P7) • Inter-LSE Staged Delivery Messages (DMS Compliant)	SAMP-SDE MSP Organizational Message UA-UA protected	Commercial-Government Grade Type2 Commercial-Government Grade Type2	Commercial-Government Type Type2 Commercial-Government Type Type2
6	• Intra-LSE FTAM • Inter-LSE Query/Response	SAMP-SDE SAMP-NLSP	Commercial Grade Type2 Commercial Grade Type2	Commercial Type Type2 Commercial Type Type2
7	• Intra-LSE FTAM • Inter-LSE Query/Response	SAMP-SDE SAMP-NLSP	Commercial Grade Type2 Commercial Grade Type2	Commercial Type Type2 Commercial Type Type2
8	Same as 2 & 4	Same as 2 & 4	Commercial Grade Type2	Commercial Type Type2
9	N/A	N/A	N/A	N/A

Information domain 4 in both the Intelligence LSE and the Tactical LSE utilizes CMIP and security associations using SAMP and SDE. There is no inter-LSE communication for information domain 4. A security administrator in the Intelligence site LSE maintains the SMIB within this LSE. Another administrator in

the Tactical LSE maintains the SMIB for information domain 4 in that LSE. Common security management attributes required for interoperability for information domain 4 intradomain communications between LSEs is coordinated between the two security administrators via X.400 secure messages, following DMS standards.

Information domain 5 in the Intelligence site LSE communicates with information domain 5 in the DEA site LSE using X.400 secure messaging, following DMS standards. This communication is unidirectional from the Intelligence site LSE to the DEA site LSE. The cryptographic algorithm used by the Message Security Protocol (MSP) is different from that used for information domains 1 and 3. The services provided by MSP and the key management approach it uses are the same for all of the information domains that engage in secure X.400 messaging.

Intradomain, intra-LSE, and inter-LSE communication for information domains 6 and 7 utilize secure voice communications, file transfers, and interactive query/response. Data transfers use SAMP for security associations and the appropriate lower layer security protocol. The grade of the cryptographic algorithm is not the same as that used in information domains 1, 2, 3, and 4. There is no cryptographic interoperability between information domains 6 and 7 and information domains 1, 2, 3, and 4, but none is required.

Intradomain communications capabilities are interoperable from the perspective of security protocols, cryptographic algorithms, and key management techniques. Common security management assures security association correctness and binding to the proper security context for all intradomain communications over local and wide area communications networks. No interdomain network communications are allowed.

Conclusively, there are no interoperability problems associated with the various system security architectures. With the exception of the DOJ-Only information domain, all local and remote security administrator interfaces are handled in the same way, but each must have the appropriate security context and security association (for remote access). For the DOJ-Only information domain, security management is integral to the information domain, not separate and associated like the other information domains. This arrangement does not constrain interoperability, but it potentially complicates interdomain transfers from the DOJ-Only information domain to the Drug Interdiction information domain. The following section discusses this interdependence.

9.7.2 Security Service and Mechanism Interdependencies

This aspect of interdependency analysis checks to ensure that the strength of protection provided within each LSE security architecture achieves information domain and end system policy requirements. It then examines how the mechanisms provide protection, and the relationships between mechanisms from LSE to LSE, end system to end system, and by intradomain and interdomain interactions. The mechanisms are viewed from three architectural perspectives (environment, end system, transfer system), and in aggregate.

Analysis related to minimum assurance requirements is important because it ensures that the concept of absolute protection will be achieved for each information domain across end systems, transfer systems, and LSEs. This is particularly critical for LSE security architectures that have rejected certain mechanisms chosen in other LSE security architectures. In these cases, an appropriate argument must be made for the variance in mechanisms and proof offered that the concept of absolute protection for each information domain will indeed be realized.

In the most general sense, the starting point for this part of the analysis is to identify the interservice relationships as they pertain to the LSE security requirements, as summarized in table 9-20.

