

AD-A145 927

MULTILEVEL SECURITY IN A LOCAL AREA NETWORK(U) NAVAL
POSTGRADUATE SCHOOL MONTEREY CA D A STRAUB MAR 84

1/1


UNCLASSIFIED

F/G 9/2

NL

END

D7HC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD-A145 927

2

NAVAL POSTGRADUATE SCHOOL

Monterey, California



DTIC
ELECTE
SEP 26 1984
S B D

THESIS

MULTILEVEL SECURITY
IN A
LOCAL AREA NETWORK
by

Debra Ann Straub

March 1984

Thesis Advisor:

Norman R. Lyons

Approved for public release; distribution unlimited

DTIC FILE COPY

84 · 09 17 071

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Multilevel Security in a Local Area Network		5. TYPE OF REPORT & PERIOD COVERED Master's Thesis March 1984
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Debra Ann Straub		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Postgraduate School Monterey, California 93943		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Postgraduate School Monterey, California 93943		12. REPORT DATE March 1984
		13. NUMBER OF PAGES 80
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) WWMCCS Information System Multilevel Security Trusted Software Local Area Network		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This thesis examines the design of a local area network that is able to simultaneously handle users at a variety of security levels, while providing full multilevel protection of the data. A major feature of the design is the use of trusted software in the network interfaces to provide security for data entering or leaving the network. This secure design was initiated to investigate options for local area network technology that could		

be incorporated into the planned near-term upgrade for the WWMCCS Information System ADP support. <

Distribution Unlimited is correct for this report per Professor Norman Lyons, thesis advisor

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

Approved for public release; distribution unlimited.

Multilevel Security in a Local Area Network

by

Debra Ann Straub
Lieutenant, United States Navy
B.S., Indiana State University, 1975

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN TELECOMMUNICATIONS SYSTEMS MANAGEMENT

from the

**NAVAL POSTGRADUATE SCHOOL
March 1984**

Author:

Debra A. Straub

Approved by:

Norman R. Lyons

Thesis Advisor

Daniel R. Dolk

Second Reader

Richard L. Elster

Chairman, Department of Administrative Science

Kenneth T. Marshall

Dean of Information and Policy Sciences

ABSTRACT

This thesis examines the design of a local area network that is able to simultaneously handle users at a variety of security levels, while providing full multilevel protection of the data. A major feature of the design is the use of trusted software in the network interfaces to provide security for data entering or leaving the network. This secure design was initiated to investigate options for local area network technology that could be incorporated into the planned near-term upgrade for the WWMCCS Information System ADP support.

TABLE OF CONTENTS

I.	INTRODUCTION	8
II.	BACKGROUND	9
III.	ARCHITECTURAL COMPONENTS	11
	A. SUBNETWORKS	11
	B. TRUSTED INTERFACE UNITS (TIUS)	11
	C. BRIDGES	12
	D. GATEWAYS	13
	E. GUARDS	13
IV.	CONCEPT	17
	A. SIMPLE MULTILEVEL LAN	18
	B. FULL MULTILEVEL LAN	20
V.	OPERATIONAL ENVIRONMENT	25
	A. SINGLE-LEVEL RESOURCE SCENARIO	26
	B. VARIABLE-LEVEL RESOURCE SCENARIO	28
	C. MULTILEVEL RESOURCE SCENARIO	30
VI.	DESIGN DETAILS	32
	A. LEVEL OF INTERCONNECTION	35
	B. ADDRESSING, SECURITY AND ROUTING	35
	C. PROTOCOLS	36
	1. Low Layer Protocols	38
	2. High Layer Protocols	42
	D. TRUSTED INTERFACE UNIT (TIU)	46
	1. Single-Level TIU	47
	2. Variable-Level TIU	54
	3. Multilevel TIU	55
	E. BRIDGES	56

1. Security Processing	59
2. Routing Concepts	61
3. Buffering	63
4. Half-Bridges	64
F. FLOW AND CONGESTION CONTROL	65
VII. SUMMARY	67
A. ADVANTAGES	68
E. DISADVANTAGES	69
C. FURTHER RECOMMENDATIONS	76
D. CONCLUSION	77
LIST OF REFERENCES	78
INITIAL DISTRIBUTION LIST	80

LIST OF FIGURES

3.1	Guard on the Secure LAN	16
4.1	Simple Multilevel LAN	19
4.2	Full Multilevel LAN	22
5.1	Single-level Resource Scenario	27
5.2	Variable-level Resource Scenario	29
5.3	Multilevel Resource Scenario	31
6.1	Local Area Network Packet Format	39
6.2	Trusted Interface Unit (TIU) Architecture	48
6.3	Bridge Architecture (Half-Duplex)	58
6.4	Fixed Routing Tables	62

I. INTRODUCTION

The Department of Defense (DoD) is currently upgrading the Worldwide Military Command and Control System (WWMCCS) in an effort to evolve the existing Automated Data Processing (ADP) capabilities into a new WWMCCS Information System (WIS).

Plans for this modernization call for a communications medium that will allow terminals to communicate with both single-level and multilevel secure functional resources that would be required at a WIS site. The WWMCCS System Engineer has proposed that the medium be a local area network (LAN) [Ref. 1]. Since the WIS is to be based around a LAN, and resources of different levels will require LAN access, the LAN must be multilevel secure from the beginning. However, a "controlled mode" LAN may be acceptable initially.

Sidhu and Gasser [Ref. 2] have designed a secure LAN that is based upon the "trusted system" concept. This thesis exams their design option for a secure LAN that can be incorporated into an initial scenario to provide multi-level security for WIS sites.

II. BACKGROUND

In Sidhu and Gassers' design it is assumed that the communication medium is as well protected as the users' work stations, safes, etc. Therefore, encryption of the data is not a requirement except where the medium must pass through an unprotected zone, such as between buildings. Their design goal is to enforce the DoD security policy [Ref. 3] with respect to accessing data on the communications medium, while enabling a wide variety of resources to be shared.

The constraints of a typical WWMCCS command center environment motivated the design. At a typical WWMCCS node data processing must be carried out at multiple security levels, with users requiring access to multiple levels of data. The cost of certifying and accrediting physical facilities for high levels of security dictates that not all facilities may be cleared to hold classified data. Therefore, an example of their design goal is to make Secret data from one network (which handles data up to the Secret level) available to a Secret subscriber on another network which handles data up to the Top Secret level (note that this design does not take into account a control for need-to-know). Also, the design must prevent Top Secret data from spilling (from a Top Secret controlled node network) into a Secret controlled node network.

The key to Sidhu and Gassers' trusted systems approach is the use of trusted interface units (TIUs) that govern secure communication between subscribers only at identical security levels. Outgoing data are marked by the TIU with the security level of its attached subscriber. All data coming into a subscriber is examined by the TIU to ensure that its security level matches that of the attached subscriber.

Another aspect of the design calls for bridges that allow data to be shared among physically separate LANs. For example, the bridges would govern the sharing of data between a Secret controlled mode network and a Top Secret controlled mode network, as previously described. The bridge is configured as a network subscriber whose task is to relay data to another network, which also acts as a network subscriber. While relaying data, the bridge ensures that no data shall be placed on a network that is not cleared to handle the security level of the data.

The WIS modernization program dictates that the secure LAN be available around 1985. Another design goal is to implement an initial secure LAN capability within this time frame. For this reason, Sidhu and Gasser stress a near-term solution in great detail, with progressively less detail provided for the more complex longer term solutions.

The remainder of this thesis discusses the architectural components, concept, operational environment and design details of this particular approach to multilevel security in a local area network.

III. ARCHITECTURAL COMPONENTS

The LAN configuration contains the following: subnetworks and their communications media, trusted interface units, bridges, gateways and guards. An overview of these components follows.

A. SUBNETWORKS

A subnetwork is a part of the LAN that fully resides within a protected environment (an area physically protected to a specified system-high security level that corresponds to a portion of a building, whole building or group of buildings). Protected environments are further defined by the proximity of a set of subscribers which operate up to a given security level. If parts of a subnetwork passes through an unprotected area, then it is assumed that encryption devices will be used that can handle the bandwidth of the subnetwork, or a separate subnetwork could be constructed with an intervening bridge. It is assumed that the protected environment will include any encrypted portions of the medium.

B. TRUSTED INTERFACE UNITS (TIUS)

To access a network a subscriber connects to a TIU that places the proper security markings on all data entering the LAN from the subscriber, and in return provides the subscriber with only that data from the LAN that has the proper security labels. It should be noted that subscribers attached to a subnetwork may operate at any security level at or below that of the subnetwork environment.

Initially, a single-level subnetwork has been proposed for the vast majority of users that will require only a single-level environment and for the evolutionary growth of the LAN [Ref. 4: p. 10]. Although this particular single-level (for example, Secret) TIU supports a single-level Secret subscriber, the subnetwork interface within the TIU is actually multilevel since it has access to all data on the subnetwork. Therefore, the single-level TIUs are considered trusted, while the single-level subscriber need not be.

Variable-level TIUs support subscribers who may operate at different security levels at different times. This particular TIU is the same as single-level TIUs, except that their security level may change over time.

For multilevel subscribers, multilevel TIUs are required. The difference between a multilevel TIU and a variable-level or single-level TIU (in terms of the trusted mechanism in the TIU) are discussed further in chapter VI.

C. BRIDGES

The subnetworks comprising the LAN are themselves local area networks that are connected to each other with bridges [Ref. 5: p. 1497-1517]. The bridges therefore provide the links between subnetworks at the same or different security levels. Bridges must be trusted to prevent packets on a higher level subnetwork from entering a lower security level subnetwork, if those packets originated with a subscriber operating above the lower level. This precaution is due to the fact that one subnetwork may be protected to a lower security level than another subnetwork to which it is bridged.

D. GATEWAYS

One of the ultimate goals of the WIS upgrade is to provide WWMCCS Intercomputer Network (WIN) access to all LAN subscribers at a site [Ref. 1], therefore necessitating a gateway connected directly to the LAN. To be fully useful, the gateway to any multilevel network (i.e., WIN, Defense Data Network (DDN), or other wide-area networks) must itself be multilevel secure, including the gateway's interface to the LAN. It should be noted that Sidhu and Gasser focus only on the LAN problem and do not address gateways to any detailed extent.

E. GUARDS

Also employed in this design are guard nodes that allow information to move from one security level to another in a controlled manner. Generally, a guard is a trusted computer that examines output from a single-level computer running at a "high" security level and transfers that data to another computer or subscriber running at a "low" level [Ref. 6]. The purpose of the guard is to allow data to flow from high to low if that data is actually classified at the low level but happens to reside in the high level environment. A guard may also be used to selectively downgrade output from a high level system. In order to prevent accidental downgrading, guards may examine data flowing from low to high for "acceptable" values. An example of the use of a guard is where there is a need for a Secret user to access Secret data residing in a Top Secret subnetwork. In this case, guards are required because existing mainframes are not trusted for multilevel or controlled mode operation. Existing mainframes must run at the level of the most sensitive data (single-level system-high), even though the system-high information is only a small portion of the total

data. Therefore the guard provides an inexpensive alternative to clearing all users to the system-high level who must access such data.

