

Determining Security Requirements for Complex
Systems with the Orange Book

Carl E. Landwehr

Computer Science and Systems Branch, Code 7593
Information Technology Division
Naval Research Laboratory
Washington, D.C. 20375

H. O. Lubbes

Computer Resources Division, Code 814T
Space and Naval Warfare Systems Command
Washington, D.C.

ABSTRACT

The DoD Trusted Computer System Evaluation Criteria define requirements corresponding to specified levels of security functions and assurance. They do not, however, help determine what level system is required for a specific environment. A simplistic technique has been proposed for this purpose that takes into account only the classification of the most sensitive information processed by a system, the clearance of its least-cleared user, and the environment in which it was developed. This paper offers a straightforward but richer technique a developer can use to map a specific system architecture and application environment to a particular requirement level as defined in the Criteria. It accounts for differences in functions provided to different users and the ways users can invoke those functions, as well as for users' clearances and the sensitivity of data. This technique is applicable throughout the system life cycle, so that security requirements can be updated as changes to system structure and function occur.

1. Introduction

This paper presents a method for determining the hardware and software security requirements of a system, based on

- (1) the local processing capability available to a system user;
- (2) the kind of communication path between the user's local device and the primary system components;
- (3) the flexibility of the processing capability the system provides to the user;
- (4) the environment in which the system was developed; and
- (5) the difference between the clearance held by the least cleared user of the system and the classification of the most sensitive data processed by the system.

This method can be understood as a risk evaluation of a system that can be conducted at a very early stage in the life cycle of a system and repeated as the structure and functions of the system change during its development and operation. Depending on the inherent risk that a system (or system design) displays, dif-

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 1985		2. REPORT TYPE		3. DATES COVERED 00-00-1985 to 00-00-1985	
4. TITLE AND SUBTITLE Determining Security Requirements for Complex Systems with the Orange Book				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Code 7593, 4555 Overlook Avenue, SW, Washington, DC, 20375				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

ferent levels of security requirements may be imposed in order to reduce the operational risk of the system to an acceptable level. Applications of this method to several environments are provided as examples. Although developed based on consideration of DoD environments, the method seems applicable to other environments to the extent that the Orange Book requirements apply to them.

The technique described here does not consider requirements for degaussing of removable storage, TEMPEST requirements, protection from physical hazards, emergency destruction, or other security requirements not related to the hardware and software architecture of the system.

2. Structure of the Orange Book

The DoD Trusted Computer System Evaluation Criteria (the ``Orange Book'' [1] provides a set of security requirements of two kinds: specific security feature requirements, which call for particular system functions in order to provide data security, and assurance requirements, which call for testing, documentation, and verification to assure that the security features are correctly implemented. A system that satisfies all requirements listed in the Orange Book would be designated A1. Systems that satisfy specified, nested subsets of the requirements are designated B3, B2, B1, C2, C1, D, in order of decreasing requirements.

The Orange Book does not provide guidance as to what level of system is appropriate for a particular operational environment. A draft application doctrine [2] has been developed, however, that defines the level of system required for a particular environment based only on the classification of the data processed by the system, the clearances of its users, and the environment in which it was developed. This simple scheme is inadequate for use in assessing security requirements of many complex systems; a more comprehensive method is proposed below.

3. Applying Technical Computer Security Guidance Effectively

Although it is imperfect in many respects, as a technical basis for specifying computer security requirements, the Orange Book is the most comprehensive and current document available. A method is needed for applying the Orange Book to the components of large scale, geographically dispersed systems, so that the appropriate requirements from the Orange Book can be identified for each host system. Such a method is defined below. As shown in Figure 1, it involves:

- (1) extracting from each system (or system design) the factors that affect the risk that its operation may lead to the unauthorized disclosure of sensitive information,
- (2) quantifying these factors, and
- (3) determining system security requirements (in terms of the levels defined in the Orange Book) that reduce the system risk to an acceptable level.

This method can be understood as a risk evaluation based on the threat of unauthorized disclosure of sensitive information. The asset of the system is sensitive information, defined in terms of its classification level, and the vulnerabilities of the system depend on the degree of control it exerts on its users. The system risk combines the value of the assets, the vulnerabilities of the system, and the clearance of the users.

+-----+




Figure 1. Steps in applying guidance.