Table 9-20. Interservice Relation to System Requirements

Requirement	Primary Security Service	Supporting Security Services
All system information must be protected from disclosure (a)	Confidentiality	I&A, Access Control
All system information must be protected from modification (b)	Integrity	I&A, Confidentiality
All mission resources must be protected against hardware and software failures (c)	Integrity	I&A, Access Control, Availability,
All mission resources must be physically and electronically protected (d)	Access Control	I&A, Integrity, Availability
All mission resources must be protected against tampering (e)	Access Control	I&A, Integrity, Availability,

Data confidentiality is the primary security service needed to meet requirement (a); all system information must be protected from disclosure. However, I&A is a supporting service because it must identify the user for a particular security context

and the end system addresses before the proper privileges and security attributes can be determined for the security association. Access control also supports confidentiality (e.g., without encryption) by providing security context strict isolation, for enforcing user privileges within the security context, and for enforcing the parameters established for security associations.

Data integrity is the primary security service needed to meet requirement (b). Both confidentially and access control support the integrity service by not allowing a user to access the data without the proper privileges. I&A is necessary to authenticate the user before the proper privileges are determined.

Component access control is the primary security service needed to meet requirement (c); all system information must be protected from modification. I&A is necessary to support both physical and electronic access control to mission components. Integrity of the mission components supports the access control service. Availability also relies on the access control to the mission components.

System integrity is the primary security service needed to meet requirement (d); all mission resources must be protected against hardware and software failures. In this case, the system needs both physical and electronic protection against tampering. Access control supports system integrity by not allowing individuals to access the parts of the system that will affect the integrity of the system. I&A assists access control in authenticating the individual who needs access. Availability also relies on system integrity to system availability.

Access control is the primary security service needed to meet requirement (d); all mission resources must be protected against tampering. I&A is necessary to support all access control services. These services protect system integrity and availability.

The next step in interdependency analysis is to weigh the strengths of the mechanisms utilized to meet the requirements. In order to assess the balance of security strengths, which is needed to perform the interdependency analysis, a scale of strengths must be developed. For this example, a hypothetical scale is used to illustrate the analysis. This hypothetical scale should in no way be construed as a recommended scale. It is envisioned that a real DoD enterprise-level scale will evolve as the DGSA begins to mature and becomes accepted, and as transition plans are developed to lead toward the goal. Table 9-21 illustrates the example hypothetical mechanism strength scale. The scale identifies six levels of strength (from little or no security strength to very high security strength). To perform the assessment, a rating is made for each mechanism or set of mechanisms that satisfies a particular strength of security service called for in the security policies.

Table 9-21. Hypothetical Mechanism Strength Scale

SCALE	Element Weights		Total Score	
	Minimum	Maximum	Minimum	Maximum
No Strength	0	.5	0	< 20
Very Low Strength	.5	2	>20	< 59
Low Strength	3	4	>59	< 99
Medium Strength	5	6	>99	< 139
High Strength	7	8	>139	< 179
Very High Strength	9	10	>179	< 191

The security mechanisms strengths required for each security service to be rated are highlighted in table 9-22. A mechanism strength is assigned to each security service for the environment, end systems, and transfer system. In end systems, non-repudiation is not rated, but the strength of the mechanism for strict isolation is assigned. Each category is rated on a scale of 1 to 10. Because of its critical importance, strict isolation is weighted by a factor of 2 (i.e., the 1-to-10 rating is multiplied by two). With this hypothetical scale, a perfect security rating for all elements would total 190 points .

The composite ratings for each information domain in an LSE are shown in table 9-23, by category (environment, end system, and transfer system) and total weight, indicating the highest measure indicator of each architectural element.

Specific end systems within LSEs may have different ratings, and these must be examined from an absolute protection perspective. In this example architecture, this situation arises in the Intelligence site LSE. In this LSE there are some end systems that implement only the Intelligence information domain and its associated security management information domain, and these end systems are based on CMW technology. All the other end systems in the Intelligence site LSE are rated in accordance with the overall LSE rating (i.e. high assurance) and reflect the aggregate LSE rating.

