

Local Area Networks

Marshall D. Abrams and Harold J. Podell

Local area network (LAN) communications security is addressed in this essay. LANs are introduced as providing: (1) a private communications facility, (2) services over a relatively limited geographic area, (3) a high data rate for computer communications, and (4) common access to a wide range of devices and services. Security issues pertinent to LANs are discussed. For example, LANs share many security problems and approaches for their solutions with point-to-point conventional communications systems. In addition, LANs have some unique problems of their own: (1) universal data availability, (2) passive and active wiretap threats, (3) end-to-end access control, and (4) security group control.

Countermeasures include physical protection, and separation by physical, logical, and encryption methods. Trusted Network Interface Units, encryption, and key distribution are also discussed.

Examples are discussed to illustrate the different approaches to LAN security. The examples in this essay are a composite of several existing product features, selected to demonstrate the use of encryption for confidentiality, and trusted system technology for a local area network.

Local area network technology/topology overview

This essay addresses LAN security from the viewpoint of open systems interconnection (OSI). That is, we focus on the seven-layer OSI protocols (illustrated in Figure 1); in fact, we concentrate on the lower layers. This focus follows the history of LANs; that is, the OSI communications problems had to be solved before open systems could be addressed.

It is usually not good form to start an essay by discussing what is not covered, but that is necessary in this case. Some people think of LANs in terms of the services they provide to users. This viewpoint is essentially looking at a LAN as a distributed system, with emphasis on the dis-

tributed operating system and the service it provides. This essay does not address this distributed processing within the terminals, workstations, and hosts connected to the LAN. That is another subject for another essay.

Figure 1. Seven-layer ISO protocol model.

Malicious software such as Trojan horses and worms can attack LANs. In fact, the physical distribution of any network increases the difficulty of protection. Malicious software is discussed in Essay 4.

LANs connect computers, terminals, workstations, and other data terminal equipment (DTE). In this essay we will use “DTE” to refer to whatever is connected to the LAN when it is not important what function it serves. The distinction between a personal computer and a workstation is not important for the purposes of this essay.

Let’s start with a functional definition. A LAN is a private communications facility, usually owned by the organization that uses it. The cost of

using the LAN is fixed, independent of level of usage. LANs provide an opportunity for the owning organization to customize its communications capabilities in many ways, such as carrying audio, video, and data traffic; providing multiple simultaneous connections; and providing security services. A LAN generally serves a limited geographic area, such as a single building or a campus, providing a high communications rate or bandwidth and common access to a wide range of devices and services. In general, LANs may be partitioned or zoned. The zones usually correspond to geographic or work units. Bridges or gateways between zones provide connectivity. Zones at physically separate locations can be connected, using wide area networks or private high-bandwidth circuits, to provide LAN services that attempt to be transparent to the physical separation.

Figure 2. Geographic separation and data rate.

A more technical definition can be found in [PADL82], which we paraphrase as follows: A LAN is a communications mechanism using a transmission technology suitable for relatively short distances (typically

a few kilometers) at relatively high bit-per-second rates (typically greater than a few hundred kilobits per second) with relatively low error rates, which exists primarily to support data communication among suitably attached computer systems and terminals (collectively, DTE). The DTE are, at least in principle, heterogeneous; that is, they are not merely multiple instances of the same product. The DTE are assumed to communicate by means of layered protocols.

Note that no assumptions are made about the particular transmission medium or the particular topology in play. LAN media can be twisted-pair wires, CATV or other coaxial-type cables, optical fibers, wireless, or whatever. LAN topologies can be “bus,” “ring,” or “star.” For our purposes, the significant properties of a LAN are the high bit transmission capacity and the good error properties.

In Figure 2 we identify three network groups from a communications viewpoint. The exact numbers are not important and have changed as various technical breakthroughs have occurred. What are important are the concepts that physical characteristics such as data transmission and error rates can be related to distance and that protocols can be optimized for an assumed operational environment. Convenience, however, may dictate use of certain protocols over a wider distance range than is optimum. From the OSI perspective, most of the protocols used in a LAN should be the same as ones used in a WAN context.

Security-relevant LAN characteristics

All data traffic is available to every node in the LAN zone. There is no routing or switching in the conventional sense; rather there is selection. There is routing and filtering among zones, provided by devices such as bridges and gateways.

