



An Assessment of the DOD Goal Security Architecture (DGSA) for Non-Military Use

Arthur E. Oldehoeft

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Gaithersburg, MD 20899

QC
100
.U56
NO.5570
1994

NIST

An Assessment of the DOD Goal Security Architecture (DGSA) for Non-Military Use

Arthur E. Oldehoeft

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems Laboratory
Gaithersburg, MD 20899

November 1994



U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Arati Prabhakar, Director

Contents

1. Introduction	1
1.1 Background	1
1.2 A Framework for Evaluation of the DGSA for Non-Military Use	2
2. Overview of the DGSA	2
2.1 Policy and Fundamental Requirements	3
2.2 High-Level Abstraction of Hardware/Software Support	4
2.3 Security Services	5
2.4 Basic Security Concepts	6
2.5 Components of an Implementation Architecture	8
3. Comparison with the International Standards Organization Security Architecture	9
3.1 Security Management Concepts Comparisons	9
3.1.1 Information Domains Concept	9
3.1.2 Security Management Information Bases (SMIBs)	10
3.2 Communication of Security Management Information	10
3.2.1 Distributed Security Management Administration	10
3.2.2 Security Management Applications Protocols	11
3.2.3 End Security Management Functions	11
3.2.4 Security Services Management	11
3.2.5 Security Mechanism Management	11
4. Comparison with the Internet Security Architecture	12
4.1 Internet Security Architecture	12
4.2 ISA Security Philosophy	13
4.3 Comparison Discussion	15
5. Comparison with Other Selected Security Architectures	17
5.1 General Comparative Remarks	17
5.2 Further Comparison with DCE Security Architecture	18
5.2.1 DCE Overview	18
5.2.2 DCE Security Services	18
5.2.3 Additional Comparative Remarks	19
6. Summary Observations	19
A. Other DOD Networking Plans	22
B. Relationship to Other Federal Networking Needs	24
C. MITRE Corporation Study of Transition Issues	25
C.1 Technology	25
C.2 Management	28
C.2.1 High-Level DOD Manager	28
C.2.2 DOD Program Manager	28
C.2.3 Site/System Security Officer	28
C.3 Policy	28
C.3.1 Multiple Security Policies	28
C.3.2 Certification and Accreditation	29

1. Introduction

1.1 Background

The purpose of this study is to assess the potential of the DOD Goal Security Architecture (DGSA) as a model and framework for the development of non-military computer and information security architectures.

The interest in the DGSA is driven by several factors. First, it is a comprehensive security model, concerned with all aspects of computer and information security. Encompassing information and entities at all levels of sensitivity (including unclassified), and relying on architectural system components of all capabilities (from highly secured to unsecured), the DGSA has the potential of meeting the needs of non-military sectors (government agencies, industry, business, education, and individuals). Second, compatibility is important to government agencies, contractors, researchers, and others who have a need to communicate and exchange information in a secure fashion with the DOD. Third, in a recent study, the federal government has called for

- a. fostering the industry-government partnership for improving security, integrity and assurance of services in public telecommunications (noting about 90 percent of DOD's telecommunications are carried by public carriers); and
- b. developing a comprehensive Internet security plan for interconnecting the federal IT community with appropriate state, local, commercial, public and private sector, and foreign government activities with
 - layered protocol standards and techniques employed with a range or set of security service standards with appropriate gateway protection devices, and
 - an architecture identifying the grades of service offered, how each is implemented and assured, how interconnections between networks should be made, and what can be done for users ¹.

To underscore the importance placed by the federal government on security for unclassified (but sensitive) information, the Office of Technology Assessment (in response to a request from the Senate Committee on Governmental Affairs and the House Subcommittee on Telecommunications and Finance) recently published a report² that focuses on policy issues in three areas: 1) national cryptography policy, including federal information processing standards and export controls; 2) guidance on safeguarding unclassified information in federal agencies; and 3) legal issues and information security, including electronic commerce, privacy, and intellectual property. The required flexibility for security of end systems in the National Information Infrastructure is emphasized in a recent assessment³ as "any end can mitigate its risk to an acceptable level."

To assess the suitability of the DGSA as a model for government agencies, industry, and the private sector, this study presents an overview of the DGSA and provides a cursory comparison of some of the features of the DGSA with several prominent, proposed non-military security architectures.

¹"Reengineering Through Information Technology" (an accompanying report of the National Performance Review), Office of the Vice President, September 1993.

²"Information Security and Privacy in Network Environments," U.S. Office of Technology Assessment, September 1994.

³"Realizing the Information Future: the Internet and Beyond," Computer Science Technology Board, National Security Telecommunications Advisory Committee, National Research Council, 1994.

1.2 A Framework for Evaluation of the DGSA for Non-Military Use

To assess the potential of the DGSA as a set of guidelines for non-military security architectures, specifications of several existing or evolving models were investigated in order to contrast philosophies and practices. Among the more publicized of non-military security architectures are the Internet Security Architecture (ISA), Distributed Computing Environment (DCE), SESAME, and the International Standards Organization (ISO).

Commonly agreed upon methods or metrics for comparing security architectures are not available. Determining a suitable basis or criteria for comparison is difficult since specifications (available from reports and papers) often deal with a limited portion of the total security architecture (e.g., communications security architecture or methods of authentication). Also, the views projected by various reports are often at incomparable levels, e.g., view of abstract architecture vs. view of implementation architecture.

Several approaches to comparing two security architectures were considered. First, one might investigate the "high-level requirements" of each architecture. This view attempts to project the larger picture of security. Judgements might be made as to whether architectures have equivalent basic objectives. The difficulty with this approach is that needed details do not surface, e.g., two architectures might both claim to preserve the integrity of data, but the extent to which this is done and their corresponding strengths cannot be determined. A second potentially useful approach might be to contrast the "security services" offered by the architectures in an effort to determine functional equivalence. Such a comparison is difficult without a suitable metric to measure and contrast the strength of these services. Furthermore, the view of security in the larger context might be lost without knowledge of how the existing services work together harmoniously to implement total security. A third approach might be to itemize the "security concepts" upon which the security services are constructed, and contrast the strength and the applicability of the concepts in each architecture. This has very strong appeal, but even if the same concepts are present in two architectures, there is no assurance of the relative strengths of their implementations. The author of this report believes that comparisons are needed at all three levels (and perhaps more) in order to accurately contrast one security architecture with another.

Because common descriptive information was difficult to locate, the scope of this report is limited to a cursory comparison of the DGSA with several other security architectures with an attempt to itemize any discovered differences or incompatibilities in the high-level requirements, security services, and security concepts.

2. Overview of the DGSA

The DGSA serves as a guide to system designers, developers, integrators and managers regarding the DOD requirements for a security technology for computers, networking and telecommunications. It is the target security architecture for Defense Information Systems (DIS) and is applicable to all individual programs and systems in order to meet the objective of interoperability. The scope of the DGSA includes both automated security services for end systems and communication media, and non-automated services for physical, environmental, procedural, and administrative security. It addresses hardware/software requirements for a wide variety of computing components, communication media, and human interfaces – involving varying degrees of bandwidth requirements, flow

characteristics, error susceptibilities, and security. Within this framework, the goal is to maximize protection, consistent with perceived threats at an affordable cost and to provide the basic security services of confidentiality, integrity, identification and access control, non-repudiation, and availability. The brief description of the DGSA in this section is derived from a more comprehensive discussion.⁴

2.1 Policy and Fundamental Requirements

The DGSA incorporates policies and concepts that are different from present DOD security practices and technologies. The most fundamental change stems from the adoption of an information domain that requires strict isolation and the accompanying principles of isolation and absolute protection in order to meet the DOD security requirements. The result is a shift in the DOD information security architecture paradigm from a kernel that supports a multi-level security policy to that of a separation kernel that enforces the isolation of domains, each with its own unique security policy. While closed system environments will still be desirable under certain conditions, increased emphasis will be placed on the use of open systems and distributed information processing, commercial and government off-the-shelf products (COTS and GOTS) to minimize the cost, and common communications carriers.