Identifying the Risk Factors

To determine a system's security requirements it is necessary to consider the environment in which that system operates. The Orange Book specifies levels of requirements independent of system environment; the draft application doctrine [2] characterizes a system's environment in terms of three parameters: the maximum clearance of the least cleared user, the maximum classification of data processed by the system, and the environment in which the system is developed and maintained (open or closed). While simple to evaluate, these parameters omit important factors that affect actual system risk.

The following paragraphs explain the factors that should be taken into account. For each factor, different levels of risk are defined so that the difference between two adjacent levels in each factor represents a roughly comparable increase (or decrease) in risk. Factors are defined so that they are roughly independent -- a change in one factor does not imply a change in another factor. These properties allow numbering the risk levels and combining them in most cases using simple addition.

Something as abstract as risk cannot be quantified precisely. Recognizing this, we have not attempted to make fine distinctions, and no doubt some systems will still fall near the boundaries of the proposed classes. Nevertheless, the scheme described below, coarse as it is, captures the intuition and experience of computer security practitioners and is preferable

to simply setting these considerations aside because they cannot be made precise.

Local Processing Capability. Some systems have receive-only terminals (e.g., stock transaction displays, airline terminal monitors); users of such terminals have no way to enter system commands directly. Such terminals represent a lower level of risk than typical interactive terminals that permit both sending and receiving information. Replacing a fixed-function interactive terminal with a programmable terminal, personal computer, or other programmable device would introduce a still higher level of risk, since the user can now program his terminal to enter commands for him. A user who accesses a system from a fixed-function terminal but via a programmable host computer would be considered to have the same local processing capability as one who uses a personal computer as a terminal. The identified risk levels for local processing capability are:

Level 1: receive-only terminal

Level 2: fixed-function interactive terminal

Level 3: programmable device (access via personal computer or programmable host)

Communication Path. The communication path between a terminal and host can also affect system risk. A terminal that has a simplex receive-only link to its host via a store-and-forward network (e.g., via radio broadcast) poses less risk than one that is connected via a duplex store-and-forward link, since the simplex path prevents the user from submitting requests to the system. Terminals that are connected to a host either directly, through a local-area network, or long-haul packet network (e.g., Telnet, DDN) offer increased possibilities for penetration and misuse (inadvertant or otherwise) over those connected only through a store-and-forward net because of the increased bandwidth and closer host-terminal interaction they permit. The identified risk levels for communication path are:

Level 1: store/forward, receive-only

Level 2: store/forward, send/receive

Level 3: interactive, via direct connection, local-area net, or long-haul packet net

User Capability. Regardless of the local processing available to a user or the communication path he uses to access a host, if that host is programmed only to provide predefined outputs regardless of the inputs the user presents, it is less risky than a system that responds to user transactions. In this sense, the system that generates the ticker tape for a stock exchange is less at risk to the terminals that display the tape than an interactive electronic banking system is to automated teller machines. Finally, a transaction-based system is less at risk from its users than a system that permits its users full programming capabilities. The identified risk levels for user capability are:

Level 1: output only

Level 2: transaction processing

Level 3: full programming

Development/Maintenance Environment. A system that has been

developed and is maintained by cleared individuals under close configuration control (closed environment) should pose less risk than one that is not developed and maintained in this way (open environment). This distinction has been proposed in the draft application doctrine [2]. It seems a reasonable one, but relatively few examples of systems developed and maintained according to the proposed definition of ``closed environment'' have been identified outside of the intelligence community. For simplicity, we assume that systems are developed and operated in an open environment. Systems that are developed and maintained in a closed environment may therefore be subject to slightly less stringent requirements than will result from our approach.

Data Exposure. A system that has a greater disparity between the clearance of its least cleared user and the classification of the most sensitive data it processes is more at risk than one that has a lesser disparity. The draft application doctrine proposes a scheme for numbering and classifying ``risk range'' we adopt this scheme but call it ``data exposure'' to distinguish it from other risk factors. Although clearance and classification levels used are based on the DoD system, they do include levels for sensitive but unclassified data and users authorized access for such data. For non-DoD environments, it seems likely that analogous clearance/classification levels could be defined. Clearance levels are identified as:

Level 0: uncleared

Level 1: uncleared, but authorized access to sensitive unclassified information

Level 2: confidential clearance

Level 3: secret clearance

Level 4: top secret/background investigation

Level 5: top secret/special background investigation

Level 6: top secret/special background investigation, with authorization for one compartment

Level 7: top secret/special background investigation, with authorization for more than one compartment

Classification levels are numbered:

Level 0: unclassified

Level 1: sensitive unclassified information

Level 2: confidential

Level 3: secret

Level 4: secret with one category

Level 5: top secret with no categories or secret with two or more categories

Level 6: top secret with one category

Level 7: top secret with two or more categories

Data exposure is computed as the difference between the level of the least cleared user of a system and the maximum level of data

processed by the system. It thus ranges from a value of 0 (all users cleared for all data) to 7 (system processes top secret data with two or more categories and some users are uncleared).