A guard usually operates with one high side and one low side, each at specific fixed levels. The guard can then be easily interfaced to the secure LAN by the use of two single-level TIUs, one for each side. Figure 3.1 illustrates one particular guard system that could be used on a secure LAN [Ref. 7: p. 5]. Data flowing from a high subnetwork would enter the guard via the high TIU and send the downgraded results via the low TIU to the appropriate subscriber. This would imply that the guard must reside on a subnetwork with a level at least that of the high subnetwork. The guard system in figure 3.1 is logically partitioned into guard trusted and untrusted functions. The guard trusted functions provide the capability to ensure the secure interchange of information across security boundaries. The guard low and high untrusted functions support the guard trusted information flow functions, provide application system specific functions that are not security relevant, and act as an interface between the guard trusted processing element and its external application system environment.

The guard system consists of three (3) computer programs: Guard Trusted Processing (GTP) element, Guard Low Processing (GLP) element, and the Guard High Processing (GHP) element. The GTP element provides all guard trusted functions. The GLP and GHP elements provide all guard low and high untrusted functions, respectively.

The GLP and GHP elements are functionally symmetric in that each acts as an interface between the GTP and its corresponding external application system environment (low or high). The GLP and GHP each consist of three logical elements: an application system interface that interprets

application system protocol; a guard protocol interface which provides a secure communications path between guard untrusted and trusted elements in support of the guard information flow functions; and guard untrusted application-specific functions which are not security relevant and, therefore need not be trusted. Since single-level TIU interfaces are used in figure 3.1, the guard would be treated by the LAN as two separate single-level subscribers. It should be noted that another option would be to use a single multilevel TIU for the guard.

Sidhu and Gasser have assumed that existing and planned guards, such as the one just described, will function on the secure LAN with little or no modification, other than that normally required to interface any host to a LAN.




Figure 3.1 Guard on the Secure LAN.

IV. CONCEPT

There are two major areas of DoD security policy which must be addressed when considering a local area network security architecture. The first is "user separation". This concept refers to the ability of a network security design to provide segregated communities of subscribers such that traffic from individual communities can be transmitted to authorized subscribers in those communities and not disclosed to subscribers outside those communities. That is, some network component must ensure that an Unclassified subscriber receives only Unclassified information over the local network. The enforcement of this policy provides protection against unauthorized disclosure to authorized but uncleared subscribers on the network.

The second policy area is the concept of "data protection". This refers to the ability of the network to provide protection against malicious attempts to access, modify or destroy information in the network. This includes subverting authorized network components and attaching unauthorized network components. In order to enforce this policy two approaches may be used--encryption (end-to-end) or Protected Wireline Distribution System (PWDS) [Ref. 8]. A PWDS provides physical protection of network components including the transmission medium, whereas encryption provides protection by making the information unintelligible.

Sidhu and Gasser describe their overall solution concept in two stages-- from the simplest LAN to the most complex LAN. In their concept it can be seen that "user separation" and "data protection" are incorporated by means of subnetworks and encryption of the LAN medium itself.

A. SIMPLE MULTILEVEL LAN

A simple multilevel LAN would appear as in figure 4.1 [Ref. 2: p. 5], with a simple LAN communications medium supporting a variety of resources or subscribers (i.e., hosts, terminals, and such devices as printers and mass storage units). Subscribers may operate at various security levels, but communication is restricted between subscribers to those of the same level. If there are subscribers of more than one security level, then only designated multilevel hosts or terminals (trusted to protect information at several classification levels) would be able to communicate with them.

Each single-level resource will be protected and controlled to process information up to a designated maximum security level for a specific resource, just as it is currently done today. All information from a single-level subscriber that enters the LAN, will be protected as if it were at that level and will not be transmitted to a resource at a lower security level. This restriction is the same as that imposed today, except the components operating at the same system-high security levels may be interconnected. The difference for the LAN is that the separation between the security levels is logical rather than physical. Multilevel resources that are assigned a set of security levels at which they are allowed to operate are restricted to communicating with other multilevel or single-level resources within their range.

The LAN itself is designed to be a passive medium such as coaxial cable or a twisted pair. The key architectural elements of the LAN are the trusted interface units (TIUs). The TIUs will maintain separation of information of different levels, with one TIU associated with each subscriber or set of subscribers operating at a given




Figure 4.1 Simple Multilevel LAN.

security level or range. The TIUs perform a single security-related function, that is, to label each packet of outgoing (subscriber to the LAN) data with the correct security level of the subscriber and to check the security labels of incoming packets and compare them against that of the subscriber. For the case of a single-level subscriber, only one level is allowed and therefore all security controls (labelling and checking) are built into the TIUs and the subscribers do not need to be "trusted" to provide the proper security labels. However, multilevel subscribers are trusted to specify their own security labels within a specified range. They must also protect incoming data within that range (the TIUs for those subscribers will enforce the correct range).

Since there is no communication across the security levels in figure 4.1, the logical effect of the LAN is that of separate LANs. This architecture provides no benefit over physically separate LANs that do not require the use of interface units. There is one main advantage with this structure and that is the upgrading capability that it possesses. If the architecture shown in figure 4.1 is used as a basis, then multilevel operation can be achieved via a series of small incremental enhancements as development progresses. Each upgrade can be accomplished without discarding any existing hardware or software or without impacting the operation of the system.

The obvious problem with the architecture of figure 4.1 concerns physical protection requirements. It is apparent that the entire LAN medium and all the TIUs must be protected to system-high since all components will contain system-high information. This implies that all hosts and terminals must be protected to system-high (the situation today) which is not practical. Keeping in mind that the subscriber-TIU interfaces usually consist of relatively short cables, the problem that arises is how to connect an unprotected subscriber to its TIU where the TIU and LAN medium must reside in a highly protected area.

The solution concept presented so far ignores very serious physical constraints on how users and hardware are protected in a secure environment. Therefore, a more flexible solution is required.

B. FULL MULTILEVEL LAN

Figure 4.2 [Ref. 2: p. 9] illustrates a full multilevel LAN. The concept is to provide a separate physical sub-LAN (subnetwork) for each community with different security protection requirements. Each subnetwork is itself a simple

multilevel IAN as depicted in figure 4.1. This is accomplished through the use of bridges constructed as interfaces between subnetworks. The bridges act as filters for classified packets of data addressed across subnetwork boundaries.

Although each subnetwork has a maximum security level associated with it, subnetworks provide full multilevel protection for all levels below their maximum. Individual users may operate at levels below the subnetwork maximum, while maintaining the same restriction that only subscribers of the same levels can communicate. It is not the intention of this design to provide a subnetwork for every combination of specific security level and compartment utilized at a WWMCCS site. The intent is just for one subnetwork for each environment that is separately physically protected to a given system-high level.

To clarify this concept, consider host A in figure 4.2. It is physically protected to the Secret level and it may run at the Unclassified level, wherein it can communicate with host B on the Unclassified subnetwork. In order for host A to run Unclassified it must be appropriately sanitized and must have no connections to any secret devices other than the TIU to the subnetwork. Host C illustrates a subscriber in an Unclassified environment connected to a Top Secret subnetwork. Because the TIU for host C must be protected to Top Secret in order to remain trusted, it is the only component for host C that needs to be within the Top Secret environment. Host D in the Confidential environment is configured as a remote from the Secret subnetwork to which it connects. Here again, the TIU for host D provides the isolation of the host from data on the subnetwork above the Confidential level.

Figure 4.2 only illustrates a subnetwork's maximum level, however the concept of a subnetwork's minimum level may also apply. A minimum level is enforced in the same




Figure 4.2 Full Multilevel LAN.

manner as a maximum subnetwork level. A minimum level is utilized to limit the damage that could be done by misclassifying and mis-routing a packet of data.

For example, if it is known that a certain subnetwork has a minimum of Secret, then interfaces to the subnetwork (bridges and TIUs) would assume that all data are at least Secret and will not, or will automatically upgrade any data labelled at a lower level. A minimum level does not limit the actual minimum level of data that users can transmit. It only limits the labels that can be placed on the data by the network components, and therefore it limits the destination of the data. Actually, the minimum level of a subnetwork would be greater than Unclassified only if the various devices protecting the data were not trusted (administratively) to provide the required protection.

For example, if part of a LAN contained highly sensitive data which the TIUs and bridges were unable to adequately protect, then that section of the LAN subnetwork could be isolated and its use limited to subscribers operating at the appropriate security levels. Users would be able to operate at all levels on the subnetwork, and be able to communicate via a bridge to other parts of the LAN. However, no data leaving the subnetwork could be labelled at a level below the minimum level of the subnetwork. If a TIU on the subnetwork were compromised, the TIUs and bridges that comprise the rest of the LAN will prevent any data originating from the subnetwork from travelling to a subscriber, elsewhere on the LAN, that is not cleared to the appropriate level, or from arriving on another subnetwork of a lower level. One point should be explained concerning this last statement. If the TIUs on the subnetwork are not trusted to avoid downgrading the information, why should the bridges or TIUs on the rest of the LAN be trusted not to do the same? They shouldn't. The lack of trust of the devices on the subnetwork refers to the type of environment the subnetwork is located in and the possibility of a lack of complete trust in the multilevel resources on that subnetwork. An

example of this reasoning is a subnetwork running in a controlled mode where data are labelled by partially trusted multilevel hosts up through Top Secret, but where the users of the subnetwork have a minimum clearance of Secret. The multilevel host is trusted only to distinguish between Secret and Top Secret.

Note that in figure 4.2 encryption plays an important role. Between the Secret and Confidential subnetworks and the Secret and Top Secret subnetworks, the encryption involves a simple serial bit stream that is link encrypted.

Between the two Top Secret subnetworks, the LAN itself is designed to be encrypted. So far, the complexities of encrypting a LAN medium have not been studied in detail. But this should not have a great impact on Sidhu and Gassers' design, since it does not depend on the ability to encrypt the LAN medium directly (the two portions could be physically separate subnetworks connected by bridges). However, encryption of the medium is attractive because it could minimize the number of separate subnetworks that would be employed at a given site.

V. OPERATIONAL ENVIRONMENT

The goals for the fully operational configuration of the secure LAN, according to Sidhu and Gasser, are that it be able to maintain separation between classified data and users that are not cleared to see that data, and that it give appropriately cleared users access to data which may have different classifications. This means that the information must be maintained by classification in the computer(s) and that the information be controlled by classification from the computer to the user. Therefore, the mature configuration of the secure LAN must be able to support multilevel computers, multilevel networks, and multilevel terminals.