Based on an underlying DOD security policy, the four fundamental DGSA security requirements are:

Multiple Information Security Policy Support

There must be information systems and communications systems support for multiple security policies of arbitrary complexity. Refinements of the requirement are:

- a. Enforce security policy.
- b. Maintain user identities.
- c. Maintain information identification .
- d. Provide data integrity service.
- e. Provide data confidentiality.
- f. Provide non-repudiation service.

Open Systems Employment

Open systems for information processing and communications systems are critical in order to facilitate the sharing and transferring of information among a collection of users with diverse missions and geographic locations. This is a distinct departure from the past practices of over-classifying information and isolating systems. International standard protocols (or at least national or DOD standards), information, and mechanisms will enable users to determine the capabilities and environment of other users or systems processes with which they will attempt to communicate.

⁴“Department of Defense DOD Goal Security Architecture (DGSA),” Draft Version 1.0, August 1, 1993, Center for Information System Security, Defense Information Systems Security Program.

Refinements of the open systems requirement include:

- a. Provide common security capability identification.
- b. Use standard security information exchanges.
- c. Use standard security information representations.
- d. Provide authentication service.
- e. Provide access control service.
- f. Provide availability service.

Appropriate Security Protection

Security mechanisms must be identified that implement security services at the level of protection dictated by the security policies. Specific means must be available to users to invoke security mechanisms appropriate for each security service, individually and collectively.

Common Security Management

Elements that must be managed are users, security policies, information, information processing systems that support one or more security policies, and security functions that support the security mechanisms that implement security services. Security management provides the foundation for security administrators to manage, in a uniform way, systems that operate under multiple policies.

Refinements of the requirement include:

- a. Identify and maintain user information managed objects.
- b. Identify and maintain information system managed objects.
- c. Identify and maintain supporting security function managed objects.
- d. Use standard managed object representations.

2.2 High-Level Abstraction of Hardware/Software Support

The highest level of architectural abstraction depicts a view of the DGSA as a collection of Local Subscriber Environments (LSEs) communicating via communication networks (CNs). The protection of an LSE is doctrinal, that is, it is provided by physical, administrative and personnel security mechanisms. At the next level of abstraction, the identifiable components within the LSE are:

End Systems (ESs)	e.g., workstations, servers, mainframes, telephones, radios
Relay Systems (RSs)	e.g., multiplexors, routers, switches, cellular nodes, message transfer agents
Local Communication Systems (LCSs)	e.g., rings, buses, wire lines

An LSE may be as simple as a workstation or a single RS such as a router, or it may be a complex interconnection of ESs and RSs through an LCS.

2.3 Security Services

The basic security services identified by the DGSA are based on standards and protocols defined by ISO and on modifications and additions for local area network (LAN) security developed by Institute of Electrical and Electronics Engineers (IEEE) standards committees.⁵ These services include authentication, access control, data integrity, data confidentiality, and non-repudiation. In addition, the DGSA provides availability as a basic service.

The ESs/RSs collectively provide such security services (both for connecting users and connecting communications) as identification and authentication, access control, confidentiality, integrity, and availability. The LCSs (controlled within the LSEs) are responsible only for availability of communication among ESs and RSs within the LSEs. The CNs (metropolitan, regional, or global) may consist of private links between LSEs – closed systems owned, leased, or operated by the DOD (e.g., Defense Communication System) with complete traffic flow security, or the CNs may use common carriers so that one end of the link terminates in a commercial zone – open system (where address information is needed by switches) with more limited traffic flow security. Components of a CN may include transmission systems, switches, routers, gateways, management systems, and network-specific servers. Some agreed-upon level of availability (responsiveness, continuity of service, and resistance to accidental and intentional threats) is assumed of CNs that utilize common carriers. The transfer system (including the communication protocols integrated into the ESs/RSs and interconnecting LCSs and CNs) is responsible for peer entity and data origin authentication, access control, non-repudiation, confidentiality, integrity, and availability of in-transit information.

Trusted LSEs are subjected to stringent evaluation during certification and accreditation. They are built on fundamental concepts of “information domain” and “separation kernel” and must enforce “strict isolation” and “absolute protection” (see Section 2.4). While the architecture of DGSA is considered theoretically feasible, these concepts transcend current state-of-the-art security technology and practices. (The use of present-day workstation technology in the DGSA would require that each workstation operate in isolation.) Because the DOD wishes to use off-the-shelf components whenever possible, a critical aspect for cost-effective implementation of the DGSA is to convince commercial vendors to construct the necessary components that allow implementation of the required security concepts.

Assuming only availability of the CN, the security architecture for communications in the DGSA is based on binding “security contexts” over distributed transfer “security associations.” For two entities communicating over the CN, security contexts are cohesively bound by security management protocols and the separation kernel (within the LSEs) to formulate the basic framework of end-to-end protection. The DGSA’s security association employs a virtual secure communications channel. This security association involves the information domain’s sensitivity label and secure communication attributes (traffic encryption key, cryptographic algorithm identifier, integrity key). A security association is established using a Security Association Management Protocol (SAMP), which is called upon by a Security Management Application Protocol (SMAP) in the information domain’s processing space.

⁵International Standard ISO 7498-2-1988(E) Security Architecture and LAN security standards developed by IEEE 802.10 committees.

2.4 Basic Security Concepts

The four fundamental security requirements, listed in Section 2.1, are supported by the following seven basic security concepts:

Information Domain

An information domain is a set of users (members), their information objects, and a security policy. Information domains are not bounded by systems or even networks of systems. They are bounded only by the presence of identifiable objects and may be supported on any system that can provide the protection requirements dictated by the security policy. Each domain is uniquely identified. All information objects in a domain have the same security-relevant attributes. Members may have different security-related attributes, e.g., read, write, read/write. A member who has an access privilege in an information domain has that privilege for every object in the domain.

Strict Isolation

Information systems that support multiple domain security policies must provide a basis for satisfying all of them. The strategy of the DGSA is to completely isolate any two domains that do not have an established explicit relationship (dictated by each of their respective policies). One consequence of strict isolation is that many aspects of covert channels, both timing and storage, are more easily controlled.

Interdomain Sharing and Transfer

The simplest method for sharing is to accept new members into an information domain and grant them access privileges. If there is a need to share only a subset of the objects in a domain, a new domain may be created to contain these objects. Objects can be transferred (moved or copied) between two domains in accordance with the security policies of each. Transferred objects must be relabeled with the label of the recipient domain. Interdomain transfers can occur only within an end system (ES or RS); they cannot occur among distributed end systems. Transfers between end systems can occur only within the same domain.

Multidomain Information Objects and Policies

Information domains are not hierarchically related nor do they infer (explicitly or implicitly) a sensitivity related to multiple categories. Users who operate in more than one information domain may need to concurrently process objects from several domains. While multidomain objects never exist on an end system (ES or RS), a user may still have a “perception” that a collection of information objects from different domains form a single composite object (e.g., for display, printing, conveying to another information system). This perception must be realized without the actual combining of real information objects, and strict isolation between domains must be maintained. Explicit multidomain security policies must state what privileges a user must have to view, print, delete, or transfer a multidomain object between systems. Implementations may allow for a “description” of a multidomain object to be held in an information domain.