Applying the Risk Factors

For a particular system, each of the risk factors needs to be evaluated in order to assess the overall ('system') risk. With minor exceptions, the system risk is simply the sum of the risks of the individual risk factors. Based on system risk and data exposure, security requirements can be determined. These requirements are characterized here in terms of the levels defined in the Orange Book because they have been published and reviewed widely. If a different subsetting of the Orange Book requirements later proves more appropriate than the current set of levels, the new subsets can be substituted. Tables 1-3 provide the necessary mappings between factor values, risk factor levels, and security requirements. The first two tables are only needed because of the exceptions mentioned above; usually, Table 3 can be used directly with the sum of the individual risk factors.

Note that in a given system, different terminals may provide different functions, lead to different levels of risk, and impose different security requirements. Security requirements for the system as a whole must be determined on the basis of the most risky part. As noted previously, the tables below assume all systems are developed/maintained under conditions of an open environment.

Local Processing Capability	Communication Path		
	1. S/F (one- way)	2.S/F (two way)	3.I/A network or direct connection (LAN, DDN)
1. Receive Only Terminal	2	3	4
2. Interactive Terminal (fixed function)	2	4	5
3. Programmable Device (access via personal computer or programmable host)	4	5	6

Table 1. Process Coupling Risk

Table 1. Together, local processing capability and communication path characterize what computer security literature refers to as the 'process coupling' risk. This term is intended to cover how well a process in one computer can maintain its integrity in the face of attempts to subvert it from outside. A high degree of coupling represents a close degree of interaction between two processes, and hence a greater vulnerability of one to the other. If there is a very limited, well-defined set of requests one process can make of the other, then the degree of process coupling will be low. Process coupling risk in a system, as shown in Table 1, is the sum of the local processing capability and communication path risks, with one exception. A fixed function interactive terminal attached to a one-way store-and-forward communication path does not increase risk over a

receive-only terminal on the same link. A programmable device increases risk over the interactive terminal, since, if improperly programmed, it might corrupt labels transmitted with data.

	+-----+					
	Process Coupling					
User Capability	2	3	4	5	6	
1. Output Only (Subscriber)	3	4	5	6	7	
2. Transaction Processing	-	5	6	7	8	
3. Full programming	-	6	7	8	9	

Table 2. System Risk

Table 2. The process coupling value from Table 1, combined with the appropriate user capability factor value yields an overall system risk independent of the data exposure. As in Table 1, the entries of Table 2 have been obtained by summing the risk factor values from each axis. The entries for a process coupling of 2 (receive-only or interactive terminal on a receive-only link) have been omitted for user capabilities of transaction processing and full programming, since a receive-only link cannot support either of these capabilities.

	+-----+							
	System Risk							
Data Exposure	3	4	5	6	7	8	9	
0 (System High)	C1	C1	C1	C1/C2	C2	C2	C2	
1	C1/C2	C2	C2	C2	C2/B1	B1	B1	
2	C2	C2/B1	B1	B1	B1	B1/B2	B2	
3	B1	B1	B1/B2	B2	B2/B3	B3	B3/A1	
4	B2	B2/B3	B3	B3/A1	A1	A1	A1	
5	B3/A1	A1	A1					
6								
7								

Table 3. Mapping System Risk and Data Exposure to Orange Book Requirements Levels

Table 3. This table relates the system risk with the data exposure to yield a level from the Orange Book that defines the security requirements for the system. As noted above, the Orange Book levels may later be replaced by related, but distinct, sets of features and assurances. The entries in this table were generated by working through examples and considering the guidance provided by the draft application doctrine [2] and current DoD

directives governing compartmented mode. Blank entries indicate that, for the specified data exposure level and system risk, it appears technically infeasible to meet the appropriate security requirements at the time.

4. Examples

A Sea Surface Surveillance System (S4)

Consider the application of the technique outlined above to a hypothetical system that keeps track of objects on the surface of the seas. The system collects information from a variety of open and secret sources and distributes it to a variety of customers. The system maintains a data base of sighting information that is both automatically and manually updated. There are two major classes of users: analysts and subscribers.