Table 9-22. Weighted Mechanisms by Domain and LSE

	Strength of Environmental Mechanisms						Strength of End System Mechanisms						Strength of Transfer System Mechanisms					
Domain -LSE	Auth	AC	Data Integrity	Conf	NR	Avail	Auth	AC	Strict Isolation	Data Integrity	Conf	Avail	Auth	AC	Data Integrity	Conf	NR	Avail
1-A	8	8	6	6	7	7	5	7	18	7	7	8	9	9	9	9	8	7
1-B	7	8	5	6	7	7	5	7	16	5	5	7	9	9	9	9	8	6
2-A	8	8	6	6	7	7	5	7	18	7	7	8	9	9	9	9	8	7
2-B	7	8	5	6	7	7	5	7	16	5	5	7	9	9	9	9	8	6
3-A	8	8	6	6	7	7	5	7	18	7	7	8	9	9	9	9	8	7
3-B	7	8	5	6	7	7	5	7	16	5	5	7	9	9	9	9	8	6
4-A	8	8	6	6	7	7	5	7	18	7	7	8	9	9	9	9	8	7
4-B	7	8	5	6	7	7	5	7	16	5	5	7	9	9	9	9	N/A	6
5-A	8	8	6	6	7	7	5	7	18	7	7	8	6.5	6.5	6.5	6.5	6.5	5
5-C	5	5	5	4	7	5	5	6	10	5	5	5	6.5	6.5	6.5	6.5	6.5	5
6-C	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
6-D	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
6-E	4	4	4	3	7	4	5	5	10	5	5	5	5	5	5	5	N/A	5
7-C	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
7-D	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
8-A	8	8	6	6	7	7	5	7	18	7	7	8	N/A	N/A	N/A	N/A	N/A	N/A
8-C	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
8-D	5	5	5	4	7	5	5	6	10	5	5	5	5	5	5	5	N/A	5
8-E	4	4	4	3	7	4	5	5	10	5	5	5	5	5	5	5	N/A	5
9-A	8	8	6	6	N/A	7	5	N/A	18	N/A	N/A	8	N/A	N/A	N/A	N/A	N/A	5
9-B	7	8	5	6	N/A	7	5	N/A	16	N/A	N/A	5	N/A	N/A	N/A	N/A	N/A	5
9-C	5	5	5	4	N/A	5	5	N/A	10	N/A	N/A	5	N/A	N/A	N/A	N/A	N/A	5
9-D	5	5	5	4	N/A	5	5	N/A	10	N/A	N/A	5	N/A	N/A	N/A	N/A	N/A	5
9-E	4	4	4	3	N/A	4	5	N/A	10	N/A	N/A	5	N/A	N/A	N/A	N/A	N/A	5

Table 9-23. Weighted Mechanism Totals, by Category and LSE-Information Domain

LSE-InformationDomains	Environment Total	End System Total	Transfer System Total	Aggregate LSE Total
A-1, 2, 3, 4, 5, 8, and 9	42 - Rating is High (low end of)	52 - Rating is High (low part of)	51 - Rating is High (low end of)	145 - Rating is High (low end of)
B-(opt 1,2) 3, 4, and 9	40 - Rating is Medium (high end of)	45 - Rating is High (low part of)	50 - Rating is High (low end of)	135 - Rating is Medium (very high end of)
C-5, 6, 7, 8, and 9	31 - Rating is Medium (low end of)	36 - Rating is Medium (low end of)	37.5 - Rating is Medium (high end of)	104.5 - Rating is Medium (low end of)
D-6, 7, 8, and 9	same as LSE-C	Same as LSE-C	Same as LSE-C (second rating)	Same as LSE-C (second rating)
E-6, 8, and 9	26 - Rating is Low (high end of)	35 - Rating is Medium (low end of)	25 (less NR) - Rating is Medium (low end of)	86 - Rating is Low (High end of)

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss each mechanism in the environment, end systems and transfer system and then will discuss the interrelationship of mechanisms among LSEs and end systems.)

9.7.3 Vulnerabilities, Residual Risk, And Recommended Countermeasures

Vulnerability analysis, residual risk assessment, and countering unacceptable residual risk must be performed before procurement activities commence to discover any security architecture flaws. Although this kind of analysis carried out before functional testing has been performed is not perfect, it can make a significant difference in the success or failure of the development process and its cost.

(This subsection is incomplete as of this draft release. It will be completed in the next draft release and will illustrate and discuss residual risk and countermeasures to unacceptable risk.)