Absolute Protection

The concept of absolute protection provides a framework for achieving uniformity of protection in all information systems supporting a particular information domain. In order to support information domains in multiple (heterogeneous) LSEs, the overall strength of protection must be consistent in

those LSEs. Strength of protection is a function of strength of mechanisms (including doctrinal) implemented in an LSE to satisfy an information domain security policy.

Uniform Accreditation

Each LSE is evaluated against the security policy of each information domain it supports. The objective is to have equivalent protection in all LSEs that support a given information domain.

Security Management

Security management is a particular instance of information system management and is central to the proper operation of protected LSEs and their component parts, separately or jointly. Security management is concerned with all aspects of protection within and among LSEs (security policy management, security service management, security mechanism management - including doctrinal, security mechanism support management - such as key distribution and management, transfer security management). In general, the DGSA adopts the ISO standards.⁶

Unlike traditional approaches, which assume that all users of an end system are subjected to the same security policy, end systems in the DGSA may be required to support multiple information domains and independent security management for each of these domains. An end system security policy must also specify how to accomplish sharing of security functions and resources among the information domains.

Security management in an end system is concerned with the installation, maintenance, and enforcement of security policy rules and the information about users, and security services and mechanisms needed to implement a security policy. Some integral portions of security are not performed within the end systems (e.g., doctrinal security).

End system security policies must specify management rules for 1) providing strict isolation among domains, 2) invoking and managing security mechanisms that implement security services, 3) developing rules for management of multidomain information objects including criteria for user access, display, labeling, and transfer between end systems, and 4) controlling and maintaining security management objects that enable security managers to control the information domain independently of others.

Example elements of security policies are 1) a description of mission and mission functions, 2) description of information objects and their attributes along with rules as to their use in multidomain objects, 3) membership criteria, 4) interdomain transfer rules, 5) security service requirements to meet risks and counter threats, 6) criteria for acceptable mechanisms to implement security services, and 7) security management-specific requirements (e.g., relationship of security management information domain to an information domain, identity and membership rights of security administrators, configuration management requirements for establishment or modification of security policy rules, identification of members responsible for accreditation of information systems that will support the information domain).

⁶OSI Management Framework, ISO 7498-4 and the security management portion ISO 7498-2.

2.5 Components of an Implementation Architecture

For implementation of the DGSA, hardware/software architectural support is required to realize

- a. the notion of a security context – combination of all LSE, hardware, system software, application software and information supporting activity in an information domain (subject to a security policy),
- b. a separation kernel that creates separate address spaces to enforce strict separation of security contexts and controls communication among security contexts, along with standard kernel interfaces for service requests.

The notion of a separation kernel and the accompanying requirements of evaluation and accreditation of the security functions impacts a substantial portion of traditional operating systems, e.g., memory management, file management, display management, interprocess communication, process scheduling (must ensure availability of service to all processes), audit, etc.

The notion of distributed security context must support both interactive and staged classes of communication. An interactive distributed security context is established through a set of mechanisms called a “security association” – the totality of communication and security mechanisms needed to bind together two security contexts in different ESs/RSs supporting the same information domain. This is an extension of the OSI notion of an “application association”. The possibility of a SAMP three-phase protocol is suggested to establish a connection, exchange keys and perform preliminary security checks, and employ encryption algorithms and keys to test the liveness of the association.

A staged delivery distributed security context is transferred from one end system to another by cryptographically wrapping/unwrapping information on the sending/receiving end. One candidate for achieving this is an existing specification for secure electronic mail, Secure Data Network Service (SDNS) Message Security Protocol (MSP) specification.

The transfer between end systems of multidomain objects requires a distributed security context for each domain component and for the relationship among the components (if the latter is contained in a separate domain).

Securing the transit system will require low-cost cryptographic devices sufficiently flexible to support the requirements of different domains (multiple cryptographic algorithms, multiple key management schemes, public key systems, key distribution center schemes) and traffic flow security.

At present, no commercially available end systems (combined ESs and RSs) provide sufficient support to satisfy the abstract DGSA requirements. However, current efforts in the academic and research communities do support some aspects and commercial operating system vendors have recently adopted design strategies that share significant aspects of the end system security architecture. Advances in technology are expected to continue to provide solutions but the evolution of the security requirements are likely to make research and development an ongoing process.

3. Comparison with the International Standards Organization Security Architecture

Working with “high-level” requirements, it should be noted that the DGSA attempts to be compliant to a large extent with applicable ISO standards for open systems along with standards for LAN security and other ISO protocols. With regard to ISO standards for communications, layers 1-3 satisfy the needs of the DOD while DOD missions address standards needed for ISO layer 4 and above. Focusing strictly on communications, and using the high-level security requirements of Section 2.1 as the comparative basis, the two architectures compare favorably in the comparison criteria 1 and 2, namely, “multiple security policy support” and “open systems employment.” In the latter area, the DGSA would interconnect with “all open systems of interest” while the OSI would presumably restrict its attention to “OSI-protocol-compatible systems.”

The ISO General Upper Layers Security (GULS) allows for the concept of subdomains and superdomains. Because the DGSA disallows subsets and supersets, and deviates in minor ways (while at the same time extending and elaborating) in security management, the OSI may not satisfy the DGSA’s high-level security requirements 3 and 4, namely “appropriate protection” and “common security management.” In other words, the DGSA is more restrictive and a system complying with the ISO requirements might not satisfy the DGSA requirements.

The DGSA significantly expands the scope of the three ISO 7498-2 areas of security management for communications (system security, security service, and security mechanism) to include “all open system areas of interest” (not just the OSI aspects of open systems). In the DGSA, the functions of these areas are performed within the LSE. The DGSA extends the OSI list of security services (authentication, access control, data confidentiality, data integrity, and non-repudiation) to also include the service of availability. The DGSA includes doctrinal mechanisms for protection of the LSEs. These expansions and extensions are not, in themselves, necessarily incompatible with ISO.

Whereas OSI employs “application associations,” the security architecture for communications in the DGSA is based on binding “security contexts” over distributed transfer “security associations.”

The discussion of ISO extensions, expansions, and incompatibilities appears in the general discussions of management of security information in the DGSA draft report. Identical titles of subsequent subsections are used in this report to aid the reader in tracking the summary discussion with the details of the draft report.

3.1 Security Management Concepts Comparisons

In the DGSA, security management includes all components of open systems. In a self-assessment, the description of the DGSA carefully itemizes the points at which it restricts, clarifies or extends ISO specifications.

3.1.1 Information Domains Concept

The DGSA and ISO 7498-2 agree in the abstract definition of an information domain. ISO leaves the elaboration of the domain concept and the interactions between domains to future extensions. The

ISO model ⁷ allows for (but does not mandate) supersets and subsets to form hierarchies of domains. This implies that information domains can have multidomain objects, or the need to express objects within a domain with different sensitivities.

The DGSA does not allow for hierarchies or orderings of domains because labels bound to objects in a domain must represent a singular sensitivity. Therefore, strict isolation of domains is a requirement. Since singular sensitivity is not mandated by the superset/subset domains in ISO, the DGSA is considered more restrictive.

A philosophical difference that would be highly visible to a user is that ISO allows for discretionary (owner-controlled) access control, e.g., ACLs and capability lists. In the DGSA, a member of a domain has the same access to all objects in the domain.