An adapter is required between the DTE and the LAN. This adapter goes by many names; in this essay, we shall call it the network interface unit (NIU). The NIU provides physical and logical conversions. For example, the voltages and ways of representing digital signals on the LAN are probably different from those used internal to the DTE. The first NIUs were external to the DTE. The NIU was connected to the DTE using the same standard that is used to connect to an external modem. External NIUs continue to be available. Efficiencies of space and power are achieved if the NIU is moved internal to the DTE. The NIU may be an internal board, whose interface is the backplane of the DTE, or it may be a chip on the main circuit board of the DTE.

Every NIU has an address. When messages are inserted on the network, the address of the destination NIU is part of the message header. As messages flow through an NIU, the destination address is examined. According to the protocol, if and only if the destination address matches the NIU doing the examining, the message is transmitted to the at-

tached DTE. By this very simple filtering mechanism, NIUs provide for pairwise communication between any two DTE attached to the network. It is also easy to provide broadcast communication to all NIUs by using a special address such as the binary value of all ones.

There are many reasons why LANs have become popular, the most salient being flexibility and cost. LAN flexibility derives from their inherent distributed control. That is, all of the active decision making takes place in individual NIUs. New NIUs may be added to the net or activated, or NIUs may be removed or deactivated without making a significant change to the overall intelligence controlling the network. This dynamic flexibility is quite valuable in environments where new DTE may be added to the network at any time or where DTE may be removed accidentally or purposefully without notification and coordination with a central authority. However, certain certification and accreditation processes for secure LANs may require notification to, and approval by, a central authority for all configuration changes.

A DTE may be either a terminal or a computer. Security vulnerabilities have been introduced by marketplace technology advances that have just about caused terminals to disappear. When computers emulate terminal protocols, it may not be possible to tell when the security assumptions about “dumb” terminals are valid. PC emulating terminals are qualitatively different, especially with regard to security. For example, a PC can record all the communications traffic. If the NIU is under program control, which might very well be the case for an internal NIU, address filtering can be turned off. The NIU can operate in “promiscuous” or “snooper” mode, passing all traffic to the PC, which in turn can record it for some future use.

The LAN provides universal access between and among devices. In particular, the LAN may be compared with the point-to-point wiring between terminals and computers that was prevalent in previous computer communications architectures. The lack of flexibility, the cost of installing point-to-point wiring, and the saturation of the physical space available for such wiring have all led to the replacement of this technology.

It is undoubtedly obvious that many of the operational advantages of LANs are also potential security liabilities. These liabilities will be discussed in some detail below.

LAN security problems

LANs share many security problems and approaches for their solutions with point-to-point conventional communications systems. In addition, they have some unique problems of their own. This section surveys these problems and leads into the section that discusses selected approaches for solution.

Universal data availability. The ready access to data anywhere along the LAN is one of its greatest security problems. Data is made available to any party whether it should have the data or not. LANs make all traffic available at or near every NIU. Covert activity is very difficult to detect. Every NIU has direct immediate access to all of the data on the network zone to which it is connected. A normal NIU is expected to ignore all data that is not addressed to itself. However, in some LANs malfunctioning or maliciously designed NIUs can be as acquisitive as they wish.

Passive and active wiretap threats. The interception of information transmission by an adversary has been traditionally referred to as wiretapping. We use the term “wiretap” when discussing LAN security even though it may not be strictly descriptive. Some network media are inherently more tap resistant, or at least may give indication when they are tapped. The media may also radiate information. Electromagnetic media, such as twisted-pair or coaxial cable, are notorious for this weakness. Optical media, such as fiber optics, are orders of magnitude better, but they still do radiate under certain conditions; this radiation can be detected if the adversary gets close enough.

Wiretapping is conventionally subdivided into passive and active categories. In passive wiretapping, the message traffic is observed but not modified. The most obvious objective of passive wiretapping is to learn the contents of messages; traffic analysis may provide the adversary with information when message content is not available. Traffic analysis could include steady-state and transient analysis of quantities of messages between parties and the lengths of these messages. A sudden change in traffic volume between national central banks, for example, might signal a change in the rate of exchange or some other financial activity that could be turned into a profit by someone.