Full multilevel operation of resources at a WIS site will not be achieved in the immediate future, no matter what the multilevel capabilities of the LAN. Therefore, to limit the risk of constructing a LAN to support full multilevel resources from the beginning, most resources in the initial installations will be single-level. In chapter IV the incremental enhancement of a secure LAN from the simple version of figure 4.1 to the multiple version of figure 4.2 was introduced. This progressive enhancement was designed to compensate for physical security installation constraints at a WIS site. From here on the LAN will be addressed as if it consisted of multiple subnetworks, with the understanding that the single-subnetwork version will be an initial capability fully compatible with the final structure.

Sidhu and Gasser have provided a series of "scenarios" that provide successively more flexible and improved multilevel secure processing capability. The three scenarios, which are centered around the three versions of the TIU

discussed in chapter III, enable capability to start with the existing single-level equipment and grow to a fully mature multilevel secure configuration. The first scenario documented in this thesis is expected to be operational around 1985 or 1986. The functions in the second and third scenarios are to be achieved via evolutionary growth. All of the scenarios are illustrated with multiple subnetworks, yet any of these scenarios can operate with either the single or multiple subnetwork versions of the secure LAN.

A. SINGLE-LEVEL RESOURCE SCENARIO

In this scenario it is assumed that all of the resources are untrusted and therefore single-level. Single-level TIUS are designed and trusted to ensure that data to and from a particular resource always have the proper label of the level of that resource. In this configuration only bridges will transmit data of more than one security level.

Figure 5.1 [Ref. 2: p. 23] depicts the single-level resource scenario: a LAN consisting of Top Secret and Secret subnetworks, supporting single-level hosts and single-level terminals. The security level of each subnetwork is the maximum level of any of the single-level subscribers on the subnetwork, and therefore the maximum level that packets will be labelled on the subnetwork. In figure 5.1, the Confidential terminal (third from the left) is able to access the Confidential host computer (lower right), the Secret terminals can access the Secret host computer, and the Top Secret terminals can access the Top Secret computer. The bridge joining the Top Secret subnetwork and the Secret subnetwork assures that only packets of a level of Secret or below are present on the Secret subnetwork. The bridge also allows all packets to pass from the Secret subnetwork (destined for the Top Secret subnetwork)




Figure 5.1 Single-level Resource Scenario.

ensuring that their labels are no higher than Secret, but it does not allow Top Secret packets to pass from the Top Secret subnetwork to the Secret subnetwork. The Top Secret subnetwork may also contain Secret or Confidential packets in addition to Top Secret packets that originated from the Secret host or Confidential terminal. Note that Secret packets may contain data below the level of Secret but cannot be so marked if the single-level resource from which it originated cannot be trusted to maintain the separation and labelling. In the remainder of this thesis there may be reference to the "level" of a packet as being equivalent to the value of some security label on the packet which is not necessarily the same as the level the user ascribes to the data.

In figure 5.1 note the connection of Confidential and Secret terminals and the Secret host to the Top Secret subnetwork. If the Secret and Confidential terminals and Secret host are connected to their TIUs via a short cable, then there may be a problem concerning different physical protection requirements where the subscriber, TIU and subnetwork are concerned (since the TIUs connected to the Top Secret subnetwork must all be protected to the Top Secret level).

B. VARIABLE-LEVEL RESOURCE SCENARIO

This scenario, which is depicted in figure 5.2 [Ref. 2: p. 25], permits the sharing of data by users at more than one security level. This is due to the introduction of variable-level trusted interface units that allow a terminal user to talk to single-level resources at different levels up to and including his terminals' classification. The interface units are trusted to the extent that they restrict the user to one level at a time since the interfacing terminal would not be trusted to simultaneously handle multilevel data. However, through the use of switching mechanisms, the level of the terminal, and therefore of the TIU, may change under user control in a static manner. For example, figure 5.2 shows two variable-level TIUs on the Top Secret network. A subscriber at the "C-TS" terminal could communicate with either the Confidential, Secret or Top Secret hosts on either subnetwork through his variable-level TIU.

A variable-level TIU can be constructed that would work with a multiple position switch to connect a user's terminal to one security level or another. However, there could be a problem in that a user would lose any context that he previously had. Findings concerning the LSI Guard system




Figure 5.2 Variable-level Resource Scenario.

[Ref. 7: p. 4-16], which allows a two-level switching capability, are such that there will be a loss of context that will prove to be extremely annoying to an operational user. The variable-level TIU does not solve this problem but does provide a more enhanced capability that would not be achieved with a single-level TIU.

The variable-level TIU permits a highly classified user to share lower classified resources with users who are operating out of a lower classified area without requiring a separate terminal for the lower classified resources. An example of this would be intelligence users accessing unclassified or Secret data maintained by logistics personnel. The shared database would run at the Secret

level, and could be physically located on either the Top Secret or Secret subnetwork. Cross-level transfer of data, however, would be prohibited.

A variable-level TIU could also be used to interface a LAN to a host that operates at different levels at different times (through periods processing).

C. MULTILEVEL RESOURCE SCENARIO

Figure 5.3 [Ref. 2: p. 26] illustrates the fully operational capability of the proposed LAN. A user with a terminal capable of maintaining the separation of data would be connected to the LAN via a multilevel TIU in order to view and modify data of different levels simultaneously in connection with terminals that have screens or windows for each security level. The terminal and the multilevel TIU coordinate the security level of each data transfer. Also, the multilevel TIUs would allow multilevel hosts to "simultaneously" communicate with various single-level or multilevel terminals and hosts.




Figure 5.3 Multilevel Resource Scenario.

VI. DESIGN DETAILS

This chapter provides the details of the design of a multilevel secure local area computer network for WIS. The design was approached with the basic requirement that the LAN allow transmission of data at different security levels but with appropriate protection of data at each security level. This requirement was further integrated with other requirements such as near-term low-risk feasibility of implementation and incremental upgrade capabilities.

The full LAN architecture depicted in figure 4.2 consists of several physically separated subnetworks. Sidhu and Gasser have attempted to maintain the logical view as that of a single local area network with the underlying substructure being totally transparent to the users. The subnetworks comprising the LAN are themselves local area networks connected to each other with bridges. The bridge, in this architecture, is somewhat similar to a gateway in the interconnection of wide-area packet-switched networks, but it is expected to be much simpler. Even though a bridge is simpler than a gateway, it is more complex than a repeater that might be used to boost signals in an Ethernet [Ref. 9]. A bridge accepts packets from one network and broadcasts it onto the medium of one of the other local area networks to which it is connected. The LAN design under review has been simplified such that the bridge connects just two local area networks to each other. A possible upgrade option would be the connection of several subnetworks to a single bridge. A brief summary follows before proceeding with the detailed design work.

Each subnetwork has a security level associated with it which is the level of the protected environment in which it physically resides. Each subnetwork will only carry information with a security level equal to, or less than, the security level of the subnetwork. It may also have a minimum level that is the minimum level at which data in the subnetwork will be labelled. The bridges and trusted interface units will enforce the minima and maxima. Data passing through an unprotected environment must be encrypted to provide protection against passive wiretapping.

The local area network shown in figure 4.2 has the following additional benefits besides meeting the basic security related requirements:

- a. It will allow reconfiguration with minimal disruption of service within a fixed security environment.
- b. It will allow user separation by communities of interest and information flow, as well as by security levels.
- c. It will allow for data security by physical separation of data flow.
- d. By limiting the effect of failures and denial-of-service attacks to a single subnetwork, it will enhance reliability.

To begin a look at the detailed design work, a specific set of LAN protocols must be examined. Most of the protocol issues in this design center around the access methods used by TIUs to interface to the LAN medium. For this access protocol, the authors chose a basic contention-type protocol similar to the "carrier-sense multiple access with collision detection" (CSMA/CD) protocols that are currently being used in local area networks such as Ethernet [Ref. 10: p. 395-404]. They have further incorporated the relevant features of the CSMA/CD access method of the proposed IEEE 802 standard for local area networks [Ref. 11].

Sidhu and Gassers' reason for choosing CSMA/CD is to show how the secure LAN architecture could be implemented using at least one well-specified and currently implemented protocol (even though the IEEE version of CSMA/CD is now only a proposal, they felt the Ethernet protocol is sufficiently similar to be considered a representative implementation). The architecture in figure 4.2 does not depend on CSMA/CD or any other specific protocol. To go beyond the superficial level of detail of figure 4.2, however, a specific protocol must be chosen around which to base further design. The authors do not preclude the use of other protocols to build a secure LAN, but the use of a protocol substantially different from CSMA/CD could possibly require significant changes to much of the design detail to follow.

It should be noted and stressed that, while slightly different versions of CSMA/CD could be used in the various subnetworks with minimal impact, the concept is not suitable for interconnection of CSMA/CD subnetworks with subnetworks using non-contention type protocols such as switched line, token passing rings, etc. If subnetworks based on non-contention technology were included it would probably add considerable complexity to the protocol architecture, particularly in the bridges which must deal with protocol conversion. However, a different design could be used based entirely on non-contention protocols.

The authors point out three features of their architecture that are not yet commonly employed in existing commercial LANs:

- (1) The use of bridges to connect subnetworks (a few commercial offerings have recently emerged in this area).
- (2) Trusted hardware and software in an interface unit and bridge.

- (3) The labelling of packets according to a security level.

For the single-level resource scenario (1985-86 time frame), the first two features will be shown to be rather straightforward to implement with only a small change to existing components while the last feature is moderately more complex. Subsequent evolution will require further development in all three areas.

The remainder of this chapter will discuss briefly certain considerations such as the level of interconnection of subnetworks, addressing, security and routing. Details of the design for the LAN protocols, interface units, bridges, and flow and congestion control will then be presented.

A. LEVEL OF INTERCONNECTION

One important issue in designing the bridges (used to interconnect subnetworks) is the protocol layer at which subnetworks are to be connected [Ref. 12: p. 1386-1407] and [Ref. 13: pp. 175-195]. A bridge can play the role as an interface unit or as a host. Since the authors are assuming a common LAN technology (suitable broadcast medium) with identical protocols implemented in all the subnetworks, the most natural choice of network interconnection is at the interface unit layer. This would in turn imply that the bridge does not implement a protocol lying at a layer higher than the protocols implemented in the interface units.

E. ADDRESSING, SECURITY AND ROUTING

A two-level hierarchical addressing scheme is specified for addressing subscribers in the LAN. Therefore an address of a subscriber will have two parts: the first part identifies a particular subnetwork, and the second part gives the

address of the subscriber on the subnetwork. All routing information will be stored in the bridges for data going across several subnetworks (this is due to the fact that information must be available to direct the data to the desired destination along some optimal path).

Addressing and routing both have implications for data security. The sender TIU inserts the destination address and security level of the data in the header part of the packet. The TIU, whether single-level, variable-level, or multilevel, is trusted to assign the correct security level of the subscriber to the packet. A data packet en route may pass through one or more bridges. Each bridge must decide if the packet should be broadcast on the second subnetwork. At least part of the routing mechanism of the bridges must be trusted since a packet of a given security level must not appear on a subnetwork of a lower security level. The receiving TIU is trusted to pass the data to its receiver only if the receiver's security level is greater than or equal to the security level of the packet.