3.1.2 Security Management Information Bases (SMIBs)

In the ISO specifications, an SMIB must have all the security-relevant information to enforce an appropriate security policy and may be distributed to the extent that a consistent policy is enforced over a (logical or physical) grouping of end systems. It may or may not be integrated with the Management Information Base.

The DGSA uses Security Management Information Bases (SMIBs) for both information domain and end system management. A distinct security management information domain may be responsible for management of single information domain or for several information domains, or it may be embedded in the information domain along with its objects. An SMIB might contain such things as information about security policy rules, information about members (registration, authentication criteria, attributes), sensitivity labels to attach to displayed information, or security service requirements for specific applications such as intradomain communications and interdomain information transfers. For an end system, an SMIB might contain such things as security policy rules, security services management information, or supporting services (alarm reporting, auditing, cryptographic key distribution).

3.2 Communication of Security Management Information

The ISO statement is that management protocols (in particular security management protocols) and the communications channels carrying the information are potentially vulnerable and should be protected.

The DGSA more specifically states that all security management information must be protected in accordance with the security policy of each information domain. Management protocols that are used to distribute security information rely on the same open systems protocol infrastructures as other applications and rely on “security associations” between security contexts for communication between distributed end systems.

3.2.1 Distributed Security Management Administration

The DGSA and ISO standards agree in language and scope.

⁷OSI Security Framework in Open Systems, ISO 7498-2.

3.2.2 Security Management Applications Protocols

The DGSA and ISO standards appear to agree in principle. Ultimately, the DGSA plans to adopt Government Open Systems Interconnection Profile (GOSIP) recommendations/mandates.

3.2.3 End Security Management Functions

ISO defines the concern of system security management to be the management of the security aspects of the overall OSI environment. A typical list of activities is overall security policy management, interaction with other OSI management functions, interaction with security service management and security mechanism management, event handling management, security audit management, and security recovery management.

The DGSA broadens the view of end systems security management to open systems environment, especially support of multiple information domains.

3.2.4 Security Services Management

ISO 7498-2 defines Security Services Management to typically include determination and assigning strength of service, assigning and maintaining rules for mechanism selection, negotiating available security mechanisms (local/remote), invoking security mechanisms, and interacting with multiple security service management functions.

The DGSA accepts the ISO definition and elaborates on specific points, noting that an information domain security policy may be very specific as to how security requirements are to be met (by mandating specific security functions) or it may be very general and allow end system management functions to select an appropriate mechanism from those available.

3.2.5 Security Mechanism Management

ISO 7498-2 specifies the following list as being typical of security mechanism management functions: key management, encipherment management, digital signature management, access control management, data integrity management, authentication management, traffic padding management, routing control management, and notarization management. The DGSA adds availability management to this list.

In ISO 7498-2, key management involves generation of suitable keys at intervals commensurate with requirements, determining which entities receive a key in accordance with access control requirements, and distributing the keys in a secure manner. Some functions will be performed outside the OSI environment, including physical distribution by trusted means. Selection of working keys may require a key distribution center or pre-distribution by management protocols. The DGSA statement clarifies the ISO statement by stating that it will incorporate all standard key management techniques, specifically the evolving Security Association Management Protocol (SAMP). Also, key management mechanisms in the DGSA extend to infrastructures beyond those stated or alluded to in the ISO statement; of particular applicability is the evolving Electronic Key Management System (EKMS), from which the majority of U.S. government keying materials are generated, distributed and accounted for.

The DGSA supports the ISO 7498-2 definitions of encipherment management and digital signature management; there is a great deal of similarity between the two. The registration of cryptographic algorithms will also be needed along with rules for changing algorithms and the audit of their use. For non-repudiation, additional security management responsibilities will be needed for archiving keys and identifying algorithms.

ISO 7498-2 notes that access control management may involve distribution of security attributes (including passwords), updates to access control lists or capability lists, and use of a protocol between communicating entities that provide these services. The DGSA also includes the initial installation of access control attributes in the SMIB and the communication of attributes between end systems (if not all SMIB information is local).

The ISO view of integrity management involves cryptographic techniques. While noting that cryptography is imperative for communications, the DGSA broadens the view of data integrity management, noting that, in some instances within a single system, data integrity can be attained/maintained by strong access control mechanisms. For communications, similarity is noted between data integrity management and encipherment management.

ISO 7498-2 states that authentication management may involve the distribution of descriptive information, passwords, or keys to entities performing authentication, and may involve a protocol for communication. In the DGSA, this information is stored in the SMIB, if necessary. It is not required by some domain policy if "physical" identification and physical controls for access to an end system are deemed sufficient.

The DGSA and ISO seemingly agree in principle on the management of traffic padding.

ISO 7498-2 states that routing control involves the definition of secured or trusted links or networks. In the DGSA, if the end systems are connected to multiples CNs, the routing control is restricted to choosing a particular network interface.

ISO 7498-2 states that notarization management involves distribution of information about notaries and the use of a protocol for communicating and interacting with notaries. In the DGSA, notarization support is combined with non-repudiation.

ISO 7498-2 does not define availability management. The DGSA view is that availability is actually managed by other management functions. For example, unavailability of a communications path for a certain period, may trigger alarms to select alternate routing mechanisms or trigger an LSE to use infrastructure capabilities to restore communications availability.

4. Comparison with the Internet Security Architecture

4.1 Internet Security Architecture

The Internet is a world wide system of interconnected computer networks that share the TCP/IP suites and name and address spaces that are specified by the Internet Advisory Board (IAB) of the Internet Society. For many years, the Internet has existed primarily through the voluntary cooperation with minimal centralized oversight. Because there has been little or no effort made to coordinate security efforts, site security ranges from essentially none to fairly stringent (especially in some industrial networks). Because of the diverse needs of the the Internet community, it does not

seem likely that uniform site-specific security can be achieved on a voluntary basis in the near future. However, the commercialization of the Internet is changing the voluntary nature and, depending on market incentives, commercial service providers could establish security rules for its subscribers.

The IAB coordinates the overall architecture of the Internet and, through its subcommittees, it currently invests a substantial effort in developing a security technology for building effective firewalls (e.g., packet filtering, circuit gateways, application gateways) to insulate sites against intrusion.⁸ The Privacy and Security Research Group (PSRG) of the Internet Engineering Task Force (IETF) is developing recommendations for an Internet Security Architecture (ISA). The group is developing a document that is intended to serve as a technical guide for designing and implementing protocols for use in the Internet.⁹

The PSRG defines the ISA as “a plan and set of principles for establishing and maintaining features and mechanisms that protect against interruption and loss of packet-switched network elements, the communication service they provide, and the data they contain and carry.” The architecture focuses on communication and computer security and does not address some of the other areas and practices that are needed for complete security, e.g., trusted systems, physical and environmental security, procedural security, emanations security, risk management, and administrative security (including configuration management). The PSRG refers the reader to other documents for these areas.¹⁰ While being explicit on principles, the document is more implementation-oriented than the DGSA document.

Three systemic vulnerabilities of Internet protocols are identified: those caused by design, those caused by its implementation, and those caused by its operational management. The architecture emphasizes the elimination of those caused by design.

4.2 ISA Security Philosophy

The underlying philosophy in specifying the design of the Internet security architecture is based on four components:

Properties of the Architecture

The architecture must seek a careful compromise between generality and specificity in order to permit the analysis of candidate protocols that meet architectural requirements. The design should be as perfect as possible (admitting to the possibility of implementation and operational flaws), be scalable in security strength (encompass a range of security technologies that may vary in degrees of resistance to attack), be truthful in advertised strength (allow the user to be confident that proper implementation and management of a protocol will provide security over a broad range of environments), and include security features as mandatory options (must be implemented as part of the protocol although its use is optional).