In active wiretapping, there are a number of different ways in which the adversary can modify the communications stream. The generic name for this threat is **message-stream modification (MSM)**. Messages can be completely deleted, they can be inserted, or their contents can be modified. Delay, reordering, duplication, and retransmission are also possible. Deliberate denial of service by temporary or permanent incapacitation of the LAN is yet another form of active wiretapping. Denial of service can also occur due to a variety of natural and accidental causes.

End-to-end access control

The term “end-to-end” is often used to (attempt to) describe the scope of control applied to a communications circuit. Unfortunately, end-to-end is one of the more overworked and less precise terms that one encounters in data communications. The problem is that one observer’s

end is another observer's midpoint. To communicate more precisely, we must identify exactly the endpoints being discussed. The communications engineer traditionally thinks of the endpoint as being the data circuit-terminating equipment (DCE) or at least the interface between the DCE and the DTE. In the LAN context, this means that the data communication end-to-end is from NIU to NIU. Figure 3 illustrates the scopes of NIU-to-NIU and DTE-to-DTE controls; these are also discussed below.

Figure 3. Different end-to-end scopes.

NIU-to-NIU access control. Access control encompasses a subset of correct operation of the LAN protocol. Access control, like flow control and error control, can occur at multiple levels in the OSI architecture. In fact, one of the criticisms of OSI implementations is the overhead of performing the same or similar functions more than once. But with LANs covering only the two or three lower layers of the OSI model, we must accept this duplication. Furthermore, access control can be useful within a LAN or set of compatible LANs, but the granularity of the access control is constrained by the nature of the identification provided by low layer addresses. Providing access control at high layers permits use of finer grained, globally meaningful identifiers, for example, directory distinguished names.

For the LAN to function, it must provide addressing and delivery services. These services require that data be delivered undamaged to its destination. For security considerations, misdelivery and nondelivery are problems associated with damage to the data communications, or at least to the address part thereof. The NIUs play an important part in the delivery mechanism. They must correctly place a destination address on each transmission and must likewise correctly identify the addresses of

those messages, and only those messages, which they are to pass through to the attached DTE.

DTE-to-DTE access control. The first extension to end-to-end is to identify the person or process embodied in the DTE. This is the personnel identification problem.

External NIUs typically support two or more attached DTEs. For access control it would be simplest to restrict that number to one. Otherwise, we will have to increase our trust in the NIU, as discussed below. If the definition of end-to-end has been extended to include person or process, then a mechanism must be built to enforce this decision and the NIU must include protocols at higher layers, since people and processes are not identifiable entities at the lower layers. Logically, this mechanism may be installed either in the DTE or in the NIU. In practice, since the local area network is a more recent development and is procured separately, it is more reasonable to expect the access control mechanisms to reside in the NIU.

Security group control

In many communication systems, the users are subdivided into multiple security groups and levels. National defense classifications have coarse level granularities of confidential, secret, and top secret, and finer granularities based on a need-to-know. Groupings supporting unclassified protection might be based on membership in a particular organization, for example, work groups such as engineering, medical, and sales. It would be entirely reasonable for individuals to belong to more than one security group. Using NIUs to enforce rule-based access control is not based on user authorization, but on device authorization. The NIUs operate at protocol layers well below layer seven, where the individual user exists.

We can increase the workload being supported by the NIU and the convenience of its users if we make the NIU responsible for the enforcement of these security group rules. The most obvious function we would like the NIU to perform is to restrict communication to those people and processes that belong to the same security group.

Security approaches

LAN security can be provided by physical protection, separation, or both. The approaches are high-level design alternatives, not specific mechanisms. Physical protection and logical separation on the LAN are discussed. This section addresses alternative schemes for providing security groups. Logical protection of messages between pairs of communicants will be discussed under encryption in a subsequent section.

Physical protection. Physical protection is the most obvious form of security. It is applicable to almost any valuable resource. A local area network is only one example. A local area network can use attack-resistant enclosures or penetration detection and alarms.

Separation. There are a number of different schemes for separating the security groups using a LAN. The most obvious and straightforward is to provide physically separate LANs for each security group. Simply providing multiple cables may not be sufficient; some mechanism may be required to make sure that an unauthorized NIU does not get plugged into the wrong LAN.