S4 analysts are the direct operators of the system: they are called on to resolve ambiguities when the system cannot associate a particular sighting with a particular platform, they can cause messages to be sent to subscribers automatically on a regular basis, and they can update the data base. They operate interactive terminals that are located in S4 spaces and connected directly to the S4 computers.

S4 subscribers are the recipients of reports generated by S4. They are located outside the S4 spaces and receive reports over a variety of different communication networks on receive-only terminals. They cannot directly enter data into the S4 system, but they can issue requests (via normal message channels) for regular updates on the location of particular objects, for example. These requests are received by S4 analysts who cause filters to be set up that automatically channel relevant reports to the subscriber. Once the appropriate filter is set up, no further human intervention is required.

Since analysts and subscribers are permitted different kinds of functions, have different clearances, and communicate with the S4 system over different paths, it is necessary to apply this technique to each class of user separately.

Local Processing Capability. Analysts operate fixed function interactive terminals, so they represent a risk level of 2. Subscribers operate receive-only terminals, yielding a risk level of 1.

Communication Path. Analysts communicate with S4 machines directly, so their risk level is 3. Subscribers communicate over a one-way store-and-forward network, so their risk level is 1.

User Capability. Analysts are permitted to issue transactions directly to S4, but they do not have full programming capability, so the risk level is 2. Subscribers have output-only capability, so the risk level is 1.

Data Exposure. S4 processes data at the TS level with multiple compartments, so the classification level is 7. S4 analysts hold TS clearances with SBI and are authorized access for all compartments that S4 processes. Consequently, their clearance level is also 7 and the data exposure for analysts is 0. Some S4 subscribers hold only Secret clearances with no compartment authorizations, so their clearance level is 3, yielding a data exposure for subscribers of 4.

Using the Tables. First, for analysts, Table 1 shows that a

local processing capability risk of 2 and communication path risk of 3 yields a process coupling risk of 5. Table 2 combines a user capability risk of 2 with a process coupling risk of 5 to yield a system risk of 7. Table 3 maps a data exposure of 0 and a system risk of 7 to a C2 level system requirement.

For subscribers, Table 1 combines a local processing capability risk of 1 with a communication path risk of 1 to yield a process coupling risk of 2. Table 2 combines a user capability risk of 1 with a process coupling risk of 2 to give a system risk of 3. Finally, Table 3 maps a data exposure of 4 and a system risk of 3 to a B2 level system requirement.

Since S4 includes both kinds of users, the more stringent of the two requirements (i.e., B2) would apply. Changes to the environments of either subscribers or analysts (such as the introduction of personal computers in place of fixed function terminals) would require the risk evaluation to be repeated, and could lead to a change in the level of security requirement.

Evolution of the S4 System

Suppose that after initial deployment of S4, its subscribers clamor for terminals more up-to-date than the original receive-only ones. The system sponsor proposes to replace them with personal computers. What are the effects on the security that the host system needs to provide? The local processing capability risk factor changes from 1 to 3, and the system risk becomes a 5; the data exposure for subscribers is unchanged. Table 3 shows that the host security should be upgraded from B2 to B3. If, in addition to the personal computers the sponsor permits subscribers to communicate with the system over a real-time network and to initiate transactions, the system risk becomes 8, and an A1 host would be indicated. By estimating the additional cost of replacing or upgrading the S4 host to the B3 or A1 level, the sponsor can quantify the cost of providing new functions while maintaining an acceptable level of risk for the system.

``Orange Book Environment''

The Orange Book does not explicitly define an environment. However, the predecessors of the Orange Book criteria were first developed in the context of an interactive computer system that provided users with directly connected, fixed-function terminals and full programming capability. The corresponding entries in Tables 1 and 2 yield a system risk of 8. Since no data exposure is defined for the Orange Book environment, it is appropriate to consider the result for the Air Force Data Services Center (AFDSC) Multics environment, which provides full programming to users at fixed function, directly connected terminals. AFDSC Multics includes non-compartmented data classified up to top secret and some users have only secret clearances, so the data exposure is 2, and the resulting security requirement from Table 3 is for a B1/B2 system. Multics is currently under evaluation by the DoD Computer Security Evaluation Center and is expected to achieve a B2 rating.

5. Discussion

Here we address some possible objections to the approach described above.

Objection: the proposed scheme imposes different requirements on a host computer based on characteristics of the user's terminal and the communication path between the terminal and the

host. These are outside the security perimeter of the host and therefore should not affect the security required of it.