⁸W.R. Cheswick and S.M. Bellovin, “Firewalls and Internet Security - Repelling the Wily Hacker,” Addison-Wesley, 1994.

⁹A partially developed set of notes, privately circulated, is entitled “The Internet Security Architecture (ISA),” Internet Privacy and Security Research Group, R.W. Shirey (ed.), Undated Draft.

¹⁰“Site Security Handbook” (RFC 1244) and “Guidelines for the Secure Operation of the Internet” (RFC 1281).

Principles for Security Systems

Long-standing and accepted underlying principles¹¹ for computer systems are adopted, including economy of mechanism, open design, separation of privilege, no single point of failure, least privilege, least common mechanism and least trust, fail-safe defaults, psychological acceptability, and built-in (not added-on).

Principles for Secure Internet Systems

Additional guideline considerations, not found in typical computer systems, are

- a. scalability of size

The design considerations should scale to the size of the world wide Internet.

- b. maximum interoperability

Multiple mechanisms for implementing a specific security service should be offered only when there are compelling reasons; and excessive number of options leads to non-interoperable implementations and confusion.

- c. preference for software

Designs that can be implemented in either hardware or software are preferred (in order to provide local option).

- d. perimeter selection

In network security, one can select a set of connection points that separate the community to be protected from the rest of the threatening world (e.g., to protect a community of host end systems, a perimeter could be established at points where hosts connect to the Internet; to protect a departmental or campus LAN, a security perimeter could be established at a router that connects the LAN to the regional wide-area network (WAN)). While end-to-end security affords the best protection against a wide range of attacks, it has the disadvantage of requiring deployment on a variety of platforms. In contrast, security at a perimeter eases the platform heterogeneity problems but is distant from end users, making it difficult to convey service requests and status between mechanism and user.

Trade Control and Patents

The architecture should not preclude the use of standard Internet protocols and the use of cryptography in Internet standards should not be dictated by existing government regulations. The use of patented technology should not be precluded but, if deployed, the patent must be licensed on a non-discriminatory basis and fees should not be onerous. Since patented technology has the potential for artificially restricting the application and use of standards, the Internet community should try to avoid standards that incorporate patented technology. Internet standards should also avoid use of technologies that are commercially proprietary, governmentally sensitive, militarily restricted, or otherwise prevented from being publicly disclosed.

¹¹J.H. Saltzer, and M.D. Schroeder, "The Protection of Information in Computer Systems," Proceedings of the IEEE, Vol. 63, No. 9, pp. 120-126.

In the area of security services, the ISA includes data confidentiality, data integrity, identification and authentication, access control, non-repudiation, and availability of communications.

Implemented in part by the use of common carriers, the Internet may be viewed by the DGSA as an open (or part of a larger) control network – with some level of assumed availability. Other security attributes will be solely dependent on the security of the communicating LSEs. By use of encryption techniques and other firewalls, the Internet can be used as an effective means for communicating unclassified or lower sensitivity information between trusted entities.

4.3 Comparison Discussion

Some corporate and nonmilitary government sites on the Internet are interested in high security. Because the DGSA is designed to provide secure end-to-end communication between trusted sites, the DGSA would seemingly in principle have some appeal. Potential impediments to its adoption are:

- a. the restrictive DGSA domain concept with built-in constraints on access control and object ownership and with the concept of strict isolation,
- b. the inability to order (partial or otherwise) domains and the inability to store a multidomain object,
- c. the additional costs imposed by high security in terms of up-front purchase price, maintenance, operation, and system performance – especially if some security features are not needed,
- d. delays in bringing to the marketplace affordable products that will provide for complete support of the DGSA requirements, and
- e. any lack of stability in the DGSA design specifications.

The following is a more direct comparison of some of the ISA and DGSA features:

- a. Major philosophical differences exist between the ISA and the DGSA. The philosophy of the DGSA is to provide a generic architecture under which one can secure an application to the extent necessary, consistent with its mission and a willingness to pay the costs. At one extreme, this would allow top-secret applications and information along with communication between trusted LSEs over proprietary networks and, on the other extreme, it would allow applications and information of lower sensitivity along with communication between untrusted LSEs over public networks (the only service provided by the network is that of availability and delivery). If availability could be guaranteed, the DGSA could confidently use the Internet for secure communication between secure LSEs.

The philosophy of the ISA is that perfect security may be unachievable but that security services should be developed and made available for participating sites or organizations to voluntarily employ according to their needs and capabilities. Many of the security services mentioned at the various protocol levels are similar in name to what the DGSA would require, e.g., confidentiality and integrity of data, authentication and identification, secure routing, secure associations (remote procedure calls and remote executions). Missing, however, from the ISA specifications are the security of end systems, compliance with the DGSA constraints on an information domain, and the accompanying principles of isolation and absolute protection.

- b. The ISA focuses on the automated aspects of computer and communications security. It recognizes three types of vulnerabilities – in design, implementation, and operational management.¹² The DGSA addresses, in addition, physical and environmental security, procedural security, and administrative security.
- c. The ISA and DGSA list the same six basic security services – data confidentiality, data integrity, identification and authentication, access control, non-repudiation, and the availability of the communications system. In the ISA, there is no discussion of security of hosts – so hosts have to be mutually trusting communicating entities. Without secure hosts, the strength of these services is admittedly diminished. The DGSA explicitly broadens the scope of the security services to include the LSEs, and adds (in addition to the above list) “security management” as a service that is central to overall security.
- d. The ISA deals only with intentional (intelligent) threats and attacks. The DGSA does not attempt to distinguish intentional attacks from other types of attacks.
- e. The ISA does not seem to incorporate the DGSA concept of an information domain. Some of the Internet literature refers to a concept of “administrative domain” at the level of host, site, or cell (see also the discussion of DCE).
- f. In the area of communications security, the ISA relates Internet communication layers to those of the OSI Reference Model (OSIRM) and discusses its own security features in the context of the Internet model. The DGSA subscribes to all of OSIRM, but assumes that the only security service provided by a (non-DOD) CN is that of availability.
- g. For end-to-end communication, the ISA adopts the view of the two hosts¹³ communicating via an Internet infrastructure. Whether or not hosts (or entities) are allowed to communicate is left to the discretion of the various sites or controlling organizations. The DGSA’s overall view is similar except the communicating entities are LSEs, and the possibility of communication is based on levels of sensitivity and proper certification and accreditation.
- h. In the ISA, the concept of “security perimeters” is introduced under a discussions of controlled entry points (gateways and firewalls) and end-to-end cryptographic security mechanisms. These allow for the implementation of security levels and the separation of a community of hosts from the rest of the Internet. In the presence of untrusted hosts, the limitations of end-to-end encryption are noted.

In the DGSA, the information domain, complete isolation, and the establishment of secure communications between domains are the most basic of all concepts. The policies of information domains describe the necessary isolation or connection to other domains. The DGSA uses SAMP to establish a secure association between trusted LSEs over an untrusted (but available) network.

¹²Other Internet reports have been published that deal with such things as recommended security practices (including doctrinal security).

¹³Ranging from PCs to supercomputers, a host is assumed to have the capability to execute relay protocols needed to connect to an intermediate system, and to execute end-to-end protocols needed to communicate with another host.