On a single LAN, there are a number of ways of providing separate channels for each security group. One very common medium for local area networks is coaxial cable, which is modulated in frequency channels. These channels are very similar to and in many cases identical with the channels used for commercial television broadcast distribution. In addition to the broadcast channels, there are a large number of additional channels used by the cable television industry (CATV). Assigning separate channels to each security group is an obvious separation mechanism; in technical terms this is frequency division multiplexing (FDM).

Channel separation has been approved by some network security managers and rejected by others. This discrepancy is based, no doubt, on different threat scenarios. Before proceeding with channel separation, you should talk with the manager in charge.

An alternative is to provide separate logical channels by including a channel identification in the header of each data communication packet. The NIU would be required to enforce the channel separation in the same way that it enforces address filtering by recognizing the logical channel and passing messages from it only to authorized DTE. This is one form of time division multiplexing (TDM).

Another way to achieve logically separate channels is to use encryption. Cryptography is discussed in Essay 15; the application of cryptography to LANs is covered in a later section.

Implications and side effects

Security comes at a price. Part of this price is fiscal. There are additional components that must be implemented and paid for. Another part of the price is decreased convenience of usage. Introduction of security measures makes a LAN less convenient to use. Some of the burden may be borne by technology (implemented in the NIUs), but part of the burden must be borne by the parties attempting to communicate.

What happens if parties belonging to different security groups need to communicate? The devices that provide interconnection between LAN

segments, between separate LANs, or between a LAN and a long distance network are variously known by a variety of terms, depending on the functions being performed. For example, repeaters are primarily associated with layer one functions; bridges, primarily with layer two functions; and routers and gateways, primarily with layer three functions. Therefore, this sequence of names generally reflects the amount of work that must be done in providing this interconnection. Outside of the security arena, there are problems of different protocols and address spaces, which have to be solved.

To a large extent, the solution to such problems is known; there are a large number of commercial devices available to provide such services. When security is an issue, we might add the name filter to the collection implying the function of enforcing security rules and allowing only some messages through. Filters may not be required where encryption is used for confidentiality. Confidential messages protected by encryption may be handled as nonsensitive information in a network. Alternatively, filters may be used in conjunction with secure protocols, such as the draft IEEE Standard for Interoperable Local Area Network Security (SILS), Part B — Secure Data Exchange (SDE) [IEEE90].

SDE is a Link Layer entity that provides services to permit the secure exchange of data. Within the Link Layer, SDE is part of the Logical Link Control Sublayer. SDE provides a connectionless service, which is located on top of the Medium Access Control Sublayer that is defined in IEEE 802 LANs and Metropolitan Area Networks.

Filters contribute to protection against denial of service by passing only properly addressed messages. Noise, “network storms,” and messages addressed outside the range set for the filter are not passed through. As shown in Figure 4, the LAN segment behind the filter is protected from malfunction or attack, and has traffic reduced as well.

Figure 4. Filter protection.

Trusted NIU

In any local area network, the NIU must be relied on to perform properly. However, when security is a concern, the NIU must be trusted.

For an unevaluated LAN NIU, the trust is based on commercial practices. This should not make us exceptionally confident, for our experience probably indicates that many commercial products are released with the implicit decision that the customer will detect some latent errors. In the security environment, this approach is clearly unsatisfactory. The problem has been addressed in the Trusted Network Interpretation (TNI) [NCSC87a]. The TNI joined network technology with trusted system technology. Probably the major contribution was understanding how the protocol layer model could incorporate the concepts from the Trusted Computer System Evaluation Criteria (TCSEC) [DOD85].

NIUs evaluated under the TNI are trusted because their hardware and software have been demonstrated to properly implement a set of security rules.

We trust these NIUs to correctly implement the LAN protocols, in particular the affixing of destination addresses and the filtering of received messages according to these addresses. In the TCSEC B evaluation class, the NIU has the responsibility of associating security labels with messages, knowing the security level of the attached DTE, and enforcing the security rules for transferring messages.

NIU protocol support

The NIUs provide protocol support at the Link Layer (layer 2). Protocols supported could include IEEE Standard 802.3 Carrier Sense Multiple Access with Collision Detect (CSMA/CD) (Ethernet), 802.4 token bus, or 802.5 token ring. By definition, NIUs operate only at the lowest two layers. At higher layers the equipment to which the NIUs are interfaced could support one or more protocol suites.