(This page is intentionally blank.)

APPENDIX A

STRAWMAN EXTENDED GSS API

(This appendix will be supplied in the next draft.)

APPENDIX B

DETAILED EXAMPLE OF SECURITY ASSOCIATION ESTABLISHMENT

(This appendix will be supplied in the next draft.)

REFERENCES

American National Standards Institute (ANSI), 1985, *Financial Institution Key Management (Wholesale)*, X9.17

Abrams, Marshall D. and Michael V. Joyce, January 1993, *On TCB Subsets and Trusted Object Management*, MITRE Technical Report 92W0000248, McLean, VA

Center for Information Management (CIM), 1992, *Technical Architecture Framework for Information Management*, Defense Information Systems Agency, Washington, DC

Case, Jeff, Mark Fedor, Martin Schoffstall and James Davin, 1989, *Simple Network Management Protocol (SNMP)*, Internet Request for Comments 1098

Case, Jeff, 1991, *SNMP Version 2*, Internet Request for Comments 1441

International Telegraph and Telephone Consultative Committee (CCITT), 1988, *Recommendations X.400-X.420: Data Communications Networks, Message Handling Systems*

_____, 1992, *Recommendations X.500-X.521 Data Communications Networks, Directory*

Department of Defense, 1985, *Department of Defense Trusted Computer System Evaluation Criteria*, DoD 5200.28-STD, Washington, DC

Institute of Electrical and Electronic Engineers (IEEE), *Standard for Interoperable LAN/MAN Security, Part 2--Secure Data Exchange Protocol Specification*, IEEE 802.10b

_____, 1993, *Standard for Interoperable LAN/MAN Security, Part 3-Key Management Protocol Specification (Draft)*, IEEE 802.10c

International Organization for Standardization (ISO), 1984, Information Processing Systems, *Open Systems Interconnection Reference Model: Basic Reference Model*, ISO 7498

_____, 1987, Information Processing Systems, *Common Management Information Protocol*, ISO 9596

_____, 1988a, Information Processing Systems, *Open Systems Interconnection - Connectionless Network Protocol*, ISO 8473

_____, 1988b, Information Processing Systems, *FTAM - File Transfer, Access, and Management*, ISO 8571-1

_____, 1988c, Information Processing Systems, *ACSE - Association Control Service Element*, IS 8649/8650

_____, 1988d, Information Processing Systems, *Abstract Syntax Notation Number 1 (ASN.1)*, ISO 8824

_____, 1989a, Information Processing Systems, *Open Systems Interconnection Reference Model, Part 2: Security Architecture*, ISO 7498-2

_____, 1989b, Information Processing Systems, *Open Systems Interconnection Reference Model, Part 4: Management Framework*, ISO 7498-4

_____, 29 July 1992a, Information Technology, Network Layer Security Protocol, ISO DIS 11577

_____, 7 December 1992b, Information Technology, Transport Layer Security Protocol, ISO 10736

_____, 1992c, Information Technology - *Open Systems Interconnection - Security Frameworks in Open Systems - Part 3: Access Control*, ISO CD 10181-3

_____, 1993a, Information Technology - *Open Systems Interconnection - Security Frameworks in Open Systems - Part 1: Overview*, ISO CD 10181-1

_____, 1993b, Information Technology - *General Upper Layers Security (GULS) - Part 3: Security Exchange Service Element Protocol (SESEP) Specification*, ISO DIS 11586-3

Linn, John, 1992, *General Security Services - Application Program Interface*, Draft Internet Request for Comments

National Institute of Standards and Technology (NIST), 1992a, *U.S. Government OSI Profile*, United States Federal Information Processing Publication 146-2

_____, 1992b, *Government Network Management Profile*
National Security Agency (NSA), 1992, *Message Security Protocol*, SDN.701 Revision 2.0

_____,22 February 1993, *Department of Defense Information Systems Security Policy*, DISSP-SP.1

Powell, General Colin L., 1992, as quoted in *C4I for the Warrior*, Joint Staff (J6)

Rivest, R. L., A. Shamir and L. Adelman, 1978, *A Method for Obtaining Digital Signatures and Public Key Cryptosystems*, Communications of the ACM, Volume 21, Number 2, pp. 120-126

Rushby, J., September, 1984, "A Trusted Computing Base for Embedded Systems", Proceedings of the 7th DOD/NBS Computer Security Symposium, pp. 294-311

(This page is intentionally blank.)