5. Comparison with Other Selected Security Architectures

5.1 General Comparative Remarks

The European Commission (CEC) commissioned a group of organizations (including Bull, DEC, HP, ICL, Perihelion, and Siemens Nixdorf Informationssysteme) to conduct a comparative study of security architectures. This project is called COMPOSITE, for COMParison Of Security Information TEchnologies.

A summary report¹⁴ discusses the similarities of three architectures: Distributed System Security Architecture (DSSA), Open System Foundation Distributed Computing Environment (OSF DCE), and Secure European System for Applications in a Multi-Vendor Environment (SESAME). According to Per Kaijser and confirmed in part from available literature on the OSF DCE^{15,16}, these architectures all provide basic security support for access control, data integrity, data confidentiality, auditability of security-related events, management of security in a secure way, cryptographic support, and key distribution.

The report is cursory with few details, thereby preventing the application of any of the comparison criteria mentioned in Section 3. However, one can contrast the three architectures with the DGSA in several respects:

- a. In the DSSA, DCE, SESAME:
 - i. A server is assumed to trust information from a limited number of authorities, set of security servers, and security facilities.
 - ii. PCs and end workstations are considered untrustworthy by the system and may or may not be trusted by a user.
 - iii. End-to-end security is employed; there is no reliance on the security of the underlying network.
 - iv. The DCE adopts the Kerberos authentication mechanism.
 - v. Availability is not listed as a security service.
- b. Correspondingly, in the DGSA:
 - i. Users of a domain are trusted by the domain according to the domain policy. Other domains are trusted only to the extent allowed by the relevant inter-domain policies.
 - ii. Both secure and insecure LSEs exist. Secure LSEs have undergone certification and accreditation. Doctrinal protection applies to secure LSEs.
 - iii. End-to-end security is employed for communication over common carriers. There is no reliance on the security of the underlying network, except for some degree of availability.
 - iv. The DGSA does not specify any particular cryptographic method of authentication.
 - v. Availability is a security service.

¹⁴Per Kaijser (Siemens Nixdorf, Germany), "Secure Open Systems," Computer Fraud & Security, November 1992.

¹⁵OSF Distributed Computing Environment, Technical Seminar.

¹⁶OSF DCE 1.0 Introduction to DCE.

5.2 Further Comparison with DCE Security Architecture

5.2.1 DCE Overview

The purpose of the DCE is to provide a “relatively seamless” distributed programming and computing environment for its users with the following benefits:

- a. interoperability and portability across heterogeneous platforms

A process running on one computer can interoperate on a second computer, thereby making applications using DCE portable to any hardware/software platform running DCE.

- b. tools and services for developing and running distributed applications (e.g., remote procedure call, directory/security/time services)
- c. integration and comprehensiveness of DCE components

In depending on each other, many DCE components are themselves distributed applications, addressing inherent problems such as data consistency and clock synchronization.

- d. global interoperability

Users within DCE are allowed to access to standard services outside of DCE.

Architecturally, the DCE is a layer between distributed applications, on one hand, and the OS and transport layers, on the other. The DCE components consist of distributed programming facilities (threads, remote procedure calls), core distributed services (directory service, security service, time service) and extended distributed services (file service, diskless support service, and personal computer integration service). The core services must be present in each DCE cell (group of machines that work together and are administered as a unit). Each of the extended services is optional and depends on the core services.

5.2.2 DCE Security Services

In DCE, a “cell” consists of user, server, and administrator machines – functioning as an administrative domain with a single security policy that incorporates the concepts of object ownership and discretionary access control. The DCE server machines provide directory, security, and time services. Intercell communication takes place via global directory services.

The DCE Security Service provides three aspects of security support in a distributed system: authentication (based on the use of Kerberos), secure communications of RPCs (for integrity and privacy), and authorization via the use of ACLs. It consists of several cooperating (sub)services:

- a. registry service – manages user, group and count information and stores administrative policies regarding characteristics of accounts that can access the distributed system;
- b. authentication service – (based on Kerberos) allows principles defined as accounts in the user registry to exchange credentials and establish mutually authenticated communications; and

- c. authorization, consisting of
 - i. access control list facility – (based on POSIX) each DCE component implements its own ACL manager to arbitrate access to its objects, and
 - ii. privilege service – a trusted certification authority that derives authorization information about principles and packages this information into a privilege attribute certificate which is placed into a ticket for use of the resource by the requester.

5.2.3 Additional Comparative Remarks

In addition to the general comparative remarks in an Section 5.1, the DCE contrasts with the DSGA in several fundamental ways:

- a. In DCE, the administrative unit is the cell (realm) which is different from that of an information domain in the DSGA. There is a single information security policy within a cell – support for multiple policies does not exist.
- b. The attending concepts of strict isolation, absolute protection, and sharing and transfer between domains do not implicitly exist in the DCE.
- c. The DCE Security Service manages the identities and rights of a user within a given cell. Authorization is controlled by individual resource managers. In the DCE, the owner of an object is the sole authority who decides the access rights of others to that object. Explicit to each host machine is the concept of a “super-user” who can gain all access rights to all objects in the resident file system. (In the DSGA, there is no concept of individual object ownership or a “super-user” unless explicitly called for by a security policy.) In DCE, there is no inherent concept of sensitivity levels of objects and a user may have different access rights to different objects. The formation, storage, manipulation, and printing of a “multi-cell” information object is possible if a user has the appropriate access.

6. Summary Observations

- a. The DSGA incorporates policies and ideas that are different from the current DOD security practices and technologies – information domains, absolute protection, and strict isolation. This forces a change in the DOD information security paradigm, a shift from multi-level security kernels to separation kernels, a displacement of closed system environments to open systems and distributed information processing, and maximizing the use of COTS equipment and common communications carriers.
- b. Although intended to address the needs of DOD, the DSGA is “security policy independent” and allows for multiple security policies. Thus, the DSGA paradigm has the potential for being applicable in both the civilian and commercial sectors. If it proves to be useful outside the military sector, vendors would have important incentives to produce COTS products that are DSGA-compliant. However, there are some attending restrictions that may counter this appeal.
 - i. The concept of an information domain may be too restrictive and, taken as it is, there would be a potential explosion of information domains.

- ii. The sharing of information between domains may be too restrictive.
- iii. The current technology is not adequate to implement the DGSA.
- iv. Accommodating diverse security policies leads to the adoption of the principle of strictly isolating the information domains. Multi-domain objects may be formed for the purpose of viewing or printing, but the formation of a multi-domain storage structure is limited to the storage of a "description" of the structure. The ordering of domains (partial or complete) to depict relationships is not allowed in the DGSA. These concepts have been traditionally important in achieving efficiency in some aspects of computer applications, including the structuring of information, the components of which may potentially be of different sensitivities (markings). Further study is needed to determine the extent to which this may be an impediment to its usefulness in non-military applications.
- v. Because much of the R&D in information processing security is fueled by funding from DOD, the DGSA is certain to influence vendor considerations. Many DGSA requirements are not satisfied by COTS components. Vendors must shift paradigms from Trusted Computer System Evaluation Criteria (TCSEC) to DGSA. But, if not compatible with needs of the non-military sectors, the cost-effectiveness may be in question since the production of DGSA-compliant components would not be leveraged by commercial and private demands.
- vii. The DGSA paradigm transcends current technology in that its implementation includes a separation kernel and evaluation and accreditation of the security functions – entailing a substantial portion of traditional operating systems, e.g., memory management, file management, display management, interprocess communication, process scheduling (must be able to ensure availability of service to all processes), audit, etc.
- viii. If the scope of the the National Research and Education Network (NREN) is limited to the federally owned and operated networks, then appropriate levels of security might be enforced through proper structuring of the network (see Appendix B).
- ix. The DGSA paradigm may well be useful in the non-military sections of the government. An important (and perhaps determining) factor will be the cost of certified and accredited systems and a secure segmented network vs. the cost of COTS.
- x. The Internet is international in scope and is managed only through the cooperation of its users. While the Internet Activities Board has been concerned with the aspects of security, the level of security proposed in the DGSA cannot be achieved in the Internet without a dramatic change in its philosophy. Reasonably strong (and perhaps acceptable) security, however, can be selectively achieved through the establishment of security perimeters and "trust-worthy" hosts. The same comments apply in general to the National Information Infrastructure (i.e. the extension of NREN to the universities, libraries, schools, laboratories, etc.).