MIL-STD-1777 [DOD83] defines one protocol suite known as TCP/IP (Transmission Control Protocol/Internet Protocol), or simply IP. Although originally developed for wide area network (WAN) applications, TCP/IP has achieved wide acceptance in LANs. The IP suite includes the following protocols that the NIUs may support: virtual terminal TELNET, DoD Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Protocol (IP), File Transfer Protocol (FTP), and Host to Front End Protocol (HFE).

The Internet community specifications for IP (RFC-791) and TCP (RFC-793) and the DoD MIL-STD specifications are intended to describe exactly the same protocols. The RFCs and the MIL-STDs for IP and TCP differ in style and level of detail.

Multilevel LAN using COMPUSEC technology

In this section we will address how a LAN could provide security services using COMPUSEC trust technology. LAN capabilities include:

1. communication among terminals and workstations,
2. terminal-to-host (end system) communications,
3. reliable host-to-host (end system-to-end system) communication,
4. host-to-host (end-system-to-end-system) datagram (connectionless) service, and
5. comprehensive network management for both performance and security.

Figure 5. MLS LAN system diagram.

Access to network media is controlled by NIUs, as shown in Figure 5. Devices controlled include terminals, hosts (end systems), workstations, gateways, and various servers. The network management workstation provides centralized management of the network. DTE can operate within a range of security levels; end systems can support multiple concurrent sessions at different sensitivity levels.

The LAN NTCB (network trusted computing base boundary shown in Figure 5 may be hard to identify in a physical implementation where the NIUs are implemented as integrated circuit cards that plug into expansion slots in other equipment, such as workstations. Essentially all ma-

for workstation vendors supply higher layer protocols as part of their product offerings.

The NIUs provide protocol processing and controlled access to the network for attached user devices at the lower layers, thereby providing complete isolation of user data streams. The NIU trusted software implements the local partition of the NTCB, which ensures that untrusted software supporting different user sessions is logically separated. To minimize the size of the NTCB, it should be designed so that untrusted software implements protocol services, including TELNET, most of TCP, most of host (end system)-to-SNS (Secure Network Server) protocol, and many data integrity features defined in TNI.

Protocol additions. To satisfy TNI evaluation requirements, it is presently necessary to develop protocols for end-to-end user identification and trusted path, since the DoD protocol suite does not support trusted path.

The NTCB partition in the NIU provides separation of all communications objects, complete logical isolation of user sessions, and all necessary support functions. Sensitivity labels may be applied to datagrams, connections, and sessions.

Applying TNI (Trusted Network Interpretation). The TNI allows a network, such as a LAN, to be evaluated as network components providing one or more of these functions:

- M: Mandatory access control (MAC)
- I: Identification and authentication (I&A)
- D: Discretionary access control (DAC)
- A: Audit

LAN subjects and objects. In developing a network security policy, it is necessary to identify the subjects and objects. The TNI emphasizes that it uses the same definition as the TCSEC, but it is still necessary to discuss the definitions. The part of the definition of subject that refers to a process-domain pair is unchanged, but the definition of “user” needs to be interpreted in the network context. There is no human user in any of the OSI protocol layers, just processes acting on behalf of the human user. In OSI terminology, this process is known as a protocol layer entity (for each protocol layer). It is the protocol layer entity that acts as the active agent, causing information to flow.

Discretionary access control in the LAN. DAC maps to the need for correct addressing and delivery. The fundamental address filtering function of the NIU implements DAC. The packet source and destination in-

formation constitutes the security attributes that support DAC. This protocol information is written and read by end system software in layer four and above. In the DoD protocol suite, source and destination addresses are part of IP. An NIU could enforce DAC based on these addresses without embodying any part of TCP. Both the address filtering functions in the NIU and IP have to be part of the NTCB.

If the DTE attached to the NIU is not capable of providing addressing information, this function must be provided by the NIU. DTE that require additional connection setup services from the LAN include terminals and terminal emulation programs in desktop computers, as well as computers that do not support the LAN protocols. NIUs typically provide dialogues or menus for loading addresses from terminal users and end system operators. Non-LAN network protocols or special host-to-front-end protocols may also be supported. In any case, the amount of trusted code in the NIU increases when the address specification function is added.