The hierarchical addressing scheme and the use of a security field in the packet have a definite effect on the protocols of the LAN and the hardware and software that will support these protocols.

C. PROTOCOLS

In order for the LAN of figure 4.2 to perform its communication service, communication protocols implemented in the TIUs and bridges must perform a variety of basic functions. Since security and bridges have been added to the "usual" structure of a LAN, the effect they will have on protocols must be considered.

The protocols have been arbitrarily divided into two groups: low layer and high layer. The protocols in the low layer group perform functions of layer 1, 2 and 3 protocols in the ISO Reference Model [Ref. 14: pp. 81-118]. For example, they provide procedures for transporting packets from a sender to receiver within the LAN. Note that protocol "layers" is used instead of "levels" simply to avoid confusion with security "levels". The layer 1 protocol is the physical layer protocol and specifies characteristics such as voltages, timing, data encoding and decoding, etc., for the transmission medium.

The layer 2 protocol (link protocol) specifies how two physically connected devices (e.g., host-TIU or TIU-TIU) communicate. For our LAN the layer 2 protocol is the TIU-LAN access protocol, implemented in the interface units, that allows interface units (and bridges) to communicate. The design of the layer 2 protocol is based on the proposed IEEE standard 802 CSMA/CD protocol. This version of CSMA/CD provides a broadcast capability and implements collision detection due to simultaneous transmissions, and retransmissions when the medium is not in use. It also retransmits when the receiver does not receive an acknowledgement. Since a packet that is successfully transmitted and delivered may still be discarded on detection of a transmission error, the service provided by this protocol is like a datagram service (e.g., there is no assurance that correct data will be delivered to the receiver(s)). Therefore, a frame check sequence field in the CSMA/CD link layer protocol is used to detect a damaged packet so that the packet is assured correct if received.

The layer 3 protocol is the network layer protocol. It provides a means for data delivery across networks or subnetworks. The layer 3 protocol is usually nonexistent for single-network LANs (or else it is merged with layer 2)

since there is no cognizance of a "network" as separate from the interface units themselves. However, this layer may include a sublayer function to provide transmission of packets through an interconnected system of computer networks. Most of this design involves only the layer 2 CSMA/CD protocol where the intersubnetworking function is subsumed in CSMA/CD and transparent to any additional inter-networking functions in layer 3. The following paragraphs illustrate the more in-depth design details concerning the low layer protocols.

1. Low Layer Protocols

The physical layer 1 protocol of the proposed CSMA/CD IEEE 802 standard can be used in this LAN unchanged. Layer 2 of the CSMA/CD protocol requires some changes in the frame format and procedure parameters to adapt the secure LAN architecture. It is in this layer where the security considerations have the greatest impact on the protocols. A data security level field will have to be added and the source and destination addresses will have to be modified to reflect the two-component nature of an address (local network number, TIU address). In figure 6.1 the IEEE CSMA/CD link layer frame format is shown for comparison with the authors' proposed changes [Ref. 2: p. 34]. The numbers at the left of each format indicate the number of bytes comprising the field. The importance here is not the exact number of bytes in each field but rather the differences between the "standard" protocol and the proposed secure LAN version. The following is an interpretation of the fields of figure 6.1(b):

Destination Subnet Number: 1 byte

function: Destination local subnetwork number

Destination: 6 bytes

function: Address on the subnetwork of the TIU
receiving the frame




Figure 6.1 Local Area Network Packet Format.

Source Subnet Number: 1 byte

function: Source local subnetwork number

Source: 6 bytes

function: Address on the subnetwork of the TIU
sending the frame

Security Level: 2 bytes

function: Security level of the data part in
the frame

Data: Variable (up to some maximum) number of bytes

function: Data in a fully transparent form,
i.e., any bit sequence is allowed

Frame Check Sequence: 4 bytes

function: Contains cyclic redundancy check
(CRC) value computed over all the fields

Note that the length of the basic header of the frame is the same as in the IEEE 802 specification except that the header has been extended by usurping part of the data field for the security level. Also the address fields have been subdivided to incorporate the subnetwork number rather than adding more fields. The particular positions of the new fields were chosen to make maximum use of anticipated off-the-shelf components and well-tested concepts in the design of the secure LAN. The reasoning for the positions of these fields will be discussed further when the TIU design details are presented.

The CSMA/CD link layer protocol allows two types of addressing called physical and multicast addressing. The physical address is the unique address of a station whereas the multicast address is either a multicast-group address or broadcast address. A broadcast address is meant to denote all the stations on the LAN, and a multicast address is associated with a group of stations on the LAN. A number of conventions may be implemented to distinguish the different types of addresses. In the secure LAN, multicast addressing

can be allowed for stations on a single subnetwork by adopting similar conventions in the destination field of the secure LAN link protocol. However, in order to allow multicast (group or broadcast) addressing throughout all the subnetworks in the secure LAN, the bridge design would be somewhat complex. Also, only unclassified sources could address a packet to all destinations in a multilevel LAN. Sidhu and Gasser do not address the "full" multicast extension at this time.

It should be noted that since the security label is part of the CSMA/CD protocol, the part of the TIU responsible for maintaining the integrity of the CSMA/CD header must be trusted. Various aspects of handling CSMA/CD or higher layer protocols (including getting data within the packet transmitted properly) need not be subject to rigorous controls placed on the development of trusted systems. This is the authors' motivation for the separation of functions that will be discussed later.

CSMA/CD alone is not generally suitable for inter-networking, however, for this secure LAN architecture the authors felt that the capabilities of CSMA/CD are adequate for communication among subscribers among a small number of subnetworks. The authors see the use of CSMA/CD alone as one of the means by which the design can be made simple and implementable in the near-term, especially considering the additional requirements that must be imposed on the implementation of trusted systems. However, in the long term as traffic load increases and the number of subnetworks grows, congestion and bottlenecks may require a more powerful protocol appropriate to internetworking.

This security architecture does not depend on any particular network layer protocol, however the authors recommend the use of the DoD Internet Protocol (IP), or an enhanced version of IP as outlined in Skeltons' report

[Ref. 15] which takes into account some characteristics of local area networks. IP provides a primitive form of congestion control. By considering a change to a new protocol now, the appropriate mechanisms can be employed to allow for future enhancement without any impact on the users of the LAN. For this upgrade of the CSMA/CD protocol it would be reasonable to merge CSMA/CD with limited portions of IP responsible for addressing, security, and congestion control, so that the separate subnetworks in the secure LAN would be treated as separated networks by the merged CSMA/CD-IP protocol. This simplifies the security issue somewhat and might possibly minimize congestion problems in the bridges as the load on the LAN increases. More comments on IP are provided in section F.

2. High Layer Protocols

This group consists of protocols at layer 4 and above in the ISO model. They use the services of the lower layer protocols and in turn provide value added services to a protocol layer or user above. For instance, if a reliable end-to-end data transfer service is required, it is provided by a data transfer protocol in this group using a suitable end-to-end acknowledgement scheme. Various other features that can be built into high layer protocols are retransmission on time-out, sequenced delivery of packets, flow and congestion control, etc. In general, the secure LAN design does not depend on or affect these high layer protocols in any way, however comments will be made regarding the use of DoD standards for the host-to-host protocols, the Transmission Control [Ref. 16] and Internet [Ref. 17] Protocols (TCP/IP). It should be noted here that IP does not fit neatly into a specific layer of the ISO reference model because it lies somewhere within or below TCP (which is at layer 4) and above the link layer, therefore making it

a "low layer" protocol. The low layer protocols previously discussed provide the basic transport mechanism for moving data between TIU's in the LAN and through the bridges. The protocols together provide a service that can be used to support a variety of high layer protocols depending on the applications. Note that lower layer protocols do not provide assurance that packets will be delivered or that they will be delivered in the order in which they are transmitted. There is also no automatic end-to-end acknowledgement for successfully delivered packets. If any of these features are desired they must be based on a suitable high layer protocol.

In order to achieve reliable and in-order data delivery at a destination in the LAN, TCP can be implemented on the network layer protocol. IP and TCP are DoD standard protocols for a "catenet" (an interconnected system of packet switched computer communication networks) [Ref. 18: pp. 287-305]. In the catenet environment, IP provides a datagram service from a source to a destination host. It also provides for fragmentation and reassembly of long datagrams for transmission through networks with small packet sizes. TCP is a connection oriented, end-to-end reliable host-to-host protocol for data delivery. It provides for recovery from lost, damaged, duplicated and out-of-order delivered data by underlying less reliable media. The sending TCP assigns a sequence number to each transmitted packet and requires a positive acknowledgement (ACK) from the receiving TCP. If this ACK is not received in a specified time out interval, then the sending TCP assumes that the packet is lost and retransmits it. The sequence numbers are also used for detecting duplicate or out-of-order packets. A checksum routine is used to detect damaged packets (note also that the CSMA/CD protocol already detects damaged packets with its frame check). TCP also has a

"window" mechanism for flow control that regulates the data flow between source and destination.

Sidhu and Gasser mention TCP, not only because it is a DoD standard and therefore likely to be used for wide-area networks, but because they feel it will be suitable for use in a LAN as well. Since users of the LAN will have a need to access TCP-based systems on wide-area networks via a gateway, great difficulties in compatibility can be avoided if the protocols used by subscribers throughout the LAN are also TCP and IP.

Sidhu and Gasser also point out that all higher layer protocols, including TCP, are unaffected by their secure LAN design, however it is important to note some potential problems that may occur with certain implementations. There is an options field in the header of the IP and one of the options may be the security label for the packet. Also, higher layer protocols may include such labelling. But, the authors have designed the secure LAN such that it uses a security label in the low layer (CSMA/CD) protocol instead. Anything above layer 2 is simply data to the CSMA/CD protocol and is ignored by the trusted portions of the TIU. If an untrusted single-level host created that data and is responsible for handling the TCP and IP protocol, then any higher layer security labels cannot be believed by the trusted CSMA/CD protocol handler in the TIU. Therefore, for the single-level and variable-level resource scenarios discussed in chapter V, which allow only single-level hosts, there is no problem in ignoring the IP or higher layer security labels and using only the security labels in the CSMA/CD layer protocol. The security level in the IP header might be used to specify the "real" level of the data contained in a packet labelled by a single-level TIU at a higher level, but then administrative controls would be necessary to actually downgrade that packet.

The multilevel resource scenario requires the ability to support multilevel hosts. In this environment a multilevel host will "choose" the security level of each packet it sends, and this level must be believed by the TIU, at least within the range of allowable values. Usually multilevel hosts support several processes running at different security levels on a single operating system. Therefore, the security level of a packet depends on the level of the process that sent that packet and is inserted into the packet by the trusted operating system of the host. For hosts that have TCP this security level would more than likely be associated with the TCP protocol layer, since the TCP is the layer at which processes are identified.