Appendices

A. Other DOD Networking Plans

This section briefly describes a DOD plan for a heretofore unimplemented Defense Data Network (DDN) and a subsequent (Near-Term) Defense Informations Systems Network (DISN-NT)¹⁷.

The DOD lists eleven basic security requirements for automated information systems (AISs): accountability, access control policy, security training and awareness, physical controls, marking, least privilege, data continuity, data integrity, contingency planning, accreditation, and risk management. These are supported by the following security services: confidentiality (data, traffic control), data integrity, identification and authentication, access control, data origin authentication, non-repudiation, and availability.

Currently, the strengths of these requirements (and corresponding services) vary across four separate networks, one for transmission of unclassified information and three for transmission of classified information:

- MILNET services for unclassified (U) and unclassified-but-sensitive (UBS) information
- DSNET1 single-level services for Secret (S) information
- DSNET2 single-level services for top-secret (TS) information
- DSNET3 single-level services for top-secret/sensitive compartmented information (TS/SCI)

There are no direct connections between DSNETs or connections from any of the four DSNETs to other long-haul packet-switched networks, but each serves hosts that connect to AUTODIN. Only MILNET has physical connections to two non-DOD backbone networks, called Federal Inter-Agency Exchanges (FIXES).

Proposed in 1985, approved in 1987, but never fully realized was the Data Defense Network (DDN). The plan was to merge the DSNETs into one, forming a Defense Integrated Secure Network (DISNET). Employing a "BLACKER" End-to-End (E³) system to separate traffic of different classifications, the idea was to provide a single network structured of non-multi-level-security (MLS) components, to handle the full-range of sensitivities. The DDN Security Architecture called for two physically separated segments – MILNET (for unclassified use - U and UBS) and DISNET (for classified use – S, TS, TS/SCI). A trusted gateway would be used in the DDN to facilitate unclassified communication between MILNET and DISNET.

The DDN proposal was not implemented for several reasons, including prohibitive cost, non-existent technologies, and the need to extensively modify existing protocols to work with BLACKER.

A subsequent plan was submitted for the employment of a Defense Information Systems Network - Near-Term (DISN-NT) for the 1992-96 time frame¹⁸. Subsequently, new operations requirements have emerged affecting implementation plans for the 1992-96 time frame:

¹⁷ "Defense Informations Systems Network Near-Term Security Architecture," Defense Information Systems Agency, December 31, 1991.

¹⁸ DISN, without the NT, is used to refer to the network beyond the year 1996.

- (1) need to facilitate internetworking with other networks,
- (2) serve Government Open Systems Interconnection Profile (GOSIP) compliant user and host systems,
- (3) incorporate smart multiplexors (SMUX), BLACKER and other approved D³ devices, and IP packet routing (IPR) services, and
- (4) serve special user communities such as the Defense Messaging System.

The constraints for DISN-NT include attainability (in cost and technology), substantial agreement with the approved DDN security architecture and implementable from existing four network system, and successful addressing of SMUX and IPR services and other evolutionary network issues.

Common goals of the DDN and DISN-NT are to transmit unclassified information and unclassified-but-sensitive (and a restricted amount of secret information) on the MILNET, combine the existing DSNETs into a single DISN, keeping levels of classified communications separated by smart multiplexors (SMUXs) E³-encryption devices, and Internet Protocol Routers (IPRs).

A 1992 plan provides for a three-layered system:

- (1) SMUX Layer:

This lower layer of SMUXs provides point-to-point encryption. It protects UBS data, but classified data must be encrypted prior to entering the lower layer.

- (2) X.25 Layer:

This middle layer consists of X.25 packet switches, arranged in a two-segment structure – retaining the names MILNET and DISNET. MILNET will continue to serve hosts that run at the UBS level, but will also (via BLACKER or other E³-approved systems) transmit S-level information. DISNET will transmit TS and TS/SCI classified data.

- (3) IPR Layer:

The upper layer will consist of IPRs connected by high-speed trunks. There will be a physically separate, system high network for each classification level that receives service.

B. Relationship to Other Federal Networking Needs

A recent report by the MITRE Corporation reviews the security services needs of the federal networks and presents a plan for an internet that spans all of the federal agencies¹⁹. The immediate community served would include all federal and executive branch departments and agencies; the greater community served would include state, local, and foreign governments, public agencies (e.g., Red Cross), and federal contractors and potential contractors in both the public and private sectors.

The report notes that all federal agencies have common needs for confidentiality, integrity, authentication, access control, non-repudiation, and availability. But different users or agencies may need different levels of assurance that these services are provided. In addition, it was noted that federal networks need to accommodate multiple policies, connect to non-government networks, respond to priority transmission requirements, survive disaster, and be easy to use and manage.

The MITRE report proposes a four-segment security architecture of trusted switches and E³-guard gateways intended to provide different levels of security service for common use networks that serve a broad range of government applications:

(1) Open:

The government uses open networks to communicate with outside parties. It exercises no control and security varies from none (e.g., some academic laboratories) to strong (e.g., some corporations). Open networks naturally carry unclassified, but could carry higher (UBS, C, S) if E³-devices were employed at trusted hosts.

(2) Controlled:

Most intra-government data would be carried by the controlled network. It could carry unclassified data and selected UBS data (e.g., NSFNET).

(3) Protected:

Some government applications require higher levels of security and some agencies have their own protected networks in which they control all hosts and network components. By adopting a more formal policy, they should be able to carry UBS and with E³ help, data classified at C and S levels.

(4) Classified:

For higher levels of security (e.g., DISNET), there is a need to protect classified and sensitive information. The network should be able to handle the full range of classifications (UBS, C, S, TS, and TS/CSI). These architectures would utilize E³ systems.

Applications with special needs not met by the above grades would use dedicated networks (e.g. CIA, NSA).

¹⁹ "Federal Internet security segments: a context for the Defense Information Systems Network," (Second Draft), R. W. Shirey, (February 22, 1993).

C. MITRE Corporation Study of Transition Issues

A recent “white paper” by the MITRE Corporation discusses what is perceived to be transitional issues faced by the DOD in implementing the DGSA.²⁰ The MITRE report is considered highly valuable for this study²¹ because, in itemizing “potential” transitional problems faced by the DOD, it also brings to the surface possible areas to be visited

- (1) in contrasting and comparing the DGSA with other security architectures, and
- (2) in evaluating the usefulness of the DGSA paradigm for non-military systems.

Transition issues are grouped into three areas – technology, management, and policy. In the area of technology, the key elements are categorized into four classes in order of critical importance to the success of the DGSA. Class 1 deals with the issue of information domains and strict isolation vs. convenient sharing and structuring of objects. Class 2 deals with the feasibility of the DGSA requirement of absolute protection. Class 3 deals with communication protocols and carrier systems – adequacy of protocols to support security associations, standardized security contexts, and standardized security management protocols, and whether common carriers might in the future be expected by the DOD to provide more services than just “availability.” Transitional problems are described for three levels of management: high-level, program, and site/security officer. In the area of policy, transitional problems are discussed for implementing multiple security policies and for certification/accreditation. The following summarizes the discussion in the MITRE report.