Response: security considerations include not only protecting data up to the point that it leaves the system but also resisting attacks on the system mounted by external users. Users with personal computers and direct connections to systems have proven a greater threat (e.g. in terms of their ability to defeat password schemes) than those who have only fixed-function terminals at their disposal. Each higher Orange Book level adds assurance requirements as well as security feature requirements. While the security features added at a particular level may or may not improve protection against threats posed by terminals and networks connected to a host, the increased assurance provided by each incremental level should decrease the likelihood of flaws that could be exploited from outside the security perimeter. It is thus appropriate to increase the Orange Book level required of a host based on the risk factors assigned to the user capability and communication path.

Objection: the proposed approach in some cases permits hosts to meet lower security requirements than would the draft application doctrine[2].

Response: the approach proposed here distinguishes aspects of application system structure that reduce its vulnerability to outside attacks. The draft application doctrine determines the level of system required based primarily on the clearances of system users and the classification of data stored in the system. There is no distinction, for example, between a system in which users can only view output and one in which users can construct and execute their own programs. Consequently, the proposed requirements must be based on the worst case assumption (user programming). By providing a more detailed model of the environment, the approach proposed here permits a more accurate assessment of the security actually required.

Objection: previous attempts to distinguish rigorously between a system that can be programmed and one to which only transactions can be submitted have failed.

Response: while a formal mathematical distinction between systems that users can program and those that perform a fixed set of functions in response to user requests may never be defined, it does not seem to be a difficult distinction to make in practice. In cases that are difficult to decide (e.g., a ``transaction-processing'' database system that permits a complex query and update capability) it is safe to assign the system the higher risk factor.

Objection: because the proposed approach determines host security requirements partly based on system architecture, changes to the architecture may lead to different security requirements.

Response: this is actually a benefit of the approach. As a system changes during its design, development, and operation, the effects of those changes on host security requirements can be easily assessed, providing a practical way to use the Orange Book requirements throughout the system life cycle. If, for example, a B2 host will not be available to support an application as originally planned and a B1 host must be used instead, the approach proposed here can help determine how system functions, user capabilities, or communication paths could be restricted to compensate for the less secure host. Conversely, if new functions or terminals are added to a system already under development, this

approach can indicate whether host security will need to be upgraded as a result. The only tradeoff that would be recognized under the draft application doctrine would be to limit the classification of the data to be processed by the system or increase the clearance of its users.

6. Conclusion

We have presented a scheme for determining an appropriate set of host security requirements using the requirements and levels identified in the Orange Book. The scheme takes into account the functions available to a user locally, the communication path used to gain access to the host, and the functions the host provides, as well as the user's clearance and the classification of data processed by the host. By including these system characteristics, this technique makes it possible to assess trade-offs among system function, system architecture, and system costs while maintaining an acceptable level of system risk.

References

1. Department of Defense Trusted Computer System Evaluation Criteria, DoD Computer Security Evaluation Center, CSC-STD-001-83, 15 August 1983.
2. Brand, S. Environmental Guidelines for Using the DoD Trusted Computer System Evaluation Criteria. Proc. Seventh DoD/NBS Computer Security Initiative Conference, Sept., 1984, Gaithersburg, MD, pp. 17-23.

Local Processing Capability	Communication Path		
	1. S/F (one- way)	2. S/F (two way)	3. I/A network or direct connection (LAN, DDN)
1. Receive Only Terminal	2	3	4
2. Interactive Terminal (fixed function)	2	4	5
3. Programmable Device (access via personal computer or programmable host)	4	5	6

Table 1. Process Coupling Risk

User Capability	Process Coupling				
	2	3	4	5	6
1. Output Only (Subscriber)	3	4	5	6	7

2. Transaction Processing	-	5	6	7	8
3. Full programming	-	6	7	8	9

Table 2. System Risk

	System Risk							
Data Exposure	3	4	5	6	7	8	9	
0 (System High)	C1	C1	C1	C1/C2	C2	C2	C2	
1	C1/C2	C2	C2	C2	C2/B1	B1	B1	
2	C2	C2/B1	B1	B1	B1	B1/B2	B2	
3	B1	B1	B1/B2	B2	B2/B3	B3	B3/A1	
4	B2	B2/B3	B3	B3/A1	A1	A1	A1	
5	B3/A1	A1	A1					
6								
7								

Table 3. Mapping System Risk and Data Exposure to Orange Book Requirements Levels