An end system address corresponds to two different DAC subject identifiers in the TCSEC operating system view. At layer three the address represents a single subject, the network service access point. The end system address also can be projected to include all the human users of the end system; in this view, the end system address is a group identifier for all of those users.

Identification and authentication. Every trusted system requires I&A. In the LAN, I&A applies to human users at terminals as well as trusted hosts. When a trusted host is connected to an NIU, it may take on the I&A responsibility.

Terminals or workstations are often the first point of contact between a human user and the automated information system. It is therefore necessary to provide trusted I&A supported by the NIU. The NIU does not provide the I&A service; it only provides the interface to that service. The following scenario is typical. The NIU is configured to make a connection to the authentication server, which is shown in Figure 5, whenever it senses a new user. Well-defined hardware and software signals from the terminal signal the presence of a new user. Security depends on the inability of the user to directly change addressing information. The authentication server provides the I&A security function for the LAN.

As a matter of improving the human interface to the automated information system, it is increasingly common to provide a single I&A function that propagates the authenticated identification to the end systems providing user services. This capability is known as unary login. Appropriate trusted protocols are required to transfer the authenticated identity to the end systems. Such a protocol may also instruct the NIU to change the destination address to the end system; this makes the network distribution of the I&A server and end system transparent to the user. Fur-

ther discussion of authentication in distributed systems is available elsewhere [WOOD92].

I&A of end system hardware and software require physical protection and administrative controls, such as configuration management. For example, disconnection of end systems must be detected by the NTCB so that network operations personnel have assurance that they will be notified of configuration changes. Unauthorized configuration changes can signal an attack.

Trusted path. The TNI requires a trusted path from users to the NTCB at the TCSEC B2 level and higher. The trusted path between the user and the first NTCB partition encountered must be supported by the NIU for users at terminals that do not include such an NTCB partition. A trusted protocol is required to extend the trusted path to other NTCB partitions.

Audit. The TCSEC requirement to audit “introduction of objects into a user’s address space” can be interpreted to require audit of individual packet delivery; this is clearly a high-overhead activity. A low-overhead alternative is to view the loading of a destination address as introducing a connection into the address space, thus drastically reducing the number of auditable events. Of course, exception events, such as all access control failures, have to be audited.

Figure 5 indicates the presence of an audit server. This server must operate in cooperation with the NIUs. The NIUs acquire the information to be placed in the audit trail and transmit it to the audit server using appropriate trusted protocols. These protocols must provide confidentiality and guaranteed delivery to satisfy the audit requirements.

Integrity. There are integrity requirements in both Part I and Part II of the TNI. The requirements for system integrity in Part I are focused on correct operation of the NTCB.

The LAN can satisfy this requirement by periodically validating the correct operation of the hardware and firmware elements of each NIU’s NTCB partition. One time when such validation is required is when an NIU is added to the network. NIU protocols, implemented within the NTCB, must be designed to provide correct operation in the case of failures of network communications or individual components.

NTCB protocols should be robust enough so that they permit the system to operate correctly in the case of localized failure. The purpose of this protection is to preserve the integrity of the NTCB itself. It is not unusual for one or more NIUs to be inoperative at any time, so it is important to minimize the effects of such failures on the rest of the network.

Some integrity service features can reside outside the NTCB. Otherwise, all software in a network would be in the NTCB.

TNI Part II security services

The security services identified in Part II of the TNI are specified as being optional. LAN vendors decide whether to provide these services as part of product definition. LAN user management decides whether specific services are required in their environment to counter identified threats.

The TNI provides criteria for evaluating the functionality, strength of mechanism, and assurance of these security services. We discuss only functionality in this essay.

Authentication. The LAN should ensure that a data exchange is established with the addressed peer entity (and not with an entity attempting a masquerade or a replay of a previous establishment). The LAN should ensure that the data source is the one claimed.

Attempts to create a session under a false identity constitute a typical threat for which peer entity authentication is an appropriate countermeasure.

Communications field integrity. Communications field integrity refers to protection of any of the fields involved in communications from unauthorized modification. Integrity service counters active threats and protects data against unauthorized alteration. The LAN should ensure that information is accurately transmitted from source to destination. The LAN should be able to counter both equipment failure and actions by persons and processes not authorized to alter the data. The LAN should also have an automated capability of testing for, detecting, and reporting errors that exceed a given threshold.