Transition issues are grouped into three areas – technology, management, and policy

C.1 Technology

The key elements of the DGSA in order of most critical importance (to the overall success of the DGSA) are:

Class 1: Information Domains, Strict Isolation

An information domain is defined by a collection of information objects, a group of identified users, and a specific security policy that protects and controls access to the information objects. While similar to ISO-defined domains, the requirement of globally unique labels for objects must be satisfied. Under strict isolation, the ISO hierarchy of domains and information is not allowed; the question arises as to how information composed of data from two domains with different sensitivity labels can be efficiently and effectively processed. Methods must be explored for an information domain to share objects. There appear to be only two ways – through a well-defined transfer policy or through the use of multi-domain objects.

Class 2: Absolute Protection

The level of protection afforded by an information domain must be consistent at all sites and must be independent of the protection afforded another domain. This implies individual evaluation and certification of domains. Also, the policy for each domain must take into account the assumption of threats from a hostile environment.

²⁰DISSP Goal Security Architecture Transition Issues, Draft, Technical White Paper by the MITRE Corporation, March 31, 1992.

²¹The NSA/DOD response to MITRE white paper has not been made available.

Class 3: Communication Protocols and Carrier Systems

- Security Protocols

- Standardized Security Association Protocols

The ISO protocols and standards appear to support the development of needed protocols. A new requirement imposed by DGSA is the overlaying of information domains on security associations – the question arises as to whether the protocols will need to be extended.

- Standardized Security Contexts

The workstation technology will need to be developed to support the concept.

- Standardized Security Management Protocols

This is an area that has been under study in the standards community. The DGSA imposes the additional requirement that these protocols support the concepts of information domains and strict isolation. LSEs must have a capability for supporting Management Application Processes (SMAPs) and a common Security Management Information Base (SMIB).

- Common-Carrier Communications Systems

The DOD's requirement is to use common carriers whenever feasible, with "availability" as the only required security attribute. The question that is raised as to whether common carriers will be able to meet all future needs of the DOD, e.g., will the DOD at some point in the future require levels of security or integrity that cannot be met by end-to-end encryption.

Class 4: Communications Encryption and Workstation Architecture

- Data Communications Encryption

- Standardized Transport and Network Layer Encryption

In the DGSA, the transport or network layers provide encryption services for data communications between information domain and transfer system contexts. Commercial products are not available that will provide the features for a common security management capability. While feasible, cost-effectiveness is a question if not useful in the broader marketplace.

- Standardized Application Layer Encryption

Commercial products exist, such as Privacy Enhanced Mail (PEM), SDNS Message Security Protocols, and Secure X.400. Since connections between information domains may have to account for context-sensitive information, analysis is required to determine the adequacy of these commercial products.

- Workstation Architecture

- Separation Kernels

Though widely discussed, few implementations exist and none are available as COTS. Vendors have focused on implementing the requirements of Trusted Computer System Evaluation Criteria (TCSEC). The question arises as to whether previous vendor investments can be leveraged.

- Security Appliques

This new architectural paradigm opens the possibility for vendors to supply security in modular fashion for classes of workstations. Open questions exist as to what extent security appliques can themselves be secured against modification and to what extent their correctness can be certified.

- Embedded Cryptographic Engines

As a security applique, the previous question arises. Performance and cost-effectiveness are issues as special-purpose hardware may be required.

Other Technology Issues

- Trusted Applications

The allowable use of trusted applications (e.g., trusted DBMS) is not clear in the DGSA specifications.

- Insider Threat

The information domain concept suggests a reasonably small number of domains that strongly address the problem of external threats. However, the absence of “need-to-know” protection within domains tends to ignore the potential problem of “insider threats.”

- Audit Requirements

The utility of audit within the context of a domain is not clear, since the need-to-know principle is not applicable. Effective implementation of inter-domain transfers and access to multi-domain objects is not clear. The ability to satisfy DGSA requirements for network auditing needs to be assessed.

- Object-Oriented Design

The concept of strict isolation (and attending architectural feature of a separation kernel) appears to run counter to the current trends toward object-oriented designs. This poses as a potential conflict as commercial vendors are re-orienting their products toward the object-oriented design marketplace.

C.2 Management

Transitional problems are described for three levels of management.

C.2.1 High-Level DOD Manager

Managerial decisions must be attuned to the fact that the DGSA requirements are expected to change over time, constituting an evolving goal architecture – transition planning is a continuing process. Focus needs to be on end objectives with consideration of specific technologies left to Program Managers. An education program will be needed for both DOD program managers and DOD suppliers regarding new policies and technological ideas. Use of COTS products must be carefully articulated to ensure maximum flexibility and consistency, in light of the fact that many COTS products do not currently satisfy DGSA requirements. DOD resources may have to be allocated to providing testbeds for new emerging technologies, the specification and prototyping and proposing to standards bodies and incorporation in vendor products new standards (e.g., the Security Management Application Protocol). Some organization will need to be designated to be responsible for the management of uniquely global names for information domains and distributing them to various DOD sites/systems.

C.2.2 DOD Program Manager

The DOD Program Manager will have to ensure the definition and number of domains a specific system will support early in the development process since later changes will impact software development costs. Program-specific requirements will have to be integrated with those of DGSA. Transitioning from present bases to DGSA will be a major concern. Dealing with internal threats (apparently not covered by information domain policies) will be a concern.

C.2.3 Site/System Security Officer

The Site/System Security Officer will be directly responsible for such things as the management of systems that support multiple domains, audits and audit reduction, addition/deletion of users, changes in user privileges. The concept of a Security Management Application Process will need to be fully understood.

C.3 Policy

C.3.1 Multiple Security Policies

- (1) Rules are needed for guidance in defining appropriate domains to satisfy mission requirements.
- (2) In transitioning to the DGSA, methodology will be needed to support the definition and development of security policies to support inter-domain information sharing.
- (3) In order to control the potential explosion of domains and associated security policies, it will be necessary to balance the number of domains appropriate for a mission against complexity and number of policies.
- (4) It will be necessary to adhere to the current national classification policy with its accompanying label requirements (for multiple compartmented information) without an exponential explosion of the number of required information domains.

- (5) The DGSA will have to support human-readable labeling requirements as mandated by national policy.
- (6) It is not clear how (or if) ad hoc discretionary policies (which allow owner definition of access) can be supported.
- (7) Study must be given to who will create, evaluate, and certify DGSA security policies.

C.3.2 Certification and Accreditation

- (1) A new assessment "valuation" will be needed to determine the value of the information in a domain and to define the level of assurance, access controls, and other security attributes, required of the system to support the domain.
- (2) A method must be established for assessing the property of "separation."
- (3) Appropriate methodologies must be developed to reduce the effort required for valuation, certification, and accreditation. To illustrate the potential effort required, the MITRE report gives an example of three systems S1, S2, S3 and four domains D1, D2, D3, D4. D1 is implemented in S1 and S3, D2 in S1 and S2 and S3, D3 in S2 and S3, and D4 in S2. Three certifications are required (for S1, S2, S3), four valuations are required (for D1, D2, D3, D4), and eight accreditations are required (for D1 in S1, D2 in S1, D2 in S2, D3 in S2, D4 in S2, D1 in S3, D2 in S3, and D3 in S3).