If a host has TCP, then the fully formed TCP packets are transferred into the TIU with some control information so that the low layer protocol envelope can be created. If TCP were to be implemented in the TIU then the same thing happens but within the TIU itself (the host just transfers "raw" data plus some control information to the TIU so that the proper TCP packets are created by the TIU). In either case, there must be some means by which security information known to the TCP implementation is transmitted to the lower layer protocols for the labels that are the basis for security markings in the LAN.

If the IP and CSMA/CD protocols were to examine security labels in the headers of higher layer protocols, then this would be a violation of the ISO concept of separation of protocol layers (lower layer protocols are not supposed to know about the TCP formats). Instead, the security label would have to be passed along as control information (an additional parameter) from the high layer software to the lowest layer interface. Sidhu and Gassers' concept appears workable, however it is still not very clean as it requires information relevant to the high layer data to become part of the low layer protocol.

For eventual upgrading to combine IP with CSMA/CD, the security label problem is somewhat simplified because the label already in the IP options field could be utilized. But this does not fully eliminate the problem where the security information originates above the TCP layer.

E. TRUSTED INTERFACE UNIT (TIU)

The single-level architecture is emphasized in this chapter. Although the most detailed design is presented for the single resource scenario, enhancements will be discussed for the remaining two scenarios. The major goal in this architecture has been to provide a distinct red/black separation within the TIU and to minimize the complexity and size of the mechanism in the red area that is responsible for maintaining security. In conventional red/black separation devices, there is a crypto unit between the two sides. Usually, neither the red nor the black side is responsible for maintaining security, and if either fails it is unlikely that the crypto unit will pass intelligible data. In our TIU, the red side must be "trusted" to work properly, or at least to prevent accidental disclosure of data, despite the possibility of hardware failure. Note that the concept of "red/black" with respect to the TIU refers to multilevel (trusted) vs. single-level (untrusted) rather than to classified vs. unclassified. The analogy between the two is useful, however, as the single-level portion may be unclassified, but it might also be classified at any level at or below the maximum of the multilevel subnetwork to which the TIU is attached. In any case, the single-level part does not need to be trusted. For many cases, such as TIUs serving Top Secret subscribers on a Top Secret network, significant cost savings may be realized by providing non-TEMPEST versions of TIUs.

Sidhu and Gasser strongly stress their motivation for minimizing the amount of mechanism in the TIU requiring trust. TIUs may be quite complex and research has demonstrated the feasibility of implementing most of the TCP/IP in the LAN interface units [Ref. 15]. They do not expect that it will be possible to adequately verify a large body of software or firmware such as the TCP/IP in the time frame required for the single-level and variable-level scenarios. The verification process is in support of certification (the technical process whereby a procedure, program, system component, or system(s) are shown to be secure; i.e., that the security design specifications are correct and have been properly implemented [Ref. 19: p. C-2]) of the TIU for a multilevel LAN application. Therefore, the less software, firmware, and overall hardware mechanisms that must be trusted, the greater the probability that certification for multilevel operation will be achieved. The feasibility of implementing a single-level TIU is based on the rather trivial increase in functionality required over that of a conventional interface unit (as previously demonstrated by the CSMA/CD protocol modifications discussed earlier). A simpler option would be to build "untrusted" interface units, implementing the full required functionality, rather than going with the trusted TIUs that could be built today. The detailed TIU is considered more of a technical challenge only because of the red/black or trusted/untrusted separation requirement.

1. Single-Level TIU

Figure 6.2 [Ref. 2: p. 43] depicts the architecture of the single-level TIU. The details in the figure have not been worked out to date. The design is not dependent on whether the LAN is broadband or baseband, but it does depend on the use of a CSMA/CD protocol. It is also not dependent

cn whether a "two cable" (separate inbound and outbound cables) or single cable system is used.




Figure 6.2 Trusted Interface Unit (TIU) Architecture.

The dotted line denotes the portions of the LAN and the LAN medium that carry multilevel data. The TIU is termed "single-level" because it allows its subscribers to send or receive data at only one specific security level, the level being determined during manufacture of the red side circuitry or by some maintenance function on the red

side. The red side of the TIU is actually multilevel secure in the sense that it sorts out the single-level data from multilevel data received from the LAN medium, and is subject to the controls necessary for trusted multilevel operation.

Four components reside within the red area; all physically protected to the highest level of the subnetwork:

LAN medium: This is in the red because it carries multilevel data in the clear. It is a passive cable where the only "failure" could be denial of service rather than compromise.

Interface: This portion of the TIU provides the physical interface to the cable. This is a passive system and could fail by denying service to the subscriber at this TIU.

CSMA/CD: This is the hardware that implements most of the CSMA/CD protocol and destination address recognition. For IEEE 802 and similar protocols such as Ethernet, the authors expect that off-the-shelf hardware will be available to provide most of the logic to recognize valid incoming packets destined for this TIU and to handle the contention and retransmission algorithms necessary to service an outbound packet. This hardware may also handle frame check on inbound and outbound packets and notify the security processor when a valid packet has been received or transmitted. Note that the CSMA/CD component must be trusted to leave unmodified the security field of the packet. If it should make an error in address recognition, the only compromise will be need-to-know. If it should make an error in recognizing the frame check (i.e., accepts a packet as valid despite an incorrect frame check), a compromise due to receipt of an

incorrectly marked classified packet is highly unlikely due to the other checks that are made by the CSMA/CD interface and security processor. This design should be fully compatible with off-the-shelf components because the overall packet format and the fields used by possible "standard" CSMA/CD interfaces are left unchanged. The source and destination addresses are treated as single fields to the CSMA/CD interface--the two-part addressing structure is interpreted only by the bridges. The security level field is ignored by the CSMA/CD interface as if it were part of the "data".

Security Processor: The sole purpose of this component is to examine the security level field of incoming packets for legal values and to insert the correct security level into that field on outgoing packets. A received packet that gets past the CSMA/CD component (because the packet has the correct destination address) but has the wrong security level, will be rejected by the security processor. The security processor can be set to accept a certain range of values for incoming levels including only one specific level. However, because the single-level TIU assumes that nothing outside the red area is trusted, secure communications can only take place at a single-level.

Contained in the black area of the TIU are a CPU, memory and I/O ports. That part of the CSMA/CD protocol not responsible for maintaining the header and security label integrity may be implemented in the black area if it can be conveniently separated from the CSMA/CD processor. Because the black area deals only with data of a single security level, it need not be trusted and can be as complex as

desired. Even if the black TIU software contained a "Trojan Horse" attempting to exploit a covert channel, [Ref. 20: pp. 613-615] it would not be able to transmit data on the LAN to a destination of the wrong level, or to receive data from the LAN of the wrong level. Proper implementation considerations should insure that even timing channels could not be exploited by untrusted software in the TIU or interfacing host.

Due to the high data rate possible on a LAN, it should be stressed that covert and timing channels, if exploitable, may provide illicit high bandwidth communications paths that may not normally be important for slower communications media such as packet-switched networks. Unless all the interfacing hosts can be certified not to contain subversive (Trojan Horse) software, adequate security is not provided by a LAN unless these covert channels are closed. The authors feel that the secure LAN architecture presented here does close channels for all practical purposes.

Packets coming into the TIU from the LAN will arrive bit-serially. The off-the-shelf CSMA/CD hardware will be designed to dump data one byte at a time into a microprocessor memory with little buffering in the CSMA/CD hardware itself. Also, the CSMA/CD hardware checks the destination address "on the fly". If the destination is incorrect, the rest of the packet is ignored (not dumped into memory) and the CPU is not notified. If the address is correct, the remainder of the packet gets deposited until the frame check at the end. The CPU is notified of correct receipt only after the frame check is determined to be correct.

The security processor must work in conjunction with the CSMA/CD hardware, but it looks only at the security level field. It will refuse to pass further bytes of incoming data from a packet if the security level in the

header is incorrect, otherwise it transparently passes all data. To find the security level field, it must recognize the start of a frame. On output, it has the option of setting the security level to a particular value or checking that the value inserted by the CPU is correct before transmitting.

If the security processor and CSMA/CD hardware are allowed to dump incoming packets into the microprocessor memory on the fly before the packet is determined to be legal, it is possible that several bytes might be dumped before the packet is recognized as not being for the current recipient of the wrong security level. If this were to occur, then a Trojan Horse in the untrusted CPU could attempt to read partially-accepted packets even if the CPU is not notified of correct receipt of data. However, the possibility of compromise would be limited to information that could be communicated via a covert channel in the header of the packet because the remainder of an unacceptable packet does not appear in memory at all. Given these concerns and the fact that the authors are assuming totally untrusted software in the black side of all TIUs, it is necessary to protect against this threat by buffering, in the read area (probably in the security processor), the first several bytes of the packet until the packet header is determined to be valid. This buffering could be accomplished in the form of a shift register the length of the header up to the security field, so that the first byte does not enter the memory until the header is read.

Note here that even if the header is buffered and loaded into memory only when the security level is valid, there is a possibility that the packet might still be in error, and that error would not be detected until the frame check is read. If this were to happen, a complete erroneous packet would be sitting in memory for the CPU to read.

However, it is highly unlikely that both the destination and security level will check out correctly if the packet was in error and not intended for the designated recipient. If a trusted mechanism inserts the security level into the original packet, it would be almost impossible for untrusted mechanisms to exploit random line noises in an attempt to send a packet to an unauthorized destination. This cannot be considered a useful information channel since it cannot be controlled in any reliable manner.

On a CSMA/CD network there may be many illegal packets received due to collisions. It is therefore possible that a packet from a low security level TIU can collide with a transmission from a high level TIU, yielding a packet containing a mixture of levels. However, if the authors' trusted CSMA/CD hardware works correctly, and if the network is correctly configured according to the distance, spacing, and other electrical requirements of the hardware and medium, all collisions will be detected before the field containing the security level is reached. A basic tenet of CSMA/CD protocols is that collisions can only occur during the transmission of the first several bytes of data. All transmitting TIUs sensing a collision should stop transmitting well before they begin to send the data field of the packet, and the "listen before talk" concept prevents a TIU from transmitting in the middle of another TIU's data field. Therefore, the maximum amount of high security level information that could be mixed in with a low level packet would be the destination and probably the source fields. Utilizing the authors' header buffering scheme in the security processor, any collision would occur well before any of the high level data was loaded into memory. Even if a malfunction prevents the transmitting TIUs from stopping at a collision, it is unlikely that the garbled security level field will be acceptable to the receiving TIU. The

use of suitable encoding of values in the security level can further minimize the chance of error. Also, auditing can detect such malfunctions.

In this initial version a TCP and IP implementation in the TIU will be located in the untrusted portions. If IP were to be further integrated into CSMA/CD, at least a portion of it would have to be contained within the security perimeter.