Nonrepudiation. Nonrepudiation (NR) service provides unforgeable proof of shipment and/or receipt of data. This service prevents the sender from disavowing a legitimate message or the recipient from denying receipt. Since ISO 7498-2 makes it clear that NR is offered only at the Application Layer, it is out of scope for a LAN security device to provide this service.

Continuity of operations. The security features providing resistance against denial-of-service attacks may include the following:

1. Use of redundancy throughout the LAN components can enhance availability.
2. Reconfiguration can provide NIU software maintenance and program downloading to NIUs for software distribution. In addition, to provide initialization and reconfiguration after removing or replacing failed or faulty NIUs, reconfiguration can assist in isolating

- and/or confining LAN failures, accommodating the addition and deletion of LAN components, and circumventing a detected fault.
3. Distribution and flexibility of LAN control functions by replication of LAN management, both ordinary and security related, can reduce or eliminate the possibility of disabling the LAN. Flexible control capability able to respond promptly to emergency needs, such as increase in traffic or quick restoration, can improve the capability to respond promptly to the changes in LAN topology and LAN throughput. Therefore, flexible control may enhance survivability and continuity of operation.
 4. Fault tolerance mechanisms provide a capability to deal with LAN failures and to maintain continuity of operations. Such mechanisms may be single point, such as filtering gateways to partition the LAN to exclude erroneous traffic, or distributed, such as a maintenance channel among the NIUs and the LAN management center.
 5. Measurement of noise and real-time statistical analysis of retransmission rates are typical LAN management services with availability implications.

LAN management. LAN management and maintenance deal with LAN viability, detecting failures and overt acts that result in denial of or reduced service. Simple throughput may not necessarily be a good measure of proper performance. Loading above capacity, flooding, replays, and protocol retry due to noise in the physical layer can reduce service below an acceptable level and/or cause selective outages. Management protocols, such as those which configure the LAN or monitor its performance, are current work items.

An availability problem may cause disruption of more than one peer entity association. The determination of a problem is an application management function, and the corrective action is a system management function.

Mechanisms for detecting denial of service are often protocol based and may involve testing or probing. A common technique is to sequentially poll all NIU addresses. Lack of response from known NIUs indicates an availability problem. Response from NIUs not configured as part of the LAN indicates another class of problem.

In addition to the LAN performance measurements mentioned above, a process may exist to measure the transmission rate between NIUs under conditions of input queuing. The measured transmission rate shall be compared with a predetermined minimum to detect an availability problem.

A request-response protocol such as “are-you-there” or “ping” message exchange may be used to detect availability problems when the connec-

tion is quiescent. The availability of entities at different protocol layers can be used to test remote NIUs, gateways, and end systems. Request-response protocols have been known to crash LANs when coupled with hardware failures and/or abnormal loading. Incompatibilities also sometimes show up when dissimilar LANs are interconnected. Any polling sequence should probably be metered to prevent creating the very condition it is designed to detect.

Note that denial of service is addressed only by detection. Methods for prevention of denial-of-service attacks have yet to be developed.

Confidentiality protection. Data confidentiality service protects data against unauthorized disclosure. Data confidentiality is mainly compromised by passive observation and cryptanalysis. Confidentiality protection is a collective term for a number of security services. These services, described below, are all concerned with the secrecy, or nondisclosure, of information transfer between peer entities through the LAN. COMPUSEC mechanisms do not provide protection against wiretapping or other physical attacks. Cryptographic protection, a COMSEC mechanism discussed in a later section, is often used to provide this protection.

The LAN must provide protection of data from unauthorized disclosure. Physical security, such as protected wireways, can also provide transmission security. The LAN manager must decide on the balance among physical, administrative, and technical security. This essay addresses only technical security.

Traffic flow confidentiality. Traffic analysis is a compromise in which analysis of message length, frequency, and protocol components (such as addresses) results in information disclosure through inference.

Traffic flow confidentiality is concerned with masking the frequency, length, and origin-destination patterns of communications between protocol entities. Encryption can effectively and efficiently restrict disclosure above the Transport Layer; that is, it can conceal the process and application but not the host computer node unless a red/black gateway is used in conjunction with a secure protocol that encrypts the end system address [DINK90].