LIST OF ACRONYMS

ACCS	Army Command and Control System
ACSE	Association Control Service Element
ADF	Access Control Decision Function
AEF	Access Control Enforcement Function
API	Application Program Interface
ASSR	Agreed Set of Security Rules
ATM	Asynchronous Transfer Mode
BISDN	Broadband ISDN
BRI	Basic Rate Interface
C2	Command and Control
C4I	Command, Control, Communications, Computers, and Intelligence
C4IFTW	C4I for the Warrior
CCITT	International Telegraph and Telephone Consultative Committee
CIM	Center for Information Management
CISS	Center for Information System Security
CLNP	Connectionless Network Layer Protocol
CMIP	Common Management Information Protocol
CMW	Compartmented Mode Workstation
CN	Communications Network
COMSEC	Communications Security
COTS	Commercial off-the-shelf
DCI	Director of Central Intelligence
DEA	Drug Enforcement Agency
DGSA	DoD Goal Security Agency
DI	Drug Interdiction
DIS	Defense Information System
DISA	Defense Information Systems Agency
DISN	Defense Information System Network
DISSP	Defense Information Systems Security Program
DMS	Defense Message System
DoD	Department of Defense
DoDIIS	DoD Intelligence Information System
DOJ	Department of Justice
DOS	Department of State

EEI	External Environment Interfaces
EKMS	Electronic Key Management System
ES	End System
FBI	Federal Bureau of Investigation
FDDI	Fiber Distributed Data Interface
FTAM	File Transfer, Access, and Management
GNMP	Government Network Management Profile
GOSIP	Government Open Systems Interconnection Profile
GOTS	Government off-the-shelf
GSS	General Security Service
GULS	General Upper Layer Security
HDTV	High Definition Television
HIPPI	High Performance Parallel Interface
I&A	Identification and Authentication
IAW	Intelligence Analyst Workstation
IM	Information Management
IPC	Interprocess Communication
ISO	International Organization for Standardization
ITSDN	Integrated Tactical/Strategic Data Network
JCALS	Joint Computer Automated Logistics System
JSAN	Joint Staff Architecture of the Nineties
LAN	Local Area Network
LCS	Local Communications System
LMD	Local Management Device
LSE	Local Subscriber Environment
MAP	Management Application Process
MIB	Management Information Base
MISSI	Multilevel Information System Security Initiative
MG	Military Grade
MLS	Multilevel Security
MSP	Message Security Protocol
MTA	Message Transfer Agent
NLSP	Network Layer Security Protocol

NSA	National Security Agency
OMB	Office of Management and Budget
OSE	Open System Environment
OSD	Office of the Secretary of Defense
OSI	Open Systems Interconnection
PABX	
PCS	Personal Communications Services
PIN	Personal Identification Number
POTS	Plain Old Telephone System
PSTN	Public Switched Telephone Network
RCAS	Reserve Component Automated System
RM	Reference Model
RS	Relay System
RSA	Rivest-Shamir-Adelman
RVM	Reference Validation Mechanism
SAID	Security Association Identifier
SAMP	Security Association Management Protocol
SATCOM	Satellite Communications
SBIS	
SCI	Special Compartmented Information
SDE	Secure Data Exchange
SESEP	Security Exchange Service Element Protocol
SILS	Secure Interoperable LAN/MAN Standard
SMAP	Security Management Application Process
SMIB	Security Management Information Base
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
SPDF	Security Policy Decision Function
SPEF	Security Policy Enforcement Function
TAC	Terminal Access Control
TAFIM	Technical Architecture Framework for Information Management
TLSP	Transport Layer Security Protocol
TFS	Traffic Flow Security
TRM	Technical Reference Model
UA	User Agent

USAF	US Air Force
USN	US Navy
USFIB	US Foreign Intelligence Board
USG	US Government
VLSI	Very Large Scale Integration