2. Variable-Level TIU

Since there will be a need in the WIS community for terminals to run at a variety of levels, a minor enhancement to the security processor in the TIU could be performed to accommodate this "variable-level" capability. The authors propose the use of a rotary switch hardwired to the red side of the TIU so that the user can manually select the level at which he intends to operate. It is assumed that the user is cleared to access any security level available on the switch. The only purpose of the switch is to allow the security processor to receive and transmit properly marked data at a level below the maximum of the TIU. Since the black side of the TIU that reads terminal input is not trusted, the security level cannot be entered from the terminal keyboard as "normal" keystrokes. Note that if the user forgets to set the rotary switch to the correct level of the destination with which he is communicating, communication will fail. Consideration should be given to the fact that any change of the switch position (to a lower security level) must result in an automatic reset of the black portion of the TIU and clearing of all its buffers and memory. It can be seen that this variable-level capability would not function well in support of the user who must rapidly switch between levels. This switching option really only supports the ability to logon and communicate with one

single-level host at a time as does the single-level TIU. It does, however, avoid the need for separate TIUs or terminals for the different levels at which a user might want to logon.

Sidhu and Gasser have ignored the issue of how to manage the terminals that are connected to these variable-level TIUs. To enforce security it is necessary that any memory in a terminal be scrubbed in a prescribed manner before lowering the level of a terminal. This may be administratively handled as a manual procedure required at each level change, but the ease of turning a rotary switch on the TIU might dictate more automatic scrubbing of the terminal linked to the switch or controlled by the TIU itself to protect against human error. The overall problem here is how to appropriately deal with the additional flexibility afforded by the variable-level TIUs without increasing the risk of accidental compromise by the user.

3. Multilevel TIU

A true multilevel TIU would allow the subscriber to make packet-by-packet decisions as to what the level should be and would receive packets of a range of levels, marking them appropriately in a trusted manner. If the security level decisions and packet markings were specified by mechanisms outside the TIU (e.g., multilevel host TCP), then the entire TIU must be considered trusted and in the red. This would then make the security processor redundant unless the LAN contains information of levels outside the multilevel host range. Formal certification of a considerable amount of TIU software might be required, depending on the complexity of the protocols implemented in the TIU. For TIUs that implement complex functions it may be possible to construct a TIU that has a form of hardware and software protection so that part of the TIU can be trusted and part

untrusted, similar to what is done for trusted operating systems. Further study is needed to determine whether building and trusting such a protection mechanism is simpler than the effort to build and trust the entire TIU. Building a multilevel TIU is the least of the problems in achieving a general multilevel computing capability because trusted hosts and multilevel terminals must be available first.

E. BRIDGES

Since a bridge is to examine the packet headers created by the TIUs, the bridge has a TIU-type interface on each of the networks to which it is attached and implements TIU protocols. Each bridge also recognizes the destination address of packets passing by on its respective subnetworks.

In general, bridges must pass multilevel data from one multilevel subnetwork to another, therefore there is little motivation for providing simple red/black separation as in the TIUs (unless one of the subnetworks is single-level). The authors envision the bridge to consist entirely of trusted hardware and software. The architecture for a bridge could be constructed that would allow at least some of the internal functions to be untrusted, but their design is simple enough so that an additional mechanism to separate the trusted and untrusted portions is probably unwarranted. This is especially true because it is expected that the bridge will work at the protocol layer in the TIU at which the trusted TIU functions (security processor and CSMA/CD protocol handler) already operate, thereby allowing the exploitation of similar mechanisms.

First the full bridge that implements only the CSMA/CD and physical layer protocols will be discussed. Figure 6.3 [Ref. 2: p. 48] represents the logical structure of the bridge, interposed between two subnetworks. The figure

shows the flow of packets as they arrive from one subnetwork on the left and are sent to the other on the right. This is a "half-duplex" illustration (the entire bridge consists of two identical structures of figure 6.3 for full duplex operation), although in reality some of the hardware (i.e., CSMA/CD interface on each side) might be shared for both directions as in a TIU. Note that the half-duplex bridge is not the same as the half-bridge as illustrated in figure 4.2.

The bridge construction is similar to two TIUs back-to-back with modified address selection and security processor mechanisms. A buffer at least large enough to hold one complete maximum size packet in each direction is necessary because the CSMA/CD protocol requires the ability to retain the packet for retransmission when a collision occurs. To ensure greater reliability, there should be several packet buffers on the transmit side so that temporary congestion on the receiving subnetwork can be smoothed out. Multiple receiving buffers could also be used to take care of temporary bursts of packets arriving faster than the bridge can process them. Note that this buffering capacity will not be able to take care of one subnetworks' consistently being unable to accept data as fast as another subnetwork is sending, since such buffers would quickly fill up. Only temporary overloads can be handled. This buffering capability does allow the use of a slower processor in the bridge that is only capable of handling an average load rather than the peak load without any noticeable degradation of service.

Packets broadcast on the sending subnetwork arrive at the bridge (upper left of figure 6.3), and are selected for acceptance into the input buffers based on the destination field in the header and the routing table within the bridge that specifies which destinations to accept. It is expected




Figure 6.3 Bridge Architecture (Half-Duplex).

that there will be a high volume of traffic local to the subnetwork (not destined for another subnetwork), therefore it seems reasonable to reject packets not requiring routing before buffering them in the bridge. The authors contend that the part of the CSMA/CD interface that checks the destination field will be different from that in a TIU, because a set of destinations, stored in the routing table, is checked. This may require CSMA/CD hardware, rather than off-the-shelf parts, contrary to the authors' premise of utilizing available materials. It should be possible to perform the routing table look "on the fly" without additional buffering, since only a small number of subnetworks are expected at any given LAN site and the subnetwork number

is an explicit part of the destination field. If the input buffers should fill to capacity, then the CSMA/CD interface could be turned off so that all packets are simply rejected, which is similar to what the TIU does when its buffers are full.

The packets that have been buffered on the input side are then processed by the bridge for security restrictions, and the acceptable packets that satisfy all the security checks are placed into the output buffers for transmission without further modification of any of their fields.

Note that because the bridge only implements CSMA/CD, requiring no acknowledgement or reply of any kind, there need not be a logical or physical connection between the two half-duplex portions of the bridge. Therefore, one option to enhance performance might be to limit each bridge to half-duplex operation; two such bridges would be required between each pair of subnetworks.

1. Security Processing

The major part of the bridge processing is in the security processing, not in the routing, which is done before input buffering by selective acceptance of packets. Packets removed from the input buffer are examined by the bridges' security processor for acceptable security level fields based on both the receiving and sending subnetworks' security levels. It enforces both minimum and maximum levels. Packets arriving at the bridge must be marked at least at the minimum level of the sending subnetwork. Packets marked below that minimum would either be upgraded by the bridge or rejected and audited as illegal. If a packet were audited as illegal then it was probably warranted as a marking that is too low as a result of a faulty TIU or configuration problem, or an indication of a hardware penetration attempt. However, auditing would

probably be more efficiently implemented by a special "security-match" TIU that scans all packets rather than by the bridge. Packets above the maximum of the sending subnetwork should also be audited, but in any case they should not be downgraded due to the possibility that a configuration error may cause data of too high a level to be placed on the sending subnetwork. The checks with respect to the minimum and maximum levels of the receiving subnetwork are similar, except that, if the packet is out of range, it is not necessarily due to a fault in a TIU or configuration problem. The sender may simply have mis-addressed the packet to a receiver on the wrong subnetwork. In some cases it may be acceptable for a low level TIU to send a packet to a high level TIU for certain applications. Such a packet may have to be upgraded by the bridge to the minimum level of the receiving subnetwork, but it should be able to be read by the destination TIU without any security problems.

Since it is possible for the bridge to know the security level of the final destination subnetwork and to compare that level with the security level of the packet, there appears to be no reason for added complexity to check any more than the level of the next subnetwork to which the packet is forwarded. This is especially true since the destination TIU or last bridge in the sequence must make the final security check anyway. Therefore, the bridges' knowledge of the subnetwork structure of the LAN need not go further than the two subnetworks to which it is connected.

Since the authors indicate that, for expedience, it probably will be necessary to trust the entire bridge, the only part of the bridge that actually enforces security is the security processor. Therefore, the routing, transmission and buffering mechanisms need only be trusted not to modify any of the packet information. Once the security

processor has accepted a packet, mis-routing (to the wrong TIU) cannot result in a security violation.

2. Routing Concepts

Many possibilities exist for multiple paths from a source to a destination at a LAN site since a subnetwork may have a number of bridges attached to it and a bridge connects two subnetworks to each other. In such a situation, a bridge might be required to decide which path to transmit a packet based on the dynamics of the load on the two subnetworks. This decision-making capability sounds desirable, but the authors have designed their bridges with static routing for simplicity. Therefore, restrictions are needed to define only one logical path between each pair of source and destination addresses (for multiple physical paths) in the LAN. Such a unique path from a source to a destination can be ensured by requiring that only one bridge on a subnetwork can receive a packet destined for any other subnetwork at a site.

The routing table in the bridge decides if a data packet on one subnetwork should be picked by it for broadcast into the other subnetwork. The bridge will read the destination subnetwork number in the header part of the packet and pick the packet for transmission to the other subnetwork if the bridge provides logical connectivity to the destination subnetwork. This means that a bridge must store information about all the destination subnetworks that lie on logical paths through it. An easy way of storing this information in a bridge is in the form of a row vector as depicted in figure 6.4(a) [Ref. 2: p. 51], where 1 in column n means that this bridge will pick packets destined for a remote subnetwork numbered n . Note that this structure would entail programming a different row vector into each bridge. In order to simplify configuration management,

it may be more desirable to make the tables in all bridges identical. To accomplish this, the row vectors could be combined for all bridges into an $M \times N$ matrix as shown in figure 6.4(b) [Ref. 2: p. 51], where M is the total number

		LOCAL NETWORK #				
		1	2	3	4	5
(a)						
		1	1	0	1	0

1 (0) in column n means that bridge will (not) pick packets destined for subnetwork n .

		LOCAL NETWORK #				
		1	2	3	4	5
(b)	B					
	R	1	1	0	0	0
	I	1	0	1	1	1
	D	3	0	1	1	1
	G	4	1	1	1	1
	E	5	1	1	1	1
#						

Figure 6.4 Fixed Routing Tables.

of bridges and N is the number of subnetworks in the LAN. To keep this design simple, the routing information in the matrix is static (does not change with time). The design

does allow for future enhancements to include dynamic updating of the routing information in the bridges, however this would be at the expense of introducing considerable complexity in the trusted bridge software.

3. Buffering

There is a total finite buffer capacity in each bridge for holding packets received from one subnetwork for transmission into the other subnetwork. Both the input and output buffers work in a first-in, first-out (FIFO) fashion, in that packets arriving first are first to be processed or transmitted. It was previously stated that if the bridge's buffers are full, then it will turn off the CSMA/CD interface so that additional packets are ignored until more buffer space becomes available. This situation would probably occur when the receiving subnetwork is overloaded with excessive collisions, or because the bridge is not processing the input load fast enough.