Selective routing. Since LAN technology substitutes selection for routing, you may wonder why this security service is included in an essay on LAN security. Selective routing applies to gateways and bridges among LANs. Selective routing is the application of rules to choose or avoid specific LANs, gateways, or links between LANs. The selection may be based on static policy or on dynamic attack or error conditions.

Multilevel LAN using communications security technology

Communications security, COMSEC, provides confidentiality services by applying an encryption mechanism. Encryption is discussed in detail in Essay 15. We limit our discussion in this essay to several issues pertaining to LANs.

Channel separation. Encryption can be used in a rather straightforward way to establish logical channels. Each logical channel is assigned a symmetric encryption key or an asymmetric pair of keys. When an NIU is going to operate on a given logical channel, it need only load the appropriate key; thus it establishes a logical channel between NIUs.

If an NIU is to be trusted to operate on multiple logical channels, then it must apply a key for each of those channels to every message that passes through it. For certain symmetric key applications, this seems to be an unnecessarily complex procedure; it would not be unreasonable to restrict an NIU to operating on a single logical channel. If the attached person or process had access to multiple logical channels, then the appropriate encryption key could be changed appropriately under the control of that person or process.

Encryption can also be used to create separate channels. Assuming for the moment that the security group encompasses the NIU and that an NIU belongs to one and only one security group, then encryption between NIUs is analogous to link encryption in conventional communications systems.

There are two ways to achieve separate channels using encryption. Only those NIUs sharing the same symmetric encryption key or matched pairs of asymmetric keys would be able to communicate, thereby creating separate logical channels for their security groups.

A message may be encrypted more than once, if the encryption algorithms permit. This serial encryption is sometimes called superencryption. We shall assume multiple encryptions, thereby making it possible to form subgroups as many times as necessary and to encrypt specific messages between two parties, both of whom belong to the same security group.

Privacy between NIUs. Encryption between DTEs may require another pair of asymmetric keys or another symmetric key. Messages between DTEs would then be doubly encrypted. Without this second DTE-to-DTE key, any DTE attached to an NIU possessing the key for the logical channel could eavesdrop on all communication on that logical channel.

Key management. Key management is a necessary managerial function. The technology is the same for LANs and WANs. However, the logistics of a truly local area LAN may influence the selection of symmetric or asymmetric keying and the manual or automated methods used for key distribution. See the key management sections in the essay on cryptography (Essay 15) for further information.

If multilevel traffic messages of different sensitivity are required, different keying is required for each sensitivity level. This is the only way to achieve the cryptographic equivalent of labeling.

End-to-end symmetric key encryption example. Sometimes end-to-end symmetric key encryption may be used in a LAN to provide logical channel separation. Consider a LAN used by hosts operating at various security levels that send and receive encrypted messages. The NIUs connected to these hosts obtain symmetric encryption keys at the level of the appropriate information to be protected from a central key distribution center (KDC) supporting the various security levels of the network. The KDC is attached to the LAN in the same way as a host. A symmetric key is sent from the KDC to an NIU upon request, using an appropriately defined protocol that authenticates both the requester and the new symmetric key.

The purpose of key distribution is really to support a trusted local service within the LAN: the ability to transform classified or sensitive messages from the host into unclassified encrypted messages suitable for transmission over the LAN.

For symmetric key management, part of the trusted network service is implemented within the KDC, which must generate new symmetric keys for the level of information being communicated, and must also decide, on the basis of an access control policy, which NIUs may share keys. The granularity of key distribution is a trade-off between convenience and protection. Fine granularity would use a unique key for each sensitivity level for each session; coarse granularity would use the same key for all sessions during a time period.

Please see Essay 17 on Privacy Enhanced Mail (PEM) for an example of asymmetric key distribution and end-to-end encryption techniques.

Summary

This essay has reviewed local area networks with special attention paid to security issues. While the LAN offers many advantages in terms of data access and flexibility, the other side of the coin is increased vulnerability to wiretapping. The problems of end-to-end access control and security groups were also discussed. The mechanisms for providing security are physical protection and separation of security groups by physical, logical, or encryption methods — all of which may adversely

affect interoperability. Most protection schemes require that the network interface units be trusted; the reasons and trusted functions were discussed. Encryption is used for creating logical channels as well as protecting sessions between two persons or processes; the distribution of symmetric encryption keys by a key distribution center was outlined.