The authors' strategy of ignoring new packets when buffers are full is only one option. Another option is for the bridge to throw out the oldest packet it has received to make room for new ones. This option might be justified on the assumption that by the time the bridge's buffers fill up, the oldest packet in the bridge is likely to be retransmitted by the sender anyway. Throwing out the oldest packet may avoid duplication of packets at the destination or flooding the LAN with duplicate packets. Note that any such retransmission would only be implemented in the high layer protocol (e.g., TCP) that has a timeout option since the authors have assumed that low layer LAN protocols will not retransmit. An extension of this may be to automatically purge any packet in a bridge that has been resident longer than a certain fixed maximum time, a maximum that is keyed to the anticipated higher layer protocol timeouts.

4. Half-Bridges

Whenever two classified subnetworks cannot be brought into close physical proximity, a bridge must be split. In other words each half-duplex bridge would be split. This situation would only occur if there were no encryption technology able to encrypt the LAN medium itself. The CSMA/CD interfaces on the transmit and receive sides of the bridge must be physically close to the LAN media due to timing constraints. Therefore, the design issue here is how to divide the internal functionality of the bridge so that enough hardware can exist on either side, adjacent to the CSMA/CD interface, to communicate with the local subnetwork, while providing a reliable encryptable serial link between the two halves.

If the half-duplex bridge (figure 6.3) were to be split, then it is clear that security considerations dictate that the security processor lie on the left (receiving) side if the sending subnetwork is at a higher level than the receiving subnetwork. If the sending subnetwork is at a lower security level, no security processor is required at all. With two subnetworks of different levels, there should be a security processor to filter packets going in one of the two directions. If minimum levels are to be enforced, security processors would be required in both directions. From these conclusions it is clear that there must be a security processor on the receiving side of each half-bridge, and that the serial encrypted link would lie between the security processor and the transmit buffers. However, a great deal of design work is required in this area.

F. FLOW AND CONGESTION CONTROL

Flow and congestion controls are mechanisms that control the traffic in the network so that network resources are not oversubscribed. Flow control regulates the rate of flow of information between two points in the network. Congestion control is inherently a multipoint mechanism that controls the total amount of traffic entering the network to prevent overload on the aggregate network resources, thereby keeping the network throughput at an acceptable level.

Flow control in the secure LAN is implemented in the higher layer protocols. Congestion issues in the secure LAN differ from that in other LANs only because bridges are the links between the subnetworks. Congestion should be a concern here because of the possibility of overload in a particular subnetwork by its own subscribers and packets arriving from the bridges. If a subnetwork is heavily congested, repeated collisions will occur on attempts by bridges (and TIUs) to transmit packets to the subnetwork. This will probably slow down the rate of packet flow through the bridge and result in backup of the bridges' buffers. In this situation, the bridge would simply ignore incoming packets it cannot buffer, therefore implementing a crude form of congestion control at the link level protocol. This is probably not a completely satisfactory solution to the potential of congestion on a subnetwork for flow problems caused by bridges and it is due to the limitations of the CSMA/CD protocol. Another option to this might be to implement a layer of protocol on or within the link level protocol in the bridges that can be used to quench the source(s) of data feeding traffic into a congested bridge or subnetwork. This option might consist of implementing part of the Internet Protocol (IP) in the bridges, either as part of the CSMA/CD protocol or as a layer above. Congestion

control of the IP involves the transmission of a control packet from the bridge back to the transmitting TIU that stops the flow of incoming packets. Implementation of a new protocol in the bridge that is not security-relevant brings added complexity to the issue. The cost of a bridge that isolates the security-relevant portions would be significantly greater than the simple "pure" CSMA/CD-based bridge which is all trusted. The other alternative, which is fully trusting the bridge containing the IP, may be infeasible due to the complexity of a full IP.

VII. SUMMARY

The sole points of access to media containing data at multiple security levels are the TIUs and bridges. Therefore, they must be "trusted" by LAN subscribers to correctly perform these functions and only those specified for them. Such trust is usually the result of a thorough verification process that examines both normal and fault-ridden processing. A basic trade-off exists between the cost and effort of conducting this process versus the assurance gained from it.

The implementation of the CSMA/CD (within the TIU) access discipline must be trusted never to modify data, nor their associated security label. The implementation of the security processor must be trusted never to modify data nor their security level, to perform its checking function correctly, and to maintain the correct security level for its attached host or terminal.

Three phases of design implementation are envisioned which bear directly on the issue of trust. The first phase calls for single-level, untrusted host and terminal subscribers. The second phase calls for untrusted host and terminal subscribers which may change security levels (following appropriate sanitization procedures) from time to time. Each such change requires re-informing the TIU of the applicable security level. The third phase calls for the trusted host and terminal subscribers capable of simultaneously supporting and governing untrusted processes running at different security levels. This phase calls for the rapid multiplexing of multilevel data exchanges. The other two phases require that the TIU enforce a single security level unless there is a manual intervention. The third

phase also implies that the attached subscribers are trusted to correctly label outgoing data.

A better approach might be to design a Data Base Management System (DBMS) that would handle data labelling. However, this protection would have to be extensive since any variables extracted from classified files would somehow have to carry that same classification level throughout.

The bridges connecting discrete LANs pose a greater requirement for trusted functions. Their implementations of the CSMA/CD discipline and the security processor require the same trust as those of the TIU. The management of the buffer space requires verification that no message intermixture nor other modification can occur. Finally, the fact that a bridge spans two networks, each holding multilevel data, means that the bridge's security processor must function as a multilevel one, with attendant complexity in the verification. Note that the aforementioned consequences were not addressed by the authors in their design.

A. ADVANTAGES

1. This design concept enforces a multilevel security policy over a collection of local networks and their subscribers, and it is intended to prevent security compromises among cleared but untrusted processes. Therefore, an untrusted but highly classified process cannot address and send classified data to a process classified at a lower level just as it cannot downgrade information that it places on the LAN medium. An untrusted process classified at a lower level cannot gain access to data on the media that are classified at a higher level.

2. The consequences of fault conditions that can occur during LAN operations are addressed. The error-checking procedures of the CSMA/CD function make it unlikely that bit

errors can simultaneously cause classified data to be mismarked in the security field and misrouted in the address field. Either of these errors will cause data to be rejected by the addressee's TIU. Also, collision detection makes it extremely unlikely that a transmission collision can result in data intermixing and a resulting compromise.

3. Convenience is the chief advantage of this "trusted system" design. It allows the simultaneous sharing and protection of data in an environment of multiple LANs. Even though it has some advantages of building on existing technology, the design costs could be quite expensive due to the data management engineering that would have to be built into the gates and bridges.

B. DISADVANTAGES

1. The design concept presented can handle multiple security levels but has generally ignored the control of need-to-know. In order to implement this requirement there must be additional mechanisms in each TIU to limit values of the source or destination fields in the packets. Such mechanisms would probably add considerable complexity to the TIU in terms of trusted software. Authorization databases could be accessible to TIUs on the LAN with the TIUs making requests for connection via these databases.

2. Each TIU attached to a terminal cannot verify the identity of the user. Each TIU must believe that anyone with physical access to the TIU therefore has authorization to access anything on the network within the range of security levels at which the TIU is initialized. Many complex authentication mechanisms (passwords, keys, etc.) could be implemented, but, as with the need-to-know, the constraints of this architecture dictate that the ultimate granter of access must be the TIU and not an external mechanism of the LAN.

3. Even though the secure LAN is multilevel in the sense that data of different security levels are kept separated with respect to subscribers in that subnetwork environment, overall security of the data on the subnetwork from threats outside the environment depends on the ability to physically protect the LAN medium, the TIUs and the bridges. Link encryption might provide protection to portions of the medium that cannot be physically protected. Except where the TIUs and bridges connect to the medium, data of all security levels must be in the red. Therefore, if there is a physical security breach and a TIU or red portion of the medium is compromised, all data on the subnetwork is accessible to the penetrator. Only techniques that encrypt all data on the medium can counter such an attack. This possible threat may be compounded by the fact that because of the nature of the broadcast medium, unauthorized receipt of data by a compromised TIU or line tap may not be detectable. A problem related to this is a possible malfunction of the TIU resulting in receipt of data by the subscriber for which he was not authorized. The use of end-to-end encryption would prevent such a compromise.

Unauthorized access to the medium or compromise of a TIU could also result in an active attack where a penetrator injects packets into the network to cause a receipt of classified data or to masquerade as a classified TIU.

4. There is no way to remotely disable a subscriber or control access between subscribers since there is no central authority capable of granting permission for two TIUs to communicate. Authority to communicate is distributed among all the TIUs on the subnetwork.

5. Since the physical subnetwork is assumed to be relatively static, it is not possible to install a new subscriber of an arbitrary security level anywhere along the medium. For example, if in a building containing only a

Secret subnetwork, one wanted to add a Top Secret terminal, it would be necessary to upgrade the entire subnetwork in the building to Top Secret, or to link that Top Secret terminal to a TIU on the nearest Top Secret subnetwork in another building with an encrypted line.

One advantage of general LAN technology is the ability to add subscribers wherever desired without disruption of service. This secure LAN architecture retains that ability for subnetworks within a given security environment.

6. The problem of congestion depends greatly on the physical layout of the LAN and its subscribers. The authors did not address the issue of multiple bridges to dynamically distribute the load on the bridges, or multiple subnetwork connections for a bridge to alter routing around a congested subnetwork. Through the exchange of control information between bridges, either of these enhancements could be easily implemented. The congestion problem should be taken into account because the subnetwork structure is more likely to be configured to accommodate the security requirements than for load distribution.

7. There will probably be a delay in packet delivery time due to the bridges. The amount of delay will depend on the power of the bridges, the load on the various subnetworks, and the number of subnetworks a packet must travel through. The need for a bridge to fully buffer a packet before retransmission alone introduces a considerable delay compared to the single-network LAN delays. The choice of a datagram CSMA/CD service for the LAN protocol makes the low layer protocols immune to delay, but higher layer protocols, such as TCF, which provide acknowledgements and may have timeouts tuned to typical LAN delays, might have problems adjusting their delays depending on the locality of the destination.

8. The authors modified the standard source and destination address fields of CSMA/CD to contain a subnetwork address along with each TIU address. This hierarchical structure of the address is irrelevant to the address-recognition hardware in the TIU, however the structure does restrict the possible destination addresses a TIU may have. In other words, all TIUs on a single subnetwork must contain addresses whose subnetwork number has a specific value, or falls within a specific range, and these ranges must be unique for each subnetwork on the LAN. The hierarchical address is no particular problem if a portion of the address-recognition mechanism in the TIUs can be "programmed" on-site for the particular subnetwork number, but the management of subnetwork numbers requires additional configuration control that is not required in other LANs. Only the subnetwork portion of the address should be configurable on-site in the TIUs, otherwise one would have to "configuration manage" the address of every TIU on the LAN.

The manner in which the addresses are determined could be another disadvantage. Manufacturers "burn-in" destination addresses at the factory such that no two devices will ever have the same address (similar to the embossing of serial numbers on products). This technique minimizes the cost of hardware in each interface unit necessary to program the address on-site, and eliminates the possibility of duplication among the thousands of interface units to be manufactured. However, this technique does prevent users from selecting specific addresses they may want.

Another method for determining addresses would be to add an additional software layer which would provide a selection for addresses for users but could possibly be more costly than the "burn-in" method.

An alternative to the hierarchical address structure is to use the standard address of CSMA/CD, but to have large tables in the bridges to identify the subnetwork of each destination TIU on the LAN. This would require continuous configuration management of these tables if new subscribers are added fairly frequently to the LAN.

9. Throughout this design concept the authors have assumed that suitable encryption devices are available that can handle both the LAN medium itself for those portions of the subnetwork that are classified but must pass through unprotected areas, and the lines between the subscribers and their TIUs for those subscribers remote from their subnetwork. Encryption for the latter is completely straightforward, as the TIU-subscriber lines will involve relatively low speeds and protocols for which encryption is commonly applied today. The authors based their secure LAN on a coaxial cable with broadband signals, implying that the cable would have to terminate at the crypto units, where the signals would be demodulated, converted to digital and encrypted. The encrypted bit stream could be transmitted using any desired communications medium while in the unprotected area, until reaching the other end at which the bit stream is decrypted and remodulated onto another coaxial cable.

Two disadvantages are apparent here. First, there is a noticeable delay involved in encryption and decryption of data and this delay would probably be noticeable to the CSMA/CD protocol. This type of delay is relevant to the CSMA/CD protocol because it is a delay in a TIUs' reading of its own transmission used to detect a collision. The major impact of a small delay is on performance, but a large delay could affect security due to assumptions that were made in the TIU design about the way the CSMA/CD protocol detects collisions.

The second disadvantage occurs when there is information other than LAN data on the medium. An example of this is when LAN data and television signals are used on the same broadband cable. To encrypt just the LAN data without destroying the TV signals (assuming these are unclassified), trusted repeaters would have to be used to capture and rebroadcast specific unclassified TV channels only.

To deal with the problem of interaction between the delay and protocols, a given subnetwork must be entirely protected and encryption would only occur between the half-bridges between the subnetworks. In other words, limit encryption to the bridges only, or better yet, incorporate end-to-end encryption.

10. Another disadvantage is the case of a low level user submitting queries to a database in a high level host. The response from the host would have to be filtered through a guard for downgrading. If the downgrading were reliable then there is no problem in allowing the query itself to go directly from the user to the host. The real problem is that the high layer TCP cannot work in a one-way fashion. A TCP acknowledgement from a high to low security level cannot be permitted in this design concept because security is enforced in the low layer protocols. Establishing a connection, even if one-way, requires two-way communications. A method of dealing with this disadvantage would again be to design a DBMS that would handle data labelling.

11. Another example related to one-way communications is the problem dealing with LANs that use control packets on the network for administrative and maintenance functions. For example, TIUs might send periodic control packets to some central destination in order to monitor the status of all TIUs; or there might be a requirement for a maintenance procedure that requires interrogating all TIUs to see if they are responding. Also, accounting information or statistics gathering may be required.

The authors have not allowed for any special type of control packet for which security restrictions do not apply. These administrative interactions will more than likely be required, and it may be possible to implement these interactions within the operational security constraints. Consideration should be given to building a special-purpose TIU for maintenance purposes that can read packets of all levels while it sends unclassified packets that can be read by any destination.

12. The last major potential disadvantage lies in the authors' choice of the protocols TCP and IP. TCP and IP are not commercial standards. To the extent that there is movement toward a commercial standard, the CCITT X.25 international standard is favored. X.25 is a network interface protocol. It is designed to interface between a host and its local packet switch. Once packets reach a local switch, it is supposed to translate requests for service into another switch-to-switch protocol (e.g., X.75) for transport to a remote switch which will reconvert it into X.25 again. X.25 does not provide for end-to-end reliability. In fact, the standard explicitly specifies several situations where a switch will close a connection when an error is detected. Furthermore, there are no mechanisms for demultiplexing or security. The design is such that a higher level transport layer protocol must provide these functions.

The European Computer Manufacturers Association (ECMA) and the National Bureau of Standards have submitted a series of transport layer protocols with different classes of capability [Ref. 21] as potential international standards. The Class 4 protocol in combination with X.25 provides most of the capabilities of TCP/IP. It is expected that 2-4 years of experience will be needed before these emerging standards will reach a state of maturity. Note that TCP/IP went through four different versions and several large

perturbations before reaching their present state. Therefore, X.25 and the ECMA Class 4 together might be considered viable alternatives to TCP/IP, but X.25 alone cannot. Since the ECMA transport protocols have only been recently made available or subject to extensive testing, then TCP/IP is probably the best product currently available.

There is one other protocol alternative worth mentioning and that is Delta T [Ref. 22]. Delta T is an end-to-end timer-based transport protocol developed at Lawrence Livermore Laboratories. It provides most of the capabilities of TCP/IP class of functions.

It is felt that TCP and IP should be used in the near-term, however, this position should be reevaluated as the Federal and International standards mature.

C. FURTHER RECOMMENDATIONS

In addition to the various alternatives and recommendations that have already been made throughout this thesis, the following alternative approaches should also be examined for use in the secure LAN design.

1. Physically separate LANs can be employed using existing commercially available LAN hardware, and therefore has minimal implementation risk. It differs from the authors' secure LAN architecture in that it does not use trusted interface units to separate traffic in each LAN, but assumes that such traffic is all of the same level. There is a problem here, in that, without trusted interface unit protection, subscribers are left with little choice but to physically protect all computers and terminals to the level of the LAN to which they connect. It also would not allow for easy implementation of the variable-level terminal concept or support for multilevel hosts.

2. The option of using different channels or time-division multiplexing to segregate security levels on a single cable should be examined for near-term use. There is the obvious advantage over the multiple cable approach above, as well as the opportunity to use commercially-available hardware. However, the number of potential channels is still rather limited, the architecture is insensitive to relative traffic load, and trusted software is still required to allow resource-sharing among levels.

3. Some DoD-sponsored research is being done in the area of end-to-end encryption (encrypting a LAN medium), permitting a single cable to pass through all WIS-supported areas of a building regardless of physical protection. The encryption would protect resident data against wiretapping by unauthorized TIUs, and access control would be accomplished through key distribution. This method would require encryption modules within each TIU, and host IUs which have open logical connections to more than one user would be required to have more complicated encryption devices.

D. CONCLUSION

This thesis has presented and examined the results of an initial design-level study to incorporate multilevel security into local area networks for upgrading the ADP support to the WWMCCS Information System. The study focused on objectives that would be achievable in the 1985 time frame and therefore make maximum use of off-the-shelf technology. The design is oriented to minimizing the near-term risks for an initial WIS scenario while laying a foundation for the "maximum" long-term capability for WIS.

The reader should be aware that many of the ideas covered in this design concept are still the subject of basic research, and before they can be put into practice, they need a more rigorous examination.

LIST OF REFERENCES

1. Modernization of The WWMCCS Information System (WIS), prepared for The Committee of Representatives, 31 July 1982, prepared by The Assistant Secretary of Defense (Communications, Command, Control and Intelligence) with the assistance of the WWMCCS System Engineer, Defense Communications Agency.
2. Gasser, M. and Sidhu, D.P., Design For A Multilevel Secure Local Area Network, The MITRE Corporation, Bedford, MA, March 1982.
3. Security Requirements For Automatic Data Processing (ADP) Systems, DoD Directive 5200.28, Change 2, 29 April 1978.
4. WIS Integration Contract LAN System Concepts Paper, Draft, 13 April 1983.
5. Clark, D.D., Pogram, K.T., and Reed, D.F., "An Introduction to Local Area Networks," Proc. IEEE, Vol. 66, No. 11, November 1978.
6. Barldauf, D.L., ACCAT Guard Overview, MTR-3861, The MITRE Corporation, Bedford, MA, November 1979.
7. Stahl, S.H., Hathaway, A., LSI Guard System Specification (Type A), MTR-80W00319, The MITRE Corporation, McLean, VA, December 1980.
8. Local Area Network Security, TM-WD-8513/200/01, Draft, System Development Corporation, McLean, VA, 4 August 1982.
9. The Ethernet: A Local Area Network Specification, Version 1.0, DEC, INTEL, XEROX, September 30, 1980.
10. Metcalfe, R.M., Boggs, D.R., "Ethernet: Distributed Packet Switching for Local Computer Networks", Comm. ACM, Vol. 19, No. 7, July 1976.
11. Local Network Standards Committee, A Status Report, Draft B, IEEE Computer Society, October 1981.
12. Cerf, V.G., Kirstein, P.T., "Issues in Packet Network Interconnection", Proc. IEEE, Vol. 66, No. 11, November 1978.

13. Surshine, C.A., "Interconnection of Computer Networks", Computer Networks, Vol. 1, 1977.
14. International Organization for Standards, "Data Processing--Open Interconnection--Basic Reference Model", ISO/TC97/SC16, Computer Networks, Vol. 5, 1981.
15. Skelton, A.P., Nabelsky, J., and Holmgren, S.F., FY80 Final Report: Cable Bus Application in Command Centers, NTR-80W00319, The NITRE Corporation, McLean, VA, December 1980.
16. Postel, J. (ed.), DoD Standard Transmission Protocol, Defense Advanced Research Projects Agency, 1981.
17. Postel, J. (ed.), DoD Standard Internet Protocol, Defense Advanced Research Projects Agency, 1981.
18. Cerf, V.G., "DARPA Activities in Packet Network Interconnection" in Interlinking of Computer Networks, 1979.
19. AF Security Accreditation Planning Model, DCA 100-79-C-0036, International Business Machines Corporation, Arlington, VA, November 1981.
20. Lampson, E., "A Note on the Confinement Problem", Comm. ACM, Vol. 16, No. 10, October 1973.
21. Burrus, J., Specification of the Transport Protocol, National Bureau of Standards, February 1981.
22. Watson, R., Delta-T Protocol Specification, Lawrence Livermore Laboratory, 01 October 1981.

INITIAL DISTRIBUTION LIST

	No. Copies
1. Defense Technical Information Center Cameron Station Alexandria, Virginia 22314	2
2. Library, Code 0142 Naval Postgraduate School Monterey, California 93943	2
3. LT Debra A. Straub, USN NAVCCMMSTA, Box 22 FPO New York 09571	1
4. Prof. Norman R. Ivons, Code 541b Department of Administrative Sciences Naval Postgraduate School Monterey, California 93